



**MIPS32® interAptiv™ Multiprocessing  
System Software User's Manual**

**Document Number: MD00904**

**Revision 02.01**

**June 15, 2016**

Public. This publication contains proprietary information which is subject to change without notice and is supplied 'as is', without any warranty of any kind.

# Table of Contents

<b>Chapter 1: Architecture Overview .....</b>	<b>27</b>
1.1: Features .....	28
1.1.1: System Level Features .....	28
1.1.2: Core Level Features.....	29
1.2: interAptiv Core.....	30
1.2.1: MIPS Multi-Thread Technology.....	31
1.2.2: Instruction Fetch Unit .....	32
1.2.3: Thread Scheduling Unit (TSU).....	32
1.2.4: Policy Manager .....	32
1.2.5: interAptiv L1 Caches.....	33
1.2.5.1: Level 1 Instruction Cache .....	33
1.2.5.2: Level 1 Data Cache .....	33
1.2.5.3: Level 1 Cache Memory Configuration .....	33
1.2.6: Memory Management Unit (MMU).....	34
1.2.6.1: Instruction TLB (ITLB) .....	34
1.2.6.2: Data TLB (DTLB).....	35
1.2.6.3: Translation Lookaside Buffer (TLB).....	35
1.2.6.4: Joint TLB (JTLB).....	35
1.2.7: Memory Protection Unit.....	36
1.2.8: Execution Pipelines.....	36
1.2.9: Instruction and Data Scratch Pad RAM .....	36
1.2.10: Inter-thread Communication Unit (ITU).....	36
1.2.11: Bus Interface (BIU).....	37
1.2.11.1: Write Buffer.....	37
1.2.11.2: SimpleBE Mode.....	37
1.2.12: System Control Coprocessor (CP0).....	37
1.2.13: Interrupt Handling.....	38
1.2.14: Modes of Operation.....	38
1.2.15: Floating Point Unit (FPU) .....	38
1.2.15.1: FPU Performance.....	38
1.2.16: interAptiv Core Power Management .....	39
1.2.17: EJTAG Debug Support .....	39
1.2.18: Fast Debug Channel .....	40
1.2.19: PDtrace™.....	40
1.2.20: MIPS16e2 Application Specific Extension .....	40
1.2.21: CorExtend® Unit .....	41
1.3: Multiprocessing System.....	41
1.3.1: Cluster Power Controller (CPC).....	41
1.3.1.1: Cluster Power Controller Reset Control .....	41
1.3.1.2: Optional CPC-basic Power Controller .....	42
1.3.2: Coherence Manager (CM2) .....	42
1.3.2.1: Request Unit (RQU) .....	42
1.3.2.2: Intervention Unit (IVU).....	43
1.3.2.3: System Memory Unit (SMU).....	43
1.3.2.4: Response Unit (RSU).....	43
1.3.2.5: Transaction Routing Unit.....	43
1.3.2.6: Level 2 Cache.....	44

1.3.2.7: CM2 Configuration Registers .....	45
1.3.2.8: PDTrace Unit .....	45
1.3.2.9: Performance Counter Unit .....	45
1.3.2.10: Coherence Manager Performance .....	45
1.3.3: I/O Coherence Unit (IOCU) .....	46
1.3.3.1: Software I/O Coherence .....	47
1.3.3.2: Hardware I/O Coherence .....	47
1.3.4: Global Interrupt Controller .....	48
1.3.5: Global Configuration Registers (GCR) .....	48
1.3.5.1: Reset Control .....	48
1.3.5.2: Inter-CPU Debug Breaks .....	49
1.3.5.3: CM2 Control Registers .....	49
1.3.6: Clocking Options .....	49
1.3.7: Design For Test (DFT) Features .....	50
1.3.7.1: Internal Scan .....	50
1.3.7.2: Memory BIST .....	50
1.4: Build-Time Configuration Options .....	51

**Chapter 2: CP0 Registers..... 53**

2.1: CP0 Register Summary .....	53
2.1.1: CP0 Registers Grouped by Function .....	53
2.1.2: CP0 Registers Grouped by Number .....	57
2.2: CP0 Register Descriptions .....	61
2.2.1: CPU Configuration and Status Registers .....	61
2.2.1.1: Device Configuration — Config (CP0 Register 16, Select 0) .....	61
2.2.1.2: Device Configuration 1 — Config1 (CP0 Register 16, Select 1) .....	64
2.2.1.3: Device Configuration 2 — Config2 (CP0 Register 16, Select 2) .....	66
2.2.1.4: Device Configuration 3 — Config3 (CP0 Register 16, Select 3) .....	69
2.2.1.5: Device Configuration 4 — Config4 (CP0 Register 16, Select 4) .....	70
2.2.1.6: Device Configuration 5 — Config5 (CP0 Register 16, Select 5) .....	71
2.2.1.7: Device Configuration 7 — Config7 (CP0 Register 16, Select 7) .....	73
2.2.1.8: Processor ID — PRId (CP0 Register 15, Select 0) .....	75
2.2.1.9: Exception Base Address — EBase (CP0 Register 15, Select 1) .....	76
2.2.1.10: Status (CP0 Register 12, Select 0) .....	77
2.2.1.11: Interrupt Control — IntCtl (CP0 Register 12, Select 1) .....	82
2.2.2: TLB Management Registers .....	85
2.2.2.1: Index (CP0 Register 0, Select 0) .....	85
2.2.2.2: Random (CP0 Register 1, Select 0) .....	86
2.2.2.3: EntryLo0 - EntryLo1 (CP0 Registers 2 and 3, Select 0) .....	87
2.2.2.4: EntryHi (CP0 Register 10, Select 0) .....	88
2.2.2.5: Context (CP0 Register 4, Select 0) .....	90
2.2.2.6: Context Configuration — ContextConfig (CP0 Register 4, Select 1) .....	90
2.2.2.7: PageMask (CP0 Register 5, Select 0) .....	91
2.2.2.8: Page Granularity — PageGrain (CP0 Register 5, Select 1) .....	92
2.2.2.9: Wired (CP0 Register 6, Select 0) .....	93
2.2.2.10: Bad Virtual Address — BadVAddr (CP0 Register 8, Select 0) .....	94
2.2.3: Memory Segmentation Registers .....	95
2.2.3.1: Segmentation Control 0 — SegCtl0 (CP0 Register 5, Select 2) .....	95
2.2.3.2: Segmentation Control 1 — SegCtl1 (CP0 Register 5, Select 3) .....	96
2.2.3.3: Segmentation Control 2 — SegCtl2 (CP0 Register 5, Select 4) .....	97
2.2.4: Exception Control Registers .....	100
2.2.4.1: Cause (CP0 Register 13, Select 0) .....	100
2.2.4.2: Exception Program Counter — EPC (CP0 Register 14, Select 0) .....	103



2.2.4.3: Error Exception Program Counter — ErrorEPC (CP0 Register 30, Select 0) .....	104
2.2.5: Timer Registers.....	105
2.2.5.1: Count (CP0 Register 9, Select 0) .....	105
2.2.5.2: Compare (CP0 Register 11, Select 0) .....	106
2.2.6: Cache Management Registers.....	106
2.2.6.1: Level 1 Instruction Cache Tag Low — ITagLo (CP0 Register 28, Select 0).....	107
2.2.6.2: Level 1 Instruction Cache Data Low — IDataLo (CP0 Register 28, Select 1).....	109
2.2.6.3: Level 1 Instruction Cache Data High — IDataHi (CP0 Register 29, Select 1).....	109
2.2.6.4: Level 1 Data Cache Tag Low — DTagLo (CP0 Register 28, Select 2) .....	110
2.2.6.5: DTagHi (CP0 Register 29, Select 2): L1 Data Cache and DSPRAM ECC.....	113
2.2.6.6: Level 1 Data Cache Data Low — DDataLo (CP0 Register 28, Select 3) .....	114
2.2.6.7: Level 2/3 Cache Tag Low — L23TagLo (CP0 Register 28, Select 4) .....	114
2.2.6.8: Level 2/3 Cache Data Low — L23DataLo (CP0 Register 28, Select 5) .....	115
2.2.6.9: Level 2/3 Cache Data High — L23DataHi (CP0 Register 29, Select 5) .....	116
2.2.6.10: ErrCtl (CP0 Register 26, Select 0).....	116
2.2.6.11: Cache Error — CacheErr (CP0 Register 27, Select 0).....	119
2.2.7: Thread Context and Shadow Control Registers.....	127
2.2.7.1: TCStatus Register (CP0 Register 2, Select 1) .....	127
2.2.7.2: TCBind Register (CP0 Register 2, Select 2) .....	129
2.2.7.3: TCRestart Register (CP0 Register 2, Select 3) .....	130
2.2.7.4: TCHalt Register (CP0 Register 2, Select 4) .....	130
2.2.7.5: TCContext Register (CP0 Register 2, Select 5) .....	131
2.2.7.6: TCSchedule Register (CP0 Register 2, Select 6).....	131
2.2.7.7: TCScheFBack Register (CP0 Register 2, Select 7) .....	133
2.2.7.8: TCOpt Register (CP0 Register 3, Select 7).....	134
2.2.7.9: SRSConf0 (CP0 Register 6, Select 1).....	134
2.2.7.10: SRSConf1-4 (CP0 Register 6, Select 2-5) .....	135
2.2.7.11: SRSCtl Register (CP0 Register 12, Select 2).....	135
2.2.7.12: SRSSMap Register (CP0 Register 12, Select 3) .....	137
2.2.8: Virtual Processing Element (VPE) Registers .....	138
2.2.8.1: VPE Control Register (CP0 Register 1, Select 1).....	138
2.2.8.2: VPE Conf0 Register (CP0 Register 1, Select 2).....	139
2.2.8.3: VPE Conf1 Register (CP0 Register 1, Select 3).....	140
2.2.8.4: VPE Schedule Register (CP0 Register 1, Select 5) .....	141
2.2.8.5: VPE ScheFBack Register (CP0 Register 1, Select 6).....	142
2.2.8.6: VPE Opt Register (CP0 Register 1, Select 7) .....	143
2.2.8.7: MVPControl Register (CP0 Register 0, Select 1) .....	144
2.2.8.8: MVPConf0 Register (CP0 Register 0, Select 2) .....	144
2.2.8.9: MVPConf1 Register (CP0 Register 0, Select 3) .....	145
2.2.9: Performance Monitoring Registers.....	146
2.2.9.1: Performance Counter 0 and 1 Control Registers (CP0 Register 25, select 0, 2) .....	146
2.2.9.2: Performance Counter 0 and 1 Count Registers (CP0 Register 25, select 1, 3) .....	156
2.2.10: Debug Registers .....	157
2.2.10.1: Debug (CP0 Register 23, Select 0) .....	157
2.2.10.2: Debug Exception Program Counter — DEPC (CP0 Register 24, Select 0) .....	160
2.2.10.3: Debug Save — DESAVE (CP0 Register 31, Select 0).....	161
2.2.10.4: Watch Low 0 - 3 — WatchLo0-3 (CP0 Register 18, Select 0-3).....	162
2.2.10.5: Watch High 0 - 3 — WatchHi0-3 (CP0 Register 19, Select 0-3) .....	162
2.2.11: PDTrace Registers.....	163
2.2.11.1: Trace Control Register — TraceControl (CP0 Register 23, Select 1) .....	164
2.2.11.2: Trace Control 2 Register — TraceControl2 (CP0 Register 23, Select 2) .....	166
2.2.11.3: Trace Control 3 Register — TraceControl3 (CP0 Register 24, Select 2) .....	167
2.2.11.4: User Trace Data 1 Register — UserTraceData1 (CP0 Register 23, Select 3) .....	168

2.2.11.5: User Trace Data 2 Register — UserDataTrace2 (CP0 Register 24, Select 3) .....	169
2.2.11.6: Trace Instruction Breakpoint Condition Register — TracelBPC (CP0 Register 23, Select 4) ..	169
2.2.11.7: Trace Data Breakpoint Condition Register — TraceDBPC (CP0 Register 23, Select 5) ...	170
2.2.12: User Mode Support Registers .....	171
2.2.12.1: YQMask Register (CP0 Register 1, Select 4) .....	172
2.2.12.2: Hardware Enable — HWREna (CP0 Register 7, Select 0) .....	172
2.2.12.3: UserLocal (CP0 Register 4, Select 2) .....	174
2.2.12.4: LLAddr Register (CP0 Register 17, Select 0) .....	174
2.2.13: Kernel Mode Support Registers .....	174
2.2.13.1: Kernel Scratch Register 1 — KScratch1 (CP0 Register 31, Select 2) .....	175
2.2.13.2: Kernel Scratch Register 2 — KScratch2 (CP0 Register 31, Select 3) .....	175
2.2.13.3: Kernel Scratch Register 3 — KScratch3 (CP0 Register 31, Select 4) .....	175
2.2.14: Memory Mapped Registers .....	176
2.2.14.1: Common Device Memory Map Base Address — CDMMBase (CP0 Register 15, Select 2) ....	176
2.2.14.2: Coherency Manager Global Configuration Register Base Address — CMGCRBase (CP0 Register 15, Select 3) .....	177

<b>Chapter 3: Memory Management Unit .....</b>	<b>179</b>
3.1: Introduction .....	179
3.2: Memory Management Unit Architecture .....	180
3.2.1: Instruction TLB (ITLB) .....	180
3.2.2: Data TLB (DTLB) .....	181
3.2.3: Joint TLB (JTLB) .....	181
3.3: Overview of Virtual-to-Physical Address Translation .....	182
3.3.1: Address Translation Flow .....	184
3.4: Relationship of TLB Entries and CP0 Registers .....	186
3.4.1: TLB Tag Entry .....	187
3.4.1.1: VPN2 Field .....	187
3.4.1.2: ASID Field .....	187
3.4.1.3: PageMask Field .....	188
3.4.1.4: Global (G) Bit .....	188
3.4.2: TLB Data Entry .....	188
3.4.2.1: Page Frame Number (PFN) .....	188
3.4.2.2: Flag Fields (C, D, V,) .....	189
3.4.3: Address Translation Examples .....	189
3.5: Enhanced Virtual Address .....	192
3.5.1: Virtual and Physical Address Maps .....	192
3.5.2: Initial EVA Configuration Parameters .....	193
3.5.3: Programmable Segmentation Control .....	193
3.5.3.1: Cache Coherency Attribute Control and the Segmentation Control Registers .....	195
3.5.3.2: Functions of the Config5.K Bit .....	195
3.5.3.3: Setting the Memory Addressing Scheme — SI_EVAReset and CONFIG5.K .....	196
3.5.3.4: Enhanced Virtual Address Detection and Support .....	197
3.5.3.5: Setting the Access Control Mode .....	197
3.5.3.6: Defining the Physical Address Range for Each Memory Segment .....	198
3.5.3.7: Enhanced Virtual Address (EVA) Instructions .....	200
3.6: Boot Exception Vector Relocation in Kernel Mode .....	201
3.6.1: Boot Configurations .....	201
3.6.2: Pins Used to Support Boot Exception Vector Relocation .....	201
3.6.3: Mapping of the Boot Exception Vector in the Legacy Configuration .....	204
3.6.4: Example Mapping of the Boot Exception Vector in the EVA Configuration .....	205

3.6.5: Defining the Boot Exception Vector Overlay Region.....	206
3.6.5.1: Setting the Type of Memory Addressing Mode .....	207
3.6.5.2: Using the SI_UseExceptionBase Pin and CONFIG5.K to Determine How to Calculate the BEV Base Address .....	207
3.6.5.3: Determining the Size and Location of the Overlay Region in Virtual Address Space .....	207
3.6.5.4: Determining the Location of the Overlay Region in Physical Memory.....	209
3.6.6: Switching the Addressing Scheme from Legacy to EVA After Boot-up .....	211
3.7: Indexing the JTLB .....	216
3.8: Hardwiring JTLB Entries.....	217
3.9: JTLB Random Replacement .....	217
3.10: TLB Exception Handling .....	219
3.10.1: Overview of TLB Exception Handling Registers .....	220
3.10.1.1: Context Register.....	220
3.10.1.2: ContextConfig Register .....	220
3.10.1.3: BadVAddr Register.....	221
3.10.2: TLB Exception Flow Examples .....	221
3.10.2.1: Single Level Table Configuration.....	221
3.10.2.2: Dual Level Table Configuration .....	223
3.11: Exception Base Address Relocation .....	226
3.12: TLB Duplicate Entries.....	227
3.13: Modes of Operation .....	227
3.13.1: Virtual Memory Segments.....	228
3.13.1.1: Unmapped Segments.....	230
3.13.1.2: Mapped Segments .....	231
3.13.2: User Mode.....	231
3.13.2.1: User Mode Legacy Configuration .....	231
3.13.2.2: User Mode EVA Configuration .....	231
3.13.3: Supervisor Mode.....	232
3.13.3.1: Supervisor Mode — Legacy Configuration.....	232
3.13.3.2: Supervisor Mode — EVA Configuration .....	234
3.13.4: Kernel Mode.....	235
3.13.4.1: Kernel Mode — Legacy Configuration.....	235
3.13.4.2: Kernel Mode — EVA Configuration .....	237
3.13.5: Debug Mode.....	238
3.13.5.1: Debug Mode, Register (drseg) .....	239
3.13.5.2: Debug Mode, Memory (dmseg).....	239
3.14: TLB Instructions.....	240
<b>Chapter 4: Memory Protection Unit .....</b>	<b>241</b>
4.1: Overview.....	241
4.2: MPU Memory Mapping.....	242
4.2.1: Default Segment Control.....	242
4.2.2: Region Control .....	244
4.2.2.1: Super-Region and Subregion Definition .....	244
4.2.2.2: Region Programming Examples.....	245
4.2.2.3: Overlapping Regions .....	248
4.2.2.4: Subregion Programming.....	249
4.3: MPU Register Mapping .....	250
4.4: MPU Registers .....	251
4.4.1: Memory Protection Unit Status Register (MPU_ACSR) — Device Byte Offset 0 .....	251
4.4.2: Memory Protection Unit Configuration Register (MPU_Config) — Device Byte Offset 0x8.....	252
4.4.3: Memory Protection Unit Segment Control Registers (SegmentControl_N) — Byte Offsets 0x10/14/18/1C.....	253

4.4.4: MPU Region Address and Control Registers .....	254
---	-----

**Chapter 5: Caches ..... 257**

5.1: Cache Configurations .....	257
5.1.1: Cacheability Attributes .....	258
5.2: L1 Instruction Cache.....	259
5.2.1: L1 Instruction Cache Virtual Aliasing .....	260
5.2.2: L1 Instruction Cache Precode Bits.....	260
5.2.3: L1 Instruction Cache Parity .....	260
5.2.4: L1 Instruction Cache Replacement Policy .....	261
5.2.5: L1 Instruction Cache Line Locking.....	261
5.2.6: L1 Instruction Cache Memory Coherence Issues .....	262
5.2.7: Software I-Cache Coherence (JVM, Self-modifying Code).....	262
5.2.8: Instruction Cache Way Prediction .....	262
5.2.9: L1 Instruction Software Cache Management.....	262
5.2.10: L1 Instruction Cache CP0 Register Interface.....	264
5.2.10.1: Config1 Register (CP0 register 16, Select 1) .....	264
5.2.10.2: CacheErr Register (CP0 register 27, Select 0).....	264
5.2.10.3: L1 Instruction Cache TagLo Register (CP0 register 28, Select 0).....	264
5.2.10.4: L1 Instruction Cache DataLo Register (CP0 register 28, Select 1) .....	265
5.2.10.5: L1 Instruction Cache DataHi Register (CP0 register 29, Select 1).....	265
5.2.11: L1 Instruction Cache Initialization .....	265
5.2.11.1: L1 Instruction Cache Initialization Routine .....	265
5.2.11.2: L1 Instruction Cache Initialization Routine Details .....	266
5.2.12: Cache Management When Writing Instructions - the “SYNCl” Instruction.....	268
5.3: L1 Data Cache.....	269
5.3.1: L1 Data Cache Virtual Aliasing .....	270
5.3.2: L1 Data Cache Parity (Optional) .....	271
5.3.3: Level 1 Data Cache Error Detection and Correction (EDC).....	271
5.3.3.1: L1 Data Cache Organization .....	271
5.3.3.2: L1 Data Cache Load/Store Operations .....	271
5.3.3.3: L1 Data Cache Error Types.....	271
5.3.3.4: Store Operations Less than 32-bits.....	272
5.3.3.5: Examples of L1 Data Cache ECC Errors .....	272
5.3.4: L1 Data Cache Replacement Policy .....	273
5.3.5: L1 Data Cache Line Locking.....	274
5.3.6: L1 Data Cache Memory Coherence Protocol .....	274
5.3.7: L1 Data Cache Initialization .....	274
5.3.7.1: L1 Data Cache Initialization Routine .....	275
5.3.7.2: L1 Data Cache Initialization Routine Details .....	275
5.3.8: Data Cache CP0 Register Interface.....	277
5.3.8.1: Config1 Register (CP0 register 16, Select 1) .....	277
5.3.8.2: CacheErr Register (CP0 register 27, Select 0).....	278
5.3.8.3: L1 Data Cache TagLo Register (CP0 register 28, Select 2).....	278
5.3.8.4: L1 Data Cache TagHi Register (CP0 register 29, Select 2) .....	278
5.3.8.5: L1 Data Cache DataLo Register (CP0 register 28, Select 3) .....	278
5.4: L1 Instruction and Data Cache Software Testing.....	278
5.4.1: L1 Instruction Cache Tag Array .....	279
5.4.2: Instruction Cache Data Array.....	279
5.4.3: Instruction Cache Way Select Array .....	279
5.4.4: L1 Data Cache Tag Array .....	279
5.4.5: Duplicate Data Cache Tag Array .....	279
5.4.6: Data Cache Data Array.....	280

5.4.7: Data Cache Way Select Array .....	280
5.5: L2 Cache .....	281
5.5.1: L2 Cache General Features.....	281
5.5.2: 0K L2 Cache Option.....	282
5.5.3: 32K and 64K L2 Cache Options.....	283
5.5.4: OCP Interface .....	283
5.5.5: L2 Replacement Policy .....	283
5.5.6: L2 Allocation Policy .....	283
5.5.6.1: Write .....	283
5.5.6.2: Read .....	284
5.5.7: Write-Through vs. Write-Back.....	285
5.5.8: Cacheable vs. Uncacheable vs. Uncached Accelerated.....	285
5.5.9: Cache Aliases .....	285
5.5.10: Performance Counters .....	285
5.5.11: Sleep Modes .....	286
5.5.11.1: Sleep Mode Using the WAIT Instruction.....	286
5.5.11.2: Internal Dynamic Sleep Mode .....	286
5.5.12: Bypass Mode .....	287
5.5.13: Reduced L2 Hit Latency.....	287
5.5.14: L2-only Sync .....	287
5.5.15: L2 Cache Initialization.....	288
5.5.15.1: init_I2u Cache Initialization Routine.....	288
5.5.15.2: init_I2c Cache Initialization Routine.....	289
5.5.15.3: init_L2u Initialization Routine Details.....	289
5.5.15.4: init_L2c Initialization Routine Details .....	290
5.5.16: L2 Cache CP0 Interface.....	291
5.5.16.1: Config2 Register (CP0 register 16, Select 2) .....	291
5.5.16.2: Error Control Register (CP0 register 26, Select 0) .....	291
5.5.16.3: Cache Error Register (CP0 register 27, Select 0) .....	292
5.5.16.4: L23TagLo Register (CP0 register 28, Select 4).....	292
5.5.16.5: L23DataHi Register(CP0 register 29, Select 5) / L23DataLo Register(CP0 register 28, Select 5) .....	292
5.5.17: L2 Cache Operations .....	292
5.5.17.1: Bus Transaction Equivalence .....	293
5.5.17.2: Details of Cache-ops .....	294
5.5.17.3: Sync in L2.....	295
5.5.17.4: L2 Cache Fetch and Lock.....	296
5.5.18: L2 Cache Error Management.....	296
5.5.18.1: ECC/Parity Support .....	296
5.5.18.2: Tag and Data Formats.....	297
5.5.18.3: Cache Parity Error Handling.....	297
5.5.18.4: Multiple Uncorrectable Errors .....	298
5.5.18.5: Bus Error Handling .....	298
5.6: The CACHE Instruction .....	299
5.6.1: Decoding the Type of Cache Operation.....	299
5.6.2: CACHE Instruction Opcodes.....	299
5.6.3: Way Selection RAM Encoding .....	299
5.7: Write Back Buffer.....	301

<b>Chapter 6: Exceptions and Interrupts.....</b>	<b>303</b>
6.1: Exception Conditions.....	303
6.2: Exception Priority.....	304
6.3: Exception Vector Locations .....	306

6.4: General Exception Processing .....	312
6.5: Debug Exception Processing .....	314
6.6: Exception Descriptions .....	316
6.6.1: Reset Exception (Reset) .....	316
6.6.2: Debug Single Step Exception (DSS).....	317
6.6.3: Debug Interrupt Exception (DINT).....	317
6.6.4: Non-Maskable Interrupt (NMI) Exception.....	318
6.6.5: Interrupt Exception (Int).....	318
6.6.6: Debug Instruction Break Exception (DIB) .....	319
6.6.7: Watch Exception — Instruction Fetch or Data Access (WATCH) .....	319
6.6.8: Address Error Exception — Instruction Fetch/Data Access (AdEL/AdES) .....	320
6.6.9: TLB Refill Exception — Instruction Fetch or Data Access (TLBL/TLBS) .....	321
6.6.10: TLB Invalid Exception — Instruction Fetch or Data Access (TLBINV).....	321
6.6.11: TLB Execute-Inhibit Exception (TLBXI).....	322
6.6.12: TLB Read-Inhibit Exception (TLBRI).....	323
6.6.13: Cache Error Exception (ICache Error/DCache Error) .....	324
6.6.14: Bus Error Exception — Instruction Fetch or Data Access (IBE) .....	325
6.6.15: Debug Software Breakpoint Exception (DBp) .....	325
6.6.16: Execution Exception — System Call (Sys) .....	325
6.6.17: Execution Exception — Breakpoint (Bp).....	326
6.6.18: Execution Exception — Coprocessor Unusable (CpU).....	326
6.6.19: Execution Exception — CorExtend Block Unusable (CEU).....	326
6.6.20: Execution Exception — DSP ASE State Disabled (DSPDis).....	327
6.6.21: Execution Exception — Reserved Instruction (RI).....	327
6.6.22: Execution Exception — Floating Point Exception (FPE).....	327
6.6.23: Execution Exception — Coprocessor2 (C2E) .....	328
6.6.24: Execution Exception — Coprocessor2 (IS1).....	328
6.6.25: Execution Exception — Integer Overflow (Ov).....	328
6.6.26: Execution Exception — Trap (Tr).....	329
6.6.27: Execution Exceptions — MT_ov, MT_under, MT_invalid, MT_yield_sched .....	329
6.6.28: Debug Data Break Exception (DDBL/DDBS).....	330
6.6.29: TLB Modified Exception (TLB Mod) .....	330
6.6.30: Thread Exceptions — MT_gs, MT_gss.....	331
6.6.31: Memory Protection Unit (MPU) Exception — MPUL/MPUS .....	331
6.7: Exception Handling and Servicing Flowcharts .....	332
6.8: Interrupts .....	338
6.8.1: Interrupt Modes.....	338
6.8.1.1: Interrupt Compatibility Mode.....	339
6.8.1.2: Vectored Interrupt Mode.....	341
6.8.1.3: External Interrupt Controller Mode .....	342
6.8.2: Generation of Exception Vector Offsets for Vectored Interrupts .....	344
6.8.3: Global Interrupt Controller.....	345

**Chapter 7: Power Management and the Cluster Power Controller .....** **347**

7.1: Introduction to the Cluster Power Controller.....	347
7.1.1: Power Domains of the interAptiv Multiprocessing System.....	348
7.1.2: Operating Level Transitions .....	348
7.1.2.1: Coherent to Non-Coherent Mode Transition .....	349
7.1.2.2: Non-Coherent to Coherent Mode Transition .....	349
7.1.2.3: Non-Coherent to Power Down Mode Transition.....	349
7.1.2.4: Non-Coherent to Clock Off Mode Transition .....	349
7.1.2.5: Clock Off to Power Down Mode Transition.....	350
7.1.2.6: Clock Off to Non-Coherent Mode Transition .....	350

7.1.2.7: PowerDown to Non-Coherent Mode Transition .....	350
7.2: CPC Register Programming .....	350
7.2.1: Requestor Access to CPC Registers .....	350
7.2.2: Global Sequence Delay Count .....	351
7.2.3: Rail Delay .....	351
7.2.4: Reset Delay .....	352
7.2.5: Executing a Power Sequence .....	353
7.2.6: Accessing Another Core .....	353
7.3: Cluster Power Controller Address Map .....	353
7.3.1: Block Offsets Relative to the Base Address .....	354
7.3.2: Register Offsets Relative to the Block Offsets .....	354
7.3.3: Global Control Block Register Map .....	356
7.3.3.1: Global CSR Access Privilege Register .....	356
7.3.3.2: Global Sequence Delay Counter .....	357
7.3.3.3: Global Rail Delay Counter .....	357
7.3.3.4: Global Reset Width Counter .....	358
7.3.3.5: Revision Register .....	358
7.3.4: Local and Core-Other Control Blocks .....	358
7.3.4.1: Command Register .....	360
7.3.4.2: Core-Other Addressing Register .....	363
7.4: Cluster Power Controller Commands .....	363
7.5: interAptiv Core Power Management Options .....	365
7.6: interAptiv Core Clock Gating .....	365
7.6.1: Designs Implementing Top Level Clock Gating .....	365
7.6.1.1: Instruction Controlled Clock Gating .....	365
7.6.1.2: Register Controlled Clock Gating .....	366
7.6.1.3: Reduction of VDD During Sleep Mode .....	366
7.6.1.4: Restart Latency Trade-Offs .....	366
7.6.2: Designs Not Implementing Top Level Clock Gating .....	367
7.6.3: Designs Implementing Fine Grain Clock Gating .....	367
7.7: interAptiv Core Power Gating .....	367
7.7.1: Hardware Suspend/Resume .....	368
7.7.2: Software Suspend/Resume .....	368
7.7.2.1: Overview of Suspend/Resume Process .....	368
7.7.3: Suspend Process .....	369
7.7.3.1: Save GPR Registers .....	370
7.7.3.2: Save CP0 Registers .....	371
7.7.3.3: Flush Dirty Lines in L1 Data Cache .....	371
7.7.3.4: Save the Resume Address .....	373
7.7.3.5: Copy Memory Power Down Sequence Into Cache .....	373
7.7.3.6: Move Memory to Low Power Mode .....	373
7.7.3.7: Shut Down Power to the interAptiv Core .....	373
7.7.4: Resume Process .....	374
7.7.4.1: System Wake-Up .....	374
7.7.4.2: Power-Up VDD to the interAptiv Core and Assert Power-On Reset .....	374
7.7.4.3: Warm/Cold Boot Detection .....	374
7.7.4.4: Exit Memory Low-Power Mode .....	375
7.7.4.5: Initialize Caches and TLB .....	375
7.7.4.6: Jump to Resume Address .....	375
7.7.4.7: Restore CP0 Registers .....	375
7.7.4.8: Restore GPR Registers .....	376
7.8: CPC-Basic .....	376
7.8.1: Core Power-up State .....	376



7.8.2: CM End Power State Control .....	376
7.8.3: Core and CM Power Up Interlock .....	377
7.8.4: Core and CM Power Up Sequences .....	377
7.8.5: Core Power Down Sequence .....	377
7.8.6: Core and CM Power Down Interlock .....	377
7.8.7: Software Control of Power-States .....	377

## **Chapter 8: Coherency Manager..... 379**

8.1: Coherence Manager Address Map .....	379
8.1.1: Block Offsets Relative to the Base Address .....	380
8.1.2: Register Offsets Relative to the Block Offsets .....	380
8.2: CM2 Programming .....	383
8.2.1: Verifying Overall System Configuration .....	383
8.2.2: Requestor Access to GCR Registers .....	383
8.2.3: CM2 Interface Ports .....	384
8.2.4: Setting the CM2 Register Block Base Address .....	384
8.2.5: Address Regions .....	385
8.2.5.1: Fixed-Size Regions .....	385
8.2.5.2: Variable-Size Regions .....	386
8.2.5.3: Address Region Priorities .....	387
8.2.5.4: Defining the Base Address Location and Size for Each Region .....	387
8.2.5.5: Defining the Target Device .....	388
8.2.5.6: Setting the Cache Coherency Attributes for Region Memory Transfers .....	389
8.2.5.7: Issue Request Protocol and Region Masking .....	389
8.2.5.8: Overlapping Regions .....	390
8.2.6: Address Map Programming Example .....	391
8.2.7: Core-Local GCRs .....	394
8.2.8: Core-Other GCRs .....	394
8.2.9: Accessing Another Cores CM2 GCR Registers .....	394
8.2.10: Boot Exception Vector Configuration .....	395
8.2.11: Coherency Domains .....	398
8.2.12: L2-Only SYNC Operation .....	399
8.2.13: Handling of Addresses Not Mapped to a Defined Region .....	400
8.2.14: Setting the Cache Coherency Attributes for Default Memory Transfers .....	400
8.2.15: In-Flight L1 and L2 Cache Operations .....	401
8.2.16: MIPS System Trace .....	402
8.2.17: Error Processing .....	402
8.2.17.1: Error Codes 1 - 15 .....	404
8.2.17.2: Error Codes 16 - 23 .....	405
8.2.17.3: Error Codes 24 - 26 .....	407
8.2.18: Custom GCR Implementation .....	408
8.2.19: Attribute-Only Regions .....	409
8.3: Global Control Block .....	410
8.3.1: Global Control Block Address Map .....	410
8.3.2: CM2 Configuration Registers .....	412
8.3.2.1: Global Config Register (GCR_CONFIG Offset 0x0000) .....	412
8.3.2.2: GCR Base Register (GCR_BASE Offset 0x0008) .....	414
8.3.2.3: Global CM2 Control Register (GCR_CONTROL Offset 0x0010) .....	415
8.3.2.4: Global CM2 Control2 Register (GCR_CONTROL2 Offset 0x0018) .....	417
8.3.2.5: Global CSR Access Privilege Register (GCR_ACCESS Offset 0x0020) .....	419
8.3.2.6: CM2 Revision Register (GCR_REV Offset 0x0030) .....	420
8.3.2.7: Global CM2 Error Mask Register (GCR_ERROR_MASK Offset 0x0040) .....	421
8.3.2.8: Global CM2 Error Cause Register (GCR_ERROR_CAUSE Offset 0x0048) .....	421



8.3.2.9: Global CM2 Error Address Register (GCR_ERROR_ADDR Offset 0x0050) .....	422
8.3.2.10: Global CM2 Error Multiple Register (GCR_ERROR_MULT Offset 0x0058) .....	422
8.3.2.11: GCR Custom Base Register (GCR_CUSTOM_BASE Offset 0x0060).....	423
8.3.2.12: GCR Custom Status Register (GCR_CUSTOM_STATUS Offset 0x0068) .....	423
8.3.2.13: L2-Only Sync Base Register (GCR_L2_ONLY_SYNC_BASE Offset 0x0070) .....	424
8.3.3: CM2 Region Address Map Registers.....	424
8.3.3.1: Global Interrupt Controller Base Address Register (GCR_GIC_BASE Offset 0x0080).....	424
8.3.3.2: Cluster Power Controller Base Address Register (GCR_CPC_BASE Offset 0x0088).....	425
8.3.3.3: CM2 Region [0 - 3] Base Address Register (GCR_REGn_BASE Offsets 0x0090, 0x00A0, 0x00B0, 0x00C0).....	425
8.3.3.4: CM2 Region [0 - 3] Address Mask Register (GCR_REGn_MASK Offsets 0x0098, 0x00A8, 0x00B8, 0x00C8).....	426
8.3.4: CM2 Status and Revision Registers .....	427
8.3.4.1: Global Interrupt Controller Status Register (GCR_GIC_STATUS Offset 0x00D0).....	427
8.3.4.2: Cache Revision Register (GCR_CACHE_REV Offset 0x00E0).....	428
8.3.4.3: Cluster Power Controller Status Register (GCR_CPC_STATUS Offset 0x00F0) .....	428
8.3.4.4: CM Arbiter Priority Register (GCR_ARB_PRI Offset 0x0160).....	429
8.3.4.5: IOCU Revision Register (GCR_IOCU1_REV Offset 0x0200) .....	429
8.3.5: CM2 Attribute-Only Region Address Map Registers.....	430
8.3.5.1: CM2 Attribute-Only Region [0 - 3] Base Address Registers (GCR_REGn_ATTR_BASE Offsets 0x0190, 0x01A0, 0x0210, 0x0220).....	430
8.3.5.2: CM Attribute-Only Region[0 - 3] Address Mask Registers (GCR_REGn_ATTR_MASK Offsets 0x0198, 0x1A8, 0x218, 0x228).....	431
8.4: Core-Local and Core-Other Control Blocks.....	432
8.4.1: Core-Local and Core-Other Control Blocks Address Map .....	432
8.4.1.1: Core Local Coherence Control Register (GCR_Cx_COHERENCE Offset 0x0008) .....	433
8.4.1.2: Core Local Config Register .....	435
8.4.1.3: Core-Other Addressing Register .....	436
8.4.1.4: Core Local Reset Exception Base Register (GCR_Cx_RESET_BASE Offset 0x0020).....	436
8.4.1.5: Core Local Identification Register (GCR_Cx_ID Offset 0x0028).....	437
8.4.1.6: Core Local Reset Exception Extended Base Register (GCR_Cx_RESET_EXT_BASE Offset 0x0030).....	438
8.4.1.7: Core Local TCID Registers (GCR_Cx_TCID_n_PRIORITY) .....	440
8.5: Global Debug Control Block .....	441
8.5.1: Global Debug Control Block Address Map.....	441
8.5.1.1: CM2 PDTrace TCB ControlB Register (GCR_DB_TCBCONTROLB Offset 0x0008) .....	442
8.5.1.2: CM2 PDTrace TCB ControlD Register (GCR_DB_TCBCONTROLD Offset 0x0010) .....	446
8.5.1.3: CM2 PDTrace TCB ControlE Register (GCR_DB_TCBCONTROLE Offset 0x0020) .....	448
8.5.1.4: CM2 PDTrace TCB Config Register (GCR_DB_TCBConfig Offset 0x0028) .....	448
8.5.1.5: CM2 Performance Counter Control Register (GCR_DB_PC_CTL Offset 0x0100) .....	449
8.5.1.6: CM2 PDTrace TCB Trace Word Read Pointer Register (GCR_DB_TCBRDP Offset 0x0108) ..	451
8.5.1.7: CM2 PDTrace TCB Trace Word Write Pointer Register (GCR_DB_TCBWRP Offset 0x0110) .	451
8.5.1.8: CM2 PDTrace TCB Trace Word Start Pointer Register (GCR_DB_TCBSTP Offset 0x0118)....	451
8.5.1.9: CM2 PDTrace TCB System Trace User Control Register ( GCR_DB_TCBSYS Offset 0x0040)	452
8.5.1.10: CM2 Performance Counter Overflow Status Register (GCR_DB_PC_OV Offset 0x120) ..	452
8.5.1.11: CM2 Performance Counter Event Select Register (GCR_DB_PC_EVENT Offset 0x130) 453	
8.5.1.12: CM2 Cycle Counter Register.....	453
8.5.1.13: CM2 Performance Counter n Qualifier Field Register (GCR_DB_PC_QUALn Offset 0x190, 0x1a0).....	454

8.5.1.14: CM2 Performance Counter n Register (GCR_DB_PC_CNTn Offset 0x198, 0x1A8) .....	454
8.5.1.15: CM2 PDTrace TCB Trace Word LO Register ( GCR_DB_TCBTW_LO Offset 0x0200) ....	454
8.5.1.16: CM2 PDTrace TCB Trace Word HI Register ( GCR_DB_TCBTW_HI Offset 0x0208) .....	455

<b>Chapter 9: Global Interrupt Controller .....</b>	<b>457</b>
9.1: GIC Terminology .....	457
9.2: GIC Features .....	458
9.3: GIC Address Map Overview .....	458
9.3.1: GIC Base Address .....	460
9.3.2: Block Offsets Relative to the Base Address .....	460
9.3.3: Register Offsets Relative to the Block Offsets .....	460
9.4: GIC Programming .....	462
9.4.1: Setting the GIC Base Address and Enabling the GIC .....	462
9.4.2: Configuring Interrupt Sources .....	462
9.4.2.1: Trigger Type Register Group .....	463
9.4.2.2: Edge Type Register Group .....	463
9.4.2.3: Polarity Type Register Group .....	463
9.4.3: Interrupt Routing .....	464
9.4.3.1: Mapping an Interrupt Source to a Processor .....	465
9.4.3.2: Mapping and Interrupt Source to a Specific Processor Pin .....	465
9.4.3.3: Mapping an Interrupt Source to a Register Set .....	465
9.4.4: Enabling, Disabling, and Polling Interrupts .....	465
9.4.4.1: Enabling External Interrupts .....	466
9.4.4.2: Disabling External Interrupts .....	466
9.4.4.3: Determining the Enabled or Disabled Interrupt State .....	466
9.4.4.4: Polling for an Active Interrupt .....	466
9.4.4.5: Programming Example .....	466
9.4.5: Inter-processor Interrupts .....	467
9.4.5.1: WEDGE Register Programming Example .....	468
9.4.5.2: Inter-Processor Interrupt Code Example .....	469
9.4.5.3: Example of Sending an Inter-Processor Interrupt .....	470
9.4.5.4: Example of Clearing an Inter-Processor Interrupt .....	471
9.4.6: Local Device Interrupt Configuration .....	471
9.4.6.1: GIC Interval Timer .....	471
9.4.6.2: GIC Watchdog Timer .....	474
9.4.7: Local Interrupt Routing .....	478
9.4.7.1: Routability of Local Interrupts .....	478
9.4.7.2: Routing Local Interrupts .....	478
9.4.7.3: Watchdog Timer Interrupts .....	481
9.4.7.4: Count and Compare Interrupts .....	482
9.4.7.5: Timer Interrupts .....	482
9.4.7.6: Performance Counter Interrupts .....	482
9.4.7.7: Software Interrupts .....	483
9.4.7.8: Fast Debug Channel Interrupts .....	483
9.4.8: EIC Mode Setting .....	484
9.4.9: Enabling, Disabling, and Polling Local Interrupts .....	484
9.4.9.1: Enabling External Interrupts .....	484
9.4.9.2: Disabling External Interrupts .....	484
9.4.9.3: Determining the Enabled or Disabled Interrupt state .....	485
9.4.9.4: Polling for an Active Interrupt .....	485
9.4.10: Debug Interrupt Generation .....	486
9.5: Shared Register Set .....	487
9.5.1: GIC Register Field Types .....	487

9.5.2: Shared Section Register Map .....	488
9.5.3: Shared Section Register Descriptions .....	492
9.5.3.1: Global Config Register .....	492
9.5.3.2: GIC CounterLo .....	494
9.5.3.3: GIC CounterHi .....	494
9.5.3.4: GIC Revision Register .....	495
9.5.3.5: Global Interrupt Polarity Registers.....	495
9.5.3.6: Global Interrupt Trigger Type Registers .....	496
9.5.3.7: Global Interrupt Dual Edge Registers.....	497
9.5.3.8: Global Interrupt Write Edge Register .....	498
9.5.3.9: Global Interrupt Reset Mask Registers.....	498
9.5.3.10: Global Interrupt Set Mask Registers.....	499
9.5.3.11: Global Interrupt Mask Registers .....	500
9.5.3.12: Global Interrupt Pending Registers .....	501
9.5.3.13: Global Interrupt Map to Pin Registers .....	502
9.5.3.14: Global Interrupt Map to VPE Registers .....	504
9.5.3.15: DINT Send to Group Register .....	505
9.6: GIC VPE-Local and VPE-Other Register Set .....	506
9.6.1: VPE-Local and VPE-Other Register Maps.....	506
9.6.2: VPE-Local and VPE-Other Section Register Description .....	508
9.6.2.1: Local Interrupt Control Register.....	508
9.6.2.2: Local Interrupt Pending Register .....	509
9.6.2.3: Local Interrupt Mask Register.....	509
9.6.2.4: Local Interrupt Reset Mask Register .....	510
9.6.2.5: Local Interrupt Set Mask Register .....	511
9.6.2.6: Local Map to Pin Registers.....	511
9.6.2.7: VPE-Other Addressing Register .....	513
9.6.2.8: VPE-Local Identification Register .....	513
9.6.3: Local Timer Register Descriptions .....	514
9.6.3.1: Watchdog Timer Config Register .....	514
9.6.3.2: Watchdog Timer Count Register .....	515
9.6.3.3: Watchdog Timer Initial Count Register.....	516
9.6.3.4: CompareLo Register .....	516
9.6.3.5: VPE-Local CompareHi Register .....	516
9.6.3.6: Local EIC Shadow Set Registers .....	517
9.6.3.7: VPE-Local DINT Group Participate Register.....	518
9.6.3.8: VPE-Local DebugBreak Group Register .....	519
9.7: GIC User-Mode Visible Section.....	520
9.8: GIC Timing Diagrams .....	521
9.8.1: Case 1: GIC Handling of Multiple Active Interrupts Assigned to the Same VPE .....	521
9.8.1.1: Case 1.1: Multiple Positive Edge-Triggered Interface Routed to the Same VPE in non-EIC Mode.....	522
9.8.1.2: Case 1.2: Multiple Positive Edge Triggered Interrupts Routed to the Same VPE in EIC Mode .	523
9.8.1.3: Case 1.3 - Multiple Active High-Level Interrupts Routed to the Same VPE in non-EIC Mode ...	524
9.8.1.4: Case 1.4: Multiple Active High-Level Interrupts Routed to the Same VPE in non-EIC Mode ....	525
9.8.2: Case 2: GIC Passing Interrupts to VPE .....	525
9.8.2.1: Case 2.1: VPE Taking Interrupts in non-EIC Mode .....	525
9.8.2.2: VPE Taking Interrupts in EIC Mode.....	527
<b>Chapter 10: Policy Manager.....</b>	<b>529</b>

10.1: Thread Scheduling Unit .....	529
10.2: Policy Managers .....	530
10.2.1: Basic Round-Robin Policy Manager .....	530
10.2.2: Weighted Round-Robin Policy Manager (WRR) .....	530
10.2.3: Enhanced Weighted Round-Robin Policy Manager (WRR2) .....	531
10.2.3.1: Throttle Functionality and Operation .....	531
10.2.4: Group Rotation Schedule .....	531
10.2.5: CP0 Register Interface .....	534
<b>Chapter 11: Inter-Thread Communication Unit .....</b>	<b>535</b>
11.1: Features Overview .....	535
11.2: ITC Storage .....	535
11.3: ITC Views .....	536
11.3.1: Bypass View .....	537
11.3.2: Control View .....	537
11.3.3: Empty/Full Synchronized View .....	537
11.3.4: Empty/Full Try View .....	537
11.3.5: P/V Synchronized View .....	537
11.3.6: P/V Try View .....	538
11.3.7: Reserved Views .....	538
11.4: ITC Address Space .....	538
<b>Chapter 12: Instruction and Data Scratch Pad RAM .....</b>	<b>541</b>
12.1: Scratchpad RAM (SPRAM) Features .....	541
12.2: SPRAM Overview .....	542
12.2.1: SPRAM Differences Versus a Cache .....	542
12.2.2: Uncacheable References to SPRAM .....	543
12.2.3: Independent Tag/Data Accesses .....	543
12.2.4: SPRAM Tag Reads and Writes .....	543
12.2.5: Multiple Cycle Data SPRAM Access .....	543
12.2.6: Multiple Cycle Instruction SPRAM Access .....	543
12.2.7: Backstalling the SPRAM Data Port .....	544
12.2.8: Access Granularity .....	544
12.2.9: Connecting I/O Devices to the Data Scratchpad Interface .....	544
12.2.10: Null Connection to Unused SPRAM Interface .....	544
12.3: SPRAM Initialization .....	545
12.3.1: ISPRAM Boot .....	545
12.3.2: ISPRAM Precode Bits .....	546
12.4: SPRAM Clocking .....	547
12.4.1: Scratchpad Reference Design .....	548
12.4.2: Tag Registers in the Reference SPRAM .....	548
12.4.3: Enabling SPRAM Access .....	549
12.4.4: SPRAM BIST Support .....	549
12.4.5: SPRAM Parity Support .....	549
<b>Chapter 13: Hardware and Software Initialization .....</b>	<b>551</b>
13.1: Hardware-Initialized Processor State .....	551
13.1.1: Coprocessor 0 State .....	551
13.1.2: TLB Initialization .....	552
13.1.3: Bus State Machines .....	552
13.1.4: Static Configuration Inputs .....	552
13.1.5: Fetch Address .....	552

13.2: Software-Initialized Processor State.....	552
13.2.1: Register File.....	552
13.2.2: TLB.....	553
13.2.3: Caches.....	553
13.2.4: Coprocessor 0 State.....	553
13.3: Boot and CMP Bringup.....	554

**Chapter 14: Floating-Point Unit..... 555**

14.1: Features Overview.....	555
14.2: IEEE Standard 754.....	556
14.3: Enabling the Floating-Point Coprocessor.....	556
14.4: Data Formats.....	556
14.4.1: Floating-Point Formats.....	556
14.4.1.1: Normalized and Denormalized Numbers.....	558
14.4.1.2: Reserved Operand Values—Infinity and NaN.....	558
14.4.1.3: Infinity and Beyond.....	558
14.4.1.4: Signalling Non-Number (SNaN).....	559
14.4.1.5: Quiet Non-Number (QNaN).....	559
14.4.2: Fixed-Point Formats.....	559
14.5: Floating-Point General Registers.....	560
14.5.1: FPRs and Formatted Operand Layout.....	560
14.5.2: Formats of Values Used in Floating Point Registers.....	561
14.5.3: Binary Data Transfers (32-Bit and 64-Bit).....	562
14.6: Floating-Point Control Registers.....	563
14.6.1: Floating-Point Implementation Register (FIR, CP1 Control Register 0).....	564
14.6.2: Floating-Point Condition Codes Register (FCCR, CP1 Control Register 25).....	566
14.6.3: Floating-Point Exceptions Register (FEXR, CP1 Control Register 26).....	566
14.6.4: Floating-Point Enables Register (FENR, CP1 Control Register 28).....	567
14.6.5: Floating-Point Control and Status Register (FCSR, CP1 Control Register 31).....	568
14.6.6: Operation of the FS/FO/FN Bits.....	570
14.6.6.1: Flush To Zero Bit.....	571
14.6.6.2: Flush Override Bit.....	571
14.6.6.3: Flush to Nearest.....	572
14.6.6.4: Recommended FS/FO/FN Settings.....	572
14.6.7: FCSR Cause Bit Update Flow.....	573
14.6.7.1: Exceptions Triggered by CTC1.....	573
14.6.7.2: Generic Flow.....	573
14.6.7.3: Multiply-Add Flow.....	573
14.6.7.4: Cause Update Flow for Input Operands.....	574
14.6.7.5: Cause Update Flow for Unimplemented Operations.....	574
14.7: Exceptions.....	574
14.7.1: Precise Exception Mode.....	574
14.7.2: Exception Conditions.....	575
14.7.2.1: Invalid Operation Exception.....	575
14.7.2.2: Division By Zero Exception.....	576
14.7.2.3: Underflow Exception.....	576
14.7.2.4: Overflow Exception.....	576
14.7.2.5: Inexact Exception.....	576
14.7.2.6: Unimplemented Operation Exception.....	577
14.8: Latency and Repeat Rates.....	577
14.9: Instruction Overview.....	578
14.9.1: Data Transfer Instructions.....	578
14.9.1.1: Data Alignment in Loads, Stores, and Moves.....	578

14.9.1.2: Addressing Used in Data Transfer Instructions .....	579
14.9.2: Arithmetic Instructions .....	579
14.9.3: Conversion Instructions .....	580
14.9.4: Formatted Operand-Value Move Instructions .....	581
14.9.5: Conditional Branch Instructions .....	582
14.9.6: Miscellaneous Instructions .....	582
14.10: Alphabetical Listing of Floating Point Instructions .....	584

## **Chapter 15: MIPS DSP-R2 Application Specific Extension ..... 587**

15.1: MIPS32® DSP-R2 ASE Features.....	587
15.1.1: Q31 and Q15 Signed 16-bit Fractions.....	588
15.1.2: Saturating Arithmetic.....	588
15.1.3: <i>Multiplying Fractions</i> .....	588
15.1.4: <i>Rounding</i> .....	588
15.1.5: Multiply-Accumulate Sequences .....	588
15.1.6: SIMD Operations.....	588
15.2: Common Applications.....	589
15.3: Software Detection of the DSP ASE Revision 2 .....	589
15.4: DSP-R2 Registers .....	589
15.4.1: DSP-R2 Accumulator Registers.....	589
15.4.2: DSP-R2 ASE Control Register.....	589
15.5: DSP-R2 Instruction Types .....	592
15.5.1: Hints in instruction names .....	592
15.5.2: Arithmetic — 64-bit.....	592
15.5.3: Arithmetic — Saturating and/or SIMD Types .....	593
15.5.4: Bit-shifts — Saturating and/or SIMD Types .....	593
15.5.5: Comparison and “Conditional-Move” Operations on SIMD Types .....	593
15.5.6: Conversions to and from SIMD Types .....	594
15.5.7: Multiplication - SIMD Types with Result in GP Register .....	594
15.5.8: Multiply Q15s from Paired-Half and Accumulate .....	595
15.5.9: Load with Register + Register Address.....	595
15.5.10: DSP-R2 Control Register Access .....	595
15.5.11: Accumulator Access Instructions .....	595
15.5.12: Dot Products and Building Blocks for Complex Multiplication.....	596
15.5.13: Other DSP-R2 ASE Instructions .....	597
15.6: DSP-R2 Code Optimization for Performance .....	597
15.6.1: Pixel Unpacking Example .....	598
15.6.2: Sum of Absolute Differences Example.....	599
15.6.3: Bit Stream Unpacking Example .....	599
15.7: Compiler Usage Model for MIPS DSP-R2.....	600
15.7.1: Enabling the DSP-R2 ASE in the Compiler.....	600
15.7.2: Data Types.....	601
15.7.3: Initialization of Q15 and Q31 Variables.....	601
15.7.4: Initialization of SIMD Variables .....	601
15.7.5: Accessing Elements in SIMD Variables .....	602
15.7.6: C Operators.....	603
15.7.7: C Intrinsics for the MIPS32 DSP-R2 ASE .....	604
15.7.8: Compiler Usage and the DSPControl Register .....	606
15.7.9: C-Based Intrinsics for the MIPS32® DSP-R2 ASE .....	607
15.7.9.1: Intrinsics for Instructions that Access and Use the DSPControl Register .....	607
15.7.9.2: Using Intrinsics for Signed and Unsigned 8-bit Integers.....	608
15.7.9.3: Using Intrinsics for Q15 Data Type.....	612
15.7.9.4: Using Intrinsics for Q31 Data Type.....	616

15.7.9.5: Using Intrinsic for Mixed Data Types: 8-bit Integers and Q15/16-bit Integers .....	618
15.7.9.6: Using Intrinsic for Mixed Data Types: Q15 and Q31 .....	620
15.7.9.7: Using Intrinsic for 64-bit Accumulators .....	620
15.7.9.8: Using Intrinsic for 32-bit Integers .....	622
15.7.9.9: Using Intrinsic for 16-bit Integers .....	625
15.7.9.10: Using Intrinsic for Mixed Data Types: 16-bit and 32-bit Integers .....	627
15.8: Code Optimizations for the MIPS DSP-R2 .....	627
15.8.1: The Use of Intrinsic Versus ASM Macros .....	627
15.8.2: Using Accumulators .....	628
15.8.3: Multiply "32-bit * 32-bit = 64-bit" .....	628
15.8.4: Multiply and Add "32-bit * 32-bit + 64-bit = 64-bit" .....	629
15.8.5: Array Alignment and Data Layout .....	629
15.8.6: GP-Relative Addressing .....	630
15.8.7: Fixed Registers and Register Variables .....	631
15.8.8: Conditional Moves .....	634
15.9: Programming Examples .....	635
15.9.1: The FIR Filter in Traditional C .....	635
15.9.2: The FIR Filter in Hand-Tuned Assembly .....	636
15.9.3: The FIR Filter in Efficient C .....	636
15.10: MIPS32 DSP-R2 Intrinsic .....	637
15.10.1: Immediate Intrinsic .....	637
15.10.2: Intrinsic for DSPControl Register .....	637
15.10.3: Intrinsic for Signed and Unsigned 8-bit Integers .....	637
15.10.4: Intrinsic for Q15 .....	638
15.10.5: Intrinsic for Q31 .....	639
15.10.6: Intrinsic for Mixed Data Types: 8-bit Integers and Q15/16-bit Integers .....	639
15.10.7: Intrinsic for Mixed Data Types: Q15 and Q31 .....	639
15.10.8: Intrinsic for 64-bit Accumulators .....	639
15.10.9: Intrinsic for 32-bit Integers .....	640
15.10.10: Intrinsic for 16-bit Integers .....	640
15.10.11: Intrinsic for Mixed Data Types: 16-bit and 32-bit Integers .....	640
15.11: DSP-R2 ASE Instruction Groups .....	641
15.12: Listing of DSP-R2 ASE Instruction Groups .....	655
15.13: DSP Instruction Latencies and Repeat Rates .....	660

<b>Chapter 16: EJTAG Debug Support .....</b>	<b>663</b>
16.1: Overview .....	664
16.2: Trace Funnel and Trace Types .....	665
16.2.1: Trace Types .....	666
16.2.2: EJTAG TAP Interface .....	666
16.2.3: EJTAGBOOT vs NORMALBOOT .....	666
16.3: Detecting Debug Mode .....	666
16.4: Ways of Entering Debug Mode .....	667
16.4.1: EJTAG Debug Single Step .....	667
16.4.2: EJTAG Debug Interrupt .....	667
16.4.3: EJTAG Hardware Data Breakpoint Match .....	667
16.4.4: EJTAG Hardware Instruction Breakpoint Match .....	668
16.4.5: EJTAG Software Breakpoint .....	668
16.5: Exiting Debug Mode .....	668
16.6: EJTAG and PDTrace Revisions .....	668
16.7: Connection Options .....	669
16.8: Hardware Breakpoints .....	670
16.8.1: Instruction Breakpoints .....	670



16.8.2: Data Breakpoints.....	670
16.8.3: Instruction Breakpoint Registers Overview .....	670
16.8.4: Data Breakpoint Registers Overview .....	671
16.8.5: Conditions for Matching Breakpoints .....	671
16.8.5.1: Conditions for Matching Instruction Breakpoints .....	671
16.8.5.2: Conditions for Matching Data Breakpoints .....	672
16.8.6: Debug Exceptions from Breakpoints.....	673
16.8.6.1: Debug Exception by Instruction Breakpoint.....	673
16.8.6.2: Debug Exception by Data Breakpoint.....	673
16.8.7: Breakpoint used as Triggerpoint .....	675
16.9: Debug Vector Addressing.....	675
16.10: Test Access Port (TAP) .....	676
16.10.1: EJTAG Internal and External Interfaces.....	677
16.10.2: Test Access Port Operation .....	677
16.10.2.1: Test-Logic-Reset State .....	678
16.10.2.2: Run-Test/Idle State.....	678
16.10.2.3: Select_DR_Scan State .....	679
16.10.2.4: Select_IR_Scan State .....	679
16.10.2.5: Capture_DR State .....	679
16.10.2.6: Shift_DR State.....	679
16.10.2.7: Exit1_DR State .....	679
16.10.2.8: Pause_DR State.....	679
16.10.2.9: Exit2_DR State .....	679
16.10.2.10: Update_DR State .....	680
16.10.2.11: Capture_IR State .....	680
16.10.2.12: Shift_IR State .....	680
16.10.2.13: Exit1_IR State.....	680
16.10.2.14: Pause_IR State .....	680
16.10.2.15: Exit2_IR State.....	680
16.10.2.16: Update_IR State .....	680
16.10.3: Test Access Port (TAP) Instructions .....	681
16.10.3.1: BYPASS Instruction.....	681
16.10.3.2: IDCODE Instruction .....	681
16.10.3.3: IMPCODE Instruction .....	682
16.10.3.4: ADDRESS Instruction.....	682
16.10.3.5: DATA Instruction .....	682
16.10.3.6: CONTROL Instruction .....	682
16.10.3.7: ALL Instruction.....	682
16.10.3.8: EJTAGBOOT Instruction .....	682
16.10.3.9: NORMALBOOT Instruction .....	683
16.10.3.10: FASTDATA Instruction .....	683
16.10.3.11: TCBCONTROLA Instruction.....	684
16.10.3.12: TCBCONTROLB Instruction.....	684
16.10.3.13: TCBCONTROLC Instruction.....	684
16.10.3.14: TCBDATA Instruction .....	684
16.10.3.15: PCSAMPLE Instruction .....	684
16.10.3.16: TCBCONTROLD Instruction.....	684
16.10.3.17: TCBCONTROLE Instruction.....	685
16.10.3.18: FDC Instruction.....	685
16.10.4: TAP Processor Accesses.....	685
16.10.4.1: Fetch/Load and Store From/To the EJTAG Probe Through dmseg .....	685
16.11: PDTrace .....	686
16.11.1: Processor Modes .....	687



16.11.2: Software Versus Hardware Control.....	687
16.11.3: Trace Information .....	687
16.11.4: Load/Store Address and Data Trace Information.....	688
16.11.5: Programmable Processor Trace Mode Options.....	689
16.11.6: Programmable Trace Information Options .....	689
16.11.6.1: User Data Trace .....	689
16.11.7: Enable Trace to Probe On-Chip Memory.....	689
16.11.8: Enabling PDtrace .....	690
16.11.8.1: Trace Trigger from EJTAG Hardware Instruction/Data Breakpoints .....	690
16.11.8.2: Turning On PDtrace™ Trace.....	690
16.11.8.3: Turning Off PDtrace™ Trace .....	692
16.12: PDtrace Cycle-by-Cycle Behavior .....	693
16.12.1: FIFO Logic in PDtrace and TCB Modules.....	693
16.12.2: Handling of FIFO Overflow in the PDtrace Module .....	693
16.12.3: Handling of FIFO Overflow in the TCB.....	694
16.12.3.1: Probe Width and Clock-ratio Settings.....	694
16.12.4: Adding Cycle Accurate Information to the Trace.....	694
16.13: PC Sampling.....	695
16.13.1: PC Sampling in Wait State.....	695
16.14: EJTAG Registers.....	695
16.14.1: General Purpose Control and Status .....	695
16.14.1.1: Debug Control Register .....	696
16.14.1.2: DebugVectorAddr Register.....	700
16.14.2: Instruction Breakpoint Registers .....	701
16.14.2.1: Instruction Breakpoint Status (IBS) Register .....	702
16.14.2.2: Instruction Breakpoint Address n (IBAn) Register .....	702
16.14.2.3: Instruction Breakpoint Address Mask n (IBMn) Register.....	703
16.14.2.4: Instruction Breakpoint ASID n (IBASIDn) Register.....	703
16.14.2.5: Instruction Breakpoint Control n (IBCn) Register .....	704
16.14.3: Data Breakpoint Registers .....	705
16.14.3.1: Data Breakpoint Status (DBS) Register .....	705
16.14.3.2: Data Breakpoint Address n (DBAn) Register .....	706
16.14.3.3: Data Breakpoint Address Mask n (DBMn) Register .....	706
16.14.3.4: Data Breakpoint ASID n (DBASIDn) Register .....	707
16.14.3.5: Data Breakpoint Control n (DBCn) Register .....	707
16.14.3.6: Data Breakpoint Value n (DBVn) Register .....	709
16.14.3.7: Data Breakpoint Value High n (DBVHn) Register .....	709
16.14.4: EJTAG TAP Registers .....	710
16.14.4.1: Instruction Register.....	710
16.14.4.2: Data Registers Overview .....	710
16.14.4.3: Bypass Register .....	710
16.14.4.4: Device Identification (ID) Register .....	710
16.14.4.5: Implementation Register.....	711
16.14.4.6: EJTAG Control Register .....	713
16.14.5: Processor Access Registers .....	718
16.14.5.1: Processor Access Address Register .....	718
16.14.5.2: Processor Access Data Register .....	718
16.14.6: Fastdata Registers .....	718
16.14.6.1: Fastdata Register (TAP Instruction FASTDATA) .....	718
16.14.7: FDC TAP Register .....	719
16.14.8: Fast Debug Channel Registers .....	720
16.14.8.1: FDC Access Control and Status (FDACSR) Register (Offset 0x0) .....	720
16.14.8.2: FDC Configuration (FDCFG) Register (Offset 0x8).....	721

16.14.8.3: FDC Status (FDSTAT) Register (Offset 0x10) .....	722
16.14.8.4: FDC Receive (FDRX) Register (Offset 0x18) .....	723
16.14.8.5: FDC Transmit n (FDTXn) Registers (Offset 0x20 + 0x8*n) .....	723
16.14.9: PDTrace™ Registers (Software Control) .....	724
16.14.10: Trace Control Block (TCB) Registers (Hardware Control) .....	724
16.14.10.1: TCBCONTROLA Register .....	725
16.14.10.2: TCBCONTROLB Register .....	728
16.14.10.3: TCBDATA Register .....	729
16.14.10.4: TCBCONTROLC Register .....	730
16.14.10.5: TCBCONTROLD Register .....	732
16.14.10.6: TCBCONTROLE Register .....	732
16.14.10.7: TCBCONFIG Register (Reg 0) .....	733
16.14.10.8: TCBTRIGx Register (Reg 16-23) .....	735
16.14.11: Register Reset State .....	737
16.15: Fast Debug Channel.....	738
16.15.1: Common Device Memory Map.....	738
16.15.2: Fast Debug Channel Interrupt.....	738
16.15.3: Core FDC Buffers.....	739
16.15.4: Sleep mode .....	741
16.16: TCB Trigger Logic .....	741
16.16.1: TCB Trace Enabling.....	741
16.16.2: Tracing a Reset Exception .....	741
16.16.3: Trigger Units Overview.....	741
16.16.4: Trigger Source Unit .....	742
16.16.5: Trigger Control Units .....	743
16.16.6: Trigger Action Unit .....	743
16.16.7: Simultaneous Triggers .....	743
16.16.7.1: Prioritized Trigger Actions .....	743
16.16.7.2: OR'ed Trigger Actions .....	744
<b>Chapter 17: Multi-CPU Debug.....</b>	<b>745</b>
17.1: CM Performance Counters.....	745
17.1.1: CM Performance Counter Functionality .....	745
17.1.2: Performance Counter Usage Models.....	746
17.1.3: CM Performance Counter Event Types and Qualifiers .....	749
17.2: Debug Mode Triggering.....	757
17.2.1: Selecting CPUs to Enter Debug Mode.....	757
17.2.2: Debug Mode Groups and Cross Triggering .....	757
17.2.3: VPE to CPU mapping in Debug Cross Trigger Facility .....	758
17.2.4: Debug Cross Trigger Facility and Power Management .....	758
17.3: PDTrace Software Architecture .....	758
17.3.1: CM Trace Functionality .....	759
17.3.1.1: CM Trace Configuration and Control .....	759
17.3.1.2: System Trace Interface Configuration and Control .....	760
17.3.1.3: Trace Funnel Enable .....	761
17.3.1.4: CM Trace Formats.....	761
17.3.1.5: CM / CPU Core Trace Correlation.....	761
17.3.2: Controlling Trace in a Multi-CPU Multiprocessing System.....	762
17.3.3: EJTAG Debug Support in the interAptiv Coherence Manager.....	762
17.3.3.1: Test Access Port (TAP) .....	762
17.3.3.2: Test Access Port (TAP) Instructions.....	767
17.3.3.3: EJTAG TAP Registers.....	768
17.3.3.4: Data Registers Overview .....	769

17.3.3.5: CM2 Trace Control Block (TCB) Registers.....	771
17.3.4: MIPS Trace Capability .....	783
17.3.5: Memory-Mapped Access to PDtrace™ Control and On-Chip Trace RAM.....	783
17.3.6: On-Chip Trace Buffer Usage.....	783
<b>Chapter 18: o32 Binary Interface and General Purpose Registers .....</b>	<b>785</b>
18.1: o32 ABI Parameters .....	785
18.2: Compiler Goals.....	786
18.3: Register Naming Conventions and Usage .....	786
18.4: Conventional Naming and General Purpose Register Usage .....	786
18.4.1: Common Name — AT (GPR \$1).....	788
18.4.2: Common Name — v0, v1 (GPR \$2, \$3).....	788
18.4.3: Common Name — a0 - a3 (GPR \$4 - \$7).....	788
18.4.4: Common Name — t0 - t9 (GPR \$8 - \$15, GPR \$24 - \$25 .....	788
18.4.5: Common Name — s0 - s8 (GPR \$16 - \$23, GPR \$30) .....	788
18.4.6: Common Name — k0, k1 (GPR \$26, \$27).....	788
18.4.7: Common Name — gp (GPR \$28) .....	789
18.4.7.1: Position Independent Code (PIC).....	789
18.4.7.2: Non-Position Independent Code (Non-PIC) .....	789
18.4.8: Common Name — sp (GPR #29) .....	789
18.4.9: Common Name — fp (GPR #30) .....	789
18.4.10: Common Name — ra (GPR #31).....	790
18.5: Floating Point Register Conventions .....	790
18.5.1: MIPS Floating Point Registers .....	790
18.5.2: Floating Point Register Software Use and Calling Conventions .....	790
18.6: Virtual Memory Layout Overview.....	791
18.6.1: Memory Regions and Object Code Naming Conventions.....	791
18.6.2: Simple Standalone Application .....	791
18.6.3: Linux Applications .....	793
18.6.4: PIC Code and the Global Offset Table.....	794
18.6.4.1: Global Offset Table (GOT) .....	794
18.6.4.2: Loading a PIC Application and its Libraries.....	795
18.6.4.3: Loading and Binding of Libraries .....	795
18.6.4.4: Dynamic (explicit) Loading of Libraries - dlopen() .....	795
18.6.4.5: PIC and GOT Constraints.....	795
18.7: Mapping Data Types into Memory.....	796
18.7.1: Sizes of C Data Types .....	796
18.7.2: Basic Type Memory Layout and Endianness.....	797
18.7.3: Memory Layout and Alignment of Structure and Array Types .....	797
18.7.4: Bit Fields in Structures .....	798
18.7.5: Alignment Rules .....	799
18.8: Calling Conventions.....	799
18.8.1: Stack Maintenance and Alignment.....	799
18.8.2: Registers and the Argument Structure.....	800
18.8.3: Returning Values from a Function.....	800
18.8.4: Calling Conventions Extended for Linux (“MIPS ABI”) PIC Code. ....	801
<b>Chapter 19: Programming Concepts .....</b>	<b>803</b>
19.1: Ordering and Synchronization .....	803
19.1.1: Consistency Model.....	803
19.1.2: LL/SC .....	803
19.1.3: Memory Barriers.....	803
19.1.4: CACHE and SYNCI Instructions .....	805

19.1.4.1: Cache Line Locking .....	805
19.1.4.2: Index Type and Optimized Routines .....	805
19.1.4.3: Completion .....	805
19.1.4.4: L2 CACHE Instructions.....	805
19.1.5: PREF Instructions .....	806
19.2: User Mode Programming .....	807
19.2.1: User Mode Accessible CP0 Registers .....	807
19.2.1.1: Programming the HWREna Register .....	807
19.2.1.2: Programming Example .....	808
19.2.2: Prefetching Data Using the pref and prefix Instructions.....	809
19.2.3: Using "SYNCl" When Writing Instructions.....	810
19.2.4: Integer Multiply and Divide.....	811
19.2.5: Tuning Software for the Pipeline .....	811
19.2.5.1: Cache Delays .....	812
19.2.6: Branch Misprediction Delays.....	812
19.2.7: Load Delayed by (Unrelated) Recent Store .....	812
19.2.8: Minimum Load-miss Penalty .....	813
19.2.9: Data Dependency Delays .....	813
19.2.10: Advice on Tuning Instruction Sequences (particularly DSP).....	814
19.2.11: Multiply/Divide Unit and Timings .....	814
19.3: Kernel Mode Programming.....	816
19.3.1: Hazard Barrier Instructions .....	816
19.3.1.1: Execution Hazards .....	816
19.3.1.2: Instruction Hazards.....	817
19.3.2: Enhanced Interrupt System.....	818
19.3.2.1: Traditional MIPS® Interrupt Signalling and Priority .....	819
19.3.2.2: VI Mode - Multiple Entry Points, Interrupt Signalling and Priority .....	819
19.3.3: External Interrupt Controller (EIC) Mode.....	819
19.4: Exception Entry Points .....	820
19.4.1: Summary of Exception Entry Points.....	821
19.5: Saving Power .....	822
<b>Chapter 20: C and C++ Efficient Programming Principles .....</b>	<b>823</b>
20.1: Basic Compiler Switches for Optimization.....	823
20.2: Data Alignment .....	824
20.2.1: Aligned Access Example.....	824
20.2.2: Unaligned Access Example .....	824
20.3: Dealing with Caches.....	826
20.3.1: Coherent Systems and Non-Coherent I/O .....	826
20.3.2: Data Prefetch .....	827
20.3.3: Mixing Cached and non-Cached Writes.....	827
20.3.4: Self-modifying Code.....	828
20.3.5: Minimizing Memory Transactions.....	829
20.4: General Porting Considerations .....	829
20.5: Migrating from ARM to MIPS.....	830
20.6: Helping the Compiler .....	830
20.7: O32 ABI Issues.....	831
20.7.1: Data sizes and alignment .....	831
20.7.2: Integer Calling Convention.....	831
20.7.3: Floating-Point Calling Convention.....	832
20.7.4: Large Structure as Return Value.....	832
20.7.5: Stack and Heap Overview.....	832
20.8: Customer Code Optimization Overview .....	832

20.8.1: Profile your Application .....	832
20.8.2: Optimize Routines.....	833
20.8.3: Intrinsics .....	833
20.9: Overview of Compiler Optimizations .....	833
20.9.1: List of Automatic Compiler Optimizations .....	833
20.9.1.1: Data flow Analysis Optimizations .....	833
20.9.1.2: Loop Optimizations.....	834
20.9.1.3: Code Generation Optimizations .....	834
20.9.2: List of Manually Chosen Compiler Optimizations.....	834
20.9.3: Compiler Optimizations for Code-Size .....	835
20.9.4: Optimized Libraries .....	835
20.9.4.1: Optimized C Library Routines.....	835
20.9.4.2: The MIPS DSP Library .....	835
20.9.4.3: GoFast Floating-Point Library.....	835
20.10: Linking to Libraries .....	835
20.10.1: Static Linking.....	836
20.10.2: Dynamic Linking.....	836
20.10.3: Position Independent Code.....	836
20.11: Thread Synchronization and Memory Ordering.....	836
20.11.1: Atomic Operations.....	836
20.11.2: Memory Barriers.....	837
20.11.3: Weak Memory Model .....	837
20.12: Programming Considerations for MIPS Virtual Processors .....	839
20.12.1: Use all Physical Cores .....	839
20.12.2: Sharing of Pipelines - Potentially Bad for Performance .....	839
20.12.3: Sharing of Caches - Good for Temporal and Spatial Locality .....	839
20.12.4: Producer-consumer Examples.....	839
20.13: Multi-threading Software - An Overview .....	840
20.13.1: Threading for Tasks .....	840
20.13.2: Fork Join Model.....	840
20.13.3: Thread Pool Model.....	840
20.13.4: PThreads.....	840
20.13.5: Thread-Safety and Reentrancy .....	840
20.13.6: SMP/MT Bugs and Issues.....	841
20.13.6.1: Race Condition .....	841
20.13.6.2: Priority Inversion.....	841
20.13.6.3: Deadlock.....	841
20.13.6.4: LiveLock .....	841
20.13.6.5: Cache thrashing .....	841
20.13.6.6: False Sharing .....	841
20.13.7: MIPS PDTrace and Performance Counters for SMP .....	842

**Chapter 21: Multithreading in the interAptiv Core..... 843**

21.1: MT ASE Definitions .....	843
21.2: Thread Context Resource Allocation in the interAptiv Core .....	844
21.3: MT ASE Instructions.....	845
21.3.1: Thread Creation .....	845
21.3.2: Multithreading Software Design Considerations and the FORK Instruction.....	845
21.3.2.1: No Implicit Context Copy .....	845
21.3.2.2: No Value Provided to Forking Thread .....	846
21.3.2.3: Absolute Virtual Thread Starting Address .....	846
21.3.3: Thread Overflow Exception.....	846
21.3.4: Thread Suspension Using the YIELD Instruction.....	847

21.3.5: Additional MT ASE Instructions.....	847
21.4: Multithreading CP0 Registers.....	848
21.4.1: Per-Core Multithreading Registers.....	848
21.4.2: Per-VPE Multithreading Registers.....	848
21.4.3: Per-TC Multithreading Registers.....	848
21.5: Thread Level Exception Processing.....	849
21.6: Fine-Grain Multithreading in the interAptiv Core.....	849
21.6.1: Dedicated Register Set.....	850
21.6.2: Automatic Fine-Grained Multithreading.....	850
21.7: Operating System Support.....	850
21.7.1: "Virtualization" of Threads and Hybrid Scheduling.....	850
21.7.2: Software Security.....	851
21.7.3: Manipulation of Dynamic Allocation Properties of TCs.....	851
21.7.4: Virtual Multiprocessor.....	852
21.7.5: Master / Slave VPE's.....	852
<b>Chapter 22: SMP Linux Programming.....</b>	<b>853</b>
22.1: Overview of Boot Sequence.....	854
22.2: Inter-Processor Interrupts.....	855
22.3: Clock Source and Events.....	855
22.4: Platform Setup.....	855
<b>Chapter 23: Instruction Set Overview.....</b>	<b>857</b>
23.1: CPU Instruction Formats.....	857
23.2: Load and Store Instructions.....	858
23.2.1: Scheduling a Load Delay Slot.....	858
23.2.2: Defining Access Types.....	858
23.3: Computational Instructions.....	859
23.3.1: Cycle Timing for Multiply and Divide Instructions.....	860
23.4: Jump and Branch Instructions.....	860
23.4.1: Overview of Jump Instructions.....	860
23.4.2: Overview of Branch Instructions.....	860
23.5: Control Instructions.....	861
23.6: Coprocessor Instructions.....	861
23.7: New Instructions for the interAptiv™ Core.....	861
23.8: Base Instruction Set for the interAptiv™ Core.....	863
<b>Chapter 24: MIPS16e Application-Specific Extension to the MIPS32® Instruction Set.....</b>	<b>871</b>
24.1: Instruction Bit Encoding.....	871
24.2: MIPS16e and MIPS16e2 Instruction Listing.....	874
24.3: MIPS16e2 Implementation Specific Instructions.....	878
<b>Chapter 25: Implementation-specific Instructions.....</b>	<b>884</b>
<b>Appendix A: References.....</b>	<b>918</b>
<b>Appendix B: Revision History.....</b>	<b>920</b>

## Architecture Overview

The interAptiv™ Multiprocessing System is a high performance multi-core microprocessor with best in class power efficiency for use in system-on-chip (SoC) applications. The interAptiv Multiprocessing System combines a multi-threading pipeline with a highly intelligent coherence manager to deliver best in class computational throughput and power efficiency. The device is fully configurable/synthesizable and can contain one to four MIPS32® interAptiv cores, system level coherence manager with L2 cache, optional coherent I/O port, and optional floating point unit.

The interAptiv Multiprocessing System is available in the following configurations. All of these configurations include a second generation Coherence Manager with integrated L2 cache (CM2).

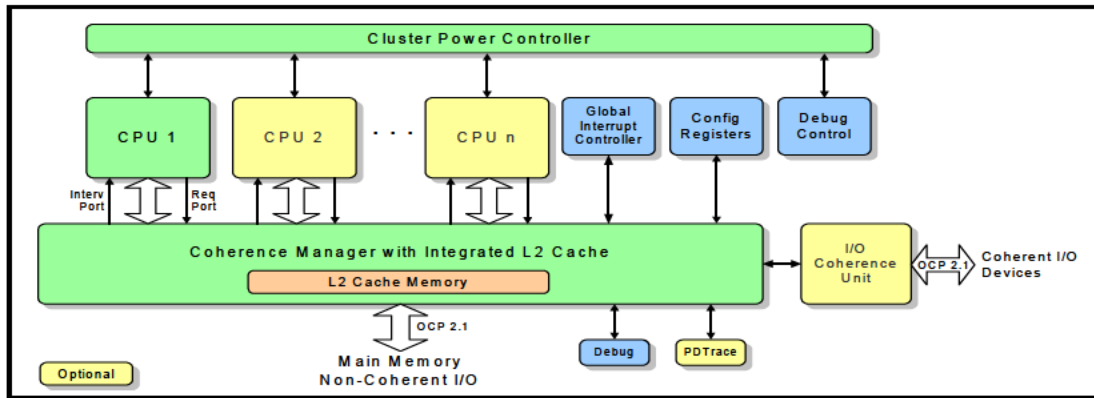
- Dual core configuration (one or two cores can be enabled)
- Quad core configuration (one, two, three, or four cores can be enabled)

The interAptiv Multiprocessing System contains the following logic blocks.

- interAptiv Cores (1 - 4)
- Coherence Manager (2nd generation) with integrated L2 cache (CM2)
- Optional Floating Point Unit per core (FPU)
- Cluster Power Controller (CPC)
- Optional Global Interrupt Controller (GIC)
- I/O Coherence Unit (IOCU)
- CM2 Global Configuration Registers (GCR)
- Multiprocessing System Debug Unit
- Optional PDTrace in-system trace debugger

Figure 1.1 shows a block diagram of the interAptiv Multiprocessing System.

**Figure 1.1 interAptiv™ Multiprocessing System Block Diagram**



In the interAptiv Multiprocessing system, multi-CPU coherence is handled in hardware by the Coherence Manager. The optional I/O Coherence Unit (IOCU) supports hardware I/O coherence by bridging a non-coherent OCP I/O interconnect to the Coherence Manager (CM2) and handling ordering requirements. The Global Interrupt Controller (GIC) handles the distribution of interrupts between and among the CPUs. Under software controlled power management, the Cluster Power Controller (CPC) can gate off the clocks and/or voltage supply to idle cores.

## 1.1 Features

The following subsections describe the features of the interAptiv Multiprocessing System. The features are divided into system level and core level.

### 1.1.1 System Level Features

- 1 - 4 coherent MIPS32 interAptiv cores
- Second generation Coherence Manager (CM2) providing L2 cache, I/O and interrupt coherence across all CPU cores
- Integrated 8-way set associative L2 cache controller supporting 128 KB to 8 MB cache sizes
- Supports 0 KB L2 cache option (no L2 cache)
- Supports 32 KB and 64 KB L2 cache options (no ECC)
- Programmable L1 data cache supporting the MESI coherence states and cache-to-cache data transfers
- Cluster Power Controller (CPC-standard) to shut down idle CPU cores
- Optional CPC-basic power controller
- Up to two hardware I/O coherence ports (IOCU) per system (optional)
- Speculative memory reads and out-of-order data return to reduce latency
- Clock ratio of 1:1 between Core, CM2, and L2 cache
- SOC system interface supports OCP version 2.1 protocol with 32-bit address and 64-bit or 256-bit data paths
- Power Control
  - Minimum frequency: 0 MHz
  - Software controlled power-down mode (triggered by WAIT instruction)



- Software-controlled clock divider
- Cluster-level dynamic clocking
- Cluster Power Controller (CPC) controlling shut down of idle CPU cores
- EJTAG Debug 5.0 port supporting multi-CPU debug
- MIPS PDtrace debug version 6.16 (optional)
  - PC, data address and data value tracing w/ trace compression
  - Includes features for correlation with CM trace
  - Support for on-chip and off-chip trace memory
  - Support for system-level trace
- User-defined global configuration registers
- Full scan design achieves test coverage in excess of 99% with optional memory BIST for internal SRAM arrays

### 1.1.2 Core Level Features

- Efficient pipeline with integer, floating point and optional CorExtend execution units shared amongst issue pipes.
- MIPS32 Release3 Instruction Set and Privileged Resource Architecture.
- Optional IEEE-754 compliant Floating Point Unit
- Instruction Fetch Unit (IFU) with up to 2 instructions fetched per cycle
- Programmable Memory Management Unit (MMU):
  - 16/32/48/64 dual-entry JTLB per VPE
  - 8-entry DTLB
  - 4 - 12 entry ITLB<sup>1</sup>
- Optional Memory Protection Unit (MPU)
- L1 Instruction and Data Caches
  - L1 instruction and data cache sizes of 4/8/16/32/64 KB, 4-way set associative
  - L1 MESI coherent data cache states
  - 32-byte cache line size
  - Virtually indexed, physically tagged
  - ECC and parity support on L1 data cache
  - Parity support on L1 instruction cache
- Data and Instruction Scratchpad RAM can be configured from 4 KB to 1 MB (optional).
- MIPS MT ASE
  - Support for 1 or 2 Virtual Processor's (VPE's) and 1 - 9 Thread Contexts (TC's)
  - Inter-Thread Communication (ITC) memory for efficient communication & data transfer
- Optional MIPS DSP-R2 ASE:
  - 3 additional pairs of accumulator registers
  - Fractional data types, Saturating arithmetic

---

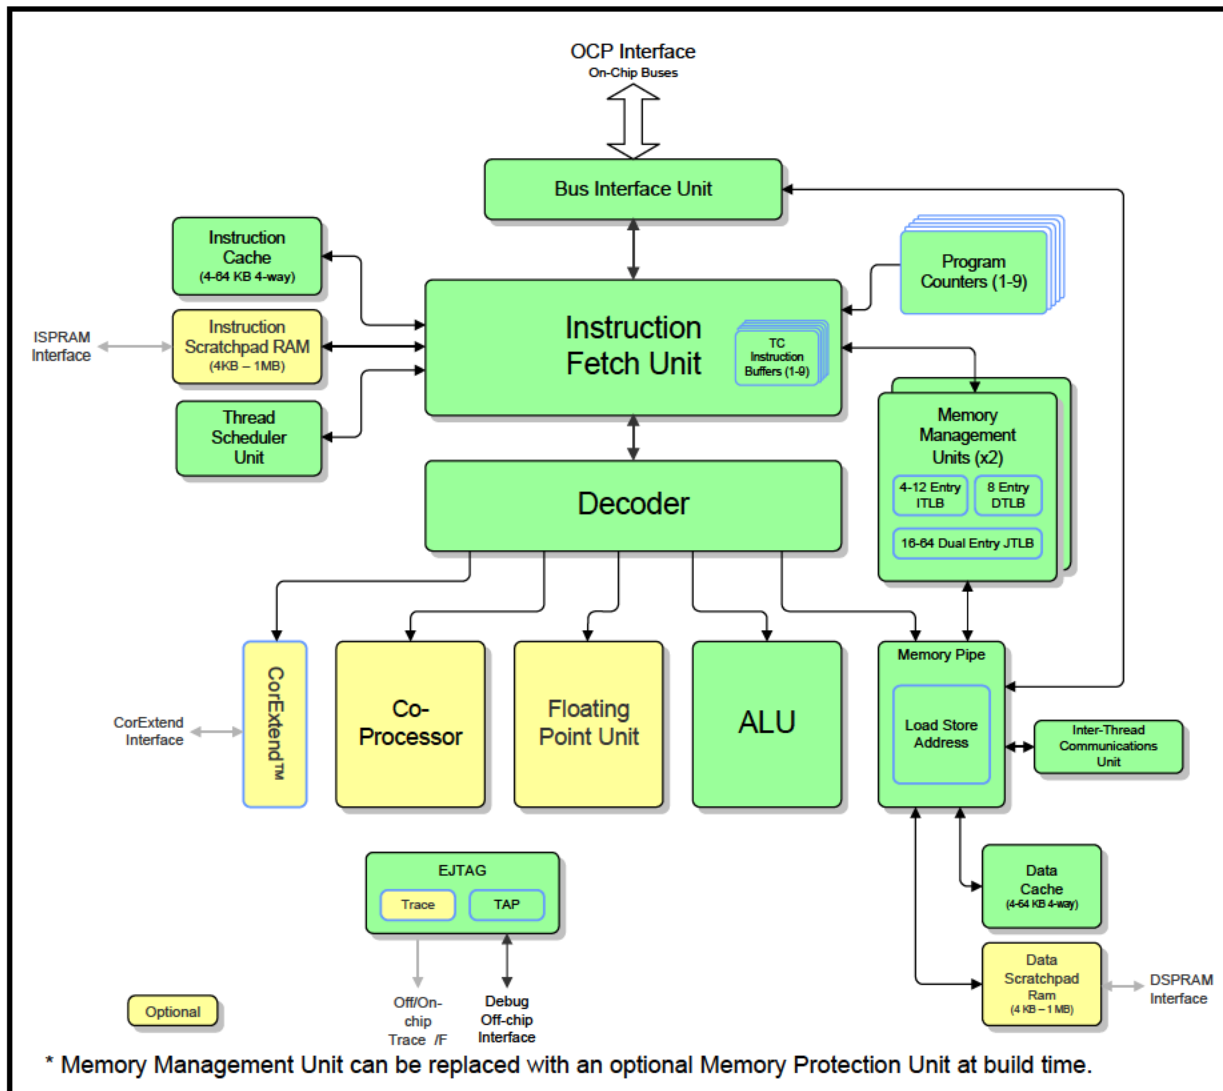
1. Three dedicated entries plus one entry per thread context (TC) up to nine threads.

- SIMD instructions operate on 2x16b or 4x8b simultaneously
- Write merging for uncached accelerated (UCA) operations
- Integrated integer Multiply/Divide Unit (MDU)
- CorExtend® MIPS32® compatible User Defined Instruction Set Extension allows user to define and add instructions to the core at build time
- Core Power Reduction Features
  - Power reduction by turning off core clock during outstanding bus requests
  - Power reduction by implementing intelligent way selection in the L1 instruction cache

## 1.2 interAptiv Core

[Figure 1.2](#) shows a block diagram of a single interAptiv core. The following subsections describe the logic blocks in this diagram.

Figure 1.2 interAptiv Core Block Diagram

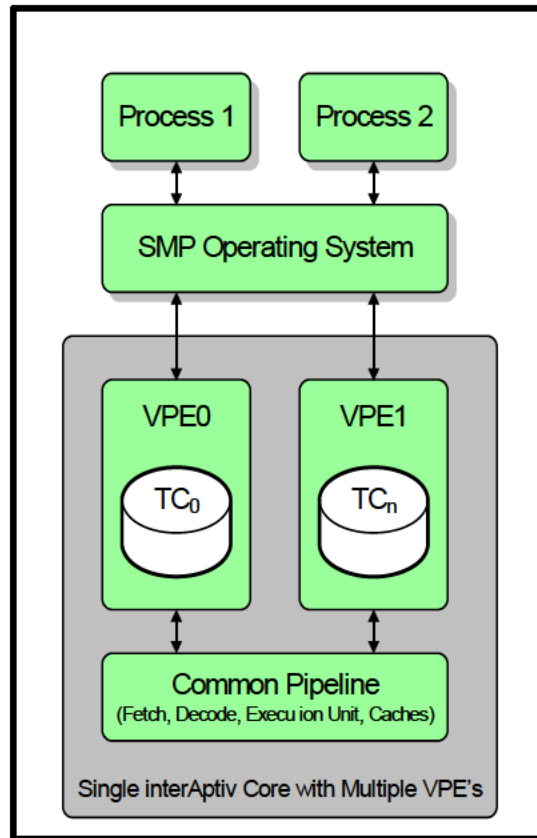


For more information on the interAptiv core in a multiprocessing environment, refer to [Section 1.3 “Multiprocessing System”](#)

### 1.2.1 MIPS Multi-Thread Technology

Building on the prior generation of MIPS multithreaded (MT) processors, the interAptiv core also implements the same multi-threaded architecture and supports the Application Specific Extensions (MT ASE) which are based on a two-layered framework involving Virtual Processing Elements (VPE's) and Thread Contexts (TCs). Each interAptiv core can support up to two VPE's. Each VPE includes a complete copy of the processor state as seen by the software system, so each VPE appears as a complete standalone processor to an SMP Linux operating system. For more fine-grained thread processing applications, each VPE is capable of supporting multiple TCs. The TCs share a common execution unit but each has its own program counter and program register files so that each can handle a thread from the software. The interAptiv core can support up to nine TCs allocated across two VPEs, optimized and partitioned at run time. [Figure 1.3](#) shows the relationship of the OS, VPEs, TCs, and the common hardware in the interAptiv core.

Figure 1.3 interAptiv™ Core with Two VPE's



### 1.2.2 Instruction Fetch Unit

This block is responsible for fetching instructions for all Thread Contexts (TCs). Each TC has an instruction buffer (IBF) that decouples the fetch unit from the execution unit. When executing instructions from multiple TCs, a portion of the IBF is used as a skid buffer. In this case, instructions are held in the IBF after being sent to the execution unit. This allows stalled instructions to be flushed from the execution pipeline without needing to be refetched.

In order to fetch instructions without intervention from the execution unit, the fetch unit contains branch prediction logic. A 512-entry Branch History Table (BHT) is used to predict the direction of branch instructions, a 4-entry Return Prediction Stack (RPS) holds the return address from the most recent subroutine calls. The link address is pushed onto the stack whenever a JAL, JALR, or BGEZAL instruction is seen. The address is popped when a JR instruction occurs.

### 1.2.3 Thread Scheduling Unit (TSU)

This unit is responsible for dispatching instructions from different Thread Contexts (TCs). A policy manager assigns priorities for each TC. The TSU determines which TCs are available and selects the highest priority one available.

### 1.2.4 Policy Manager

The policy manager is a configurable block. Simple round-robin or fixed priority policies can be selected during design. The interAptiv core includes a reference policy manager that implements a weighted round-robin algorithm for long-term distribution of execution bandwidth.

## 1.2.5 interAptiv L1 Caches

The interAptiv core contains L1 instruction and data caches as described in the following subsections.

### 1.2.5.1 Level 1 Instruction Cache

The Level-1 (L1) instruction cache is configurable at 4/8/16/32/64 KB in size and is organized as 4-way set associative. Up to four instruction cache misses can be outstanding. The instruction cache is virtually indexed and physically tagged to make the data access independent of virtual to physical address translation. Instruction cache tag and data access are staggered across 2 cycles, with up to 4 instructions fetched per cycle.

An instruction tag entry holds 20 or 21 bits of physical address, a valid bit, a lock bit, and an optional parity bit. The instruction data entry holds two instructions (64 bits), 6 bits of pre-decode information to speed the decode of branch and jump instructions, and 9 optional parity bits (one per data byte plus one more for the pre-decode information). There are four data entries for each tag entry. The tag and data entries exist for each way of the cache. The LRU replacement bits (6-bit) are shared among the 4 ways and are stored in a separate array.

The interAptiv core supports instruction-cache locking. Cache locking allows critical code segments to be locked into the cache on a “per-line” basis, enabling the system programmer to maximize the performance of the system cache.

The cache-locking function is always available on all instruction-cache entries. Entries can be marked as locked or unlocked on a per entry basis using the CACHE instruction.

### 1.2.5.2 Level 1 Data Cache

The Level 1 (L1) data cache is configurable at 4/8/16/32/64 in size. It is also organized as 4-way set associative. Data cache misses are non-blocking and up to nine misses may be outstanding. The data cache is virtually indexed and physically tagged to make the data access independent of virtual-to-physical address translation. To achieve the highest possible frequencies using commercially available SRAM generators, cache access and hit determination is spread across three pipeline stages, dedicating an entire cycle for the SRAM access.

A data cache tag entry holds 20 or 21 bits of physical address, a valid bit, a state bit, and an optional parity bit. The data entry holds 64 bits of data per way, with optional parity per byte. There are 4 data entries for each tag entry. The tag and data entries exist for each way of the cache.

The interAptiv core supports a data-cache locking mechanism identical to that used in the instruction cache. Critical data segments are locked into the cache on a “per-line” basis. The locked contents can be updated on a store hit, but are not selected for replacement on a cache miss. Locked lines do not participate in the coherence scheme so processes which lock lines into a particular cache should be locked to that processor and prevented from migrating.

The cache-locking function is always available on all data-cache entries. Entries can then be marked as locked or unlocked on a per entry basis using the CACHE instruction.

The data cache supports ECC dual-bit error detection and single-bit error correction.

### 1.2.5.3 Level 1 Cache Memory Configuration

As described above, the interAptiv core incorporates on-chip L1 instruction and data caches that are typically implemented from readily available single-port synchronous SRAMs and accessed in two cycles: one cycle for the actual SRAM read and another cycle for the tag comparison, hit determination, and way selection. The instruction and data

caches each have their own 64-bit data paths and can be accessed simultaneously. [Table 1.1](#) lists the interAptiv core instruction and data cache attributes.

**Table 1.1 interAptiv™ Core L1 Instruction and Data Cache Attributes**

Parameter	Instruction	Data
Size <sup>1</sup>	4, 8, 16, 32, or 64 KB	4, 8, 16, 32, or 64 KB
Organization	4-way set associative	4-way set associative
Line Size	32 Bytes	32 Bytes
Read Unit	64 bits	64 bits
Write Policies	N/A	coherent and non-coherent write-back with write allocate
Miss restart after transfer of	miss word	miss word
Cache Locking	per line	per line
Error Detection Mechanism	1. No parity on I or D 2. Parity on I, parity on D 3. Parity in I, ECC on D	

1. For Linux based applications, MIPS recommends a cache size of 64 KB, with a minimum size of 32 KB.

## 1.2.6 Memory Management Unit (MMU)

Each interAptiv VPE contains an optional Memory Management Unit (MMU) that is primarily responsible for converting virtual addresses to physical addresses and providing attribute information for different segments of memory. Each VPE contains a separate JTLB so that translations for each VPE are independent from the other VPE. The interAptiv core also can be configured with an optional Memory Protection Unit (MPU) as described in the following section. The MMU contains the following Translation Buffer Lookaside (TLB) types:

- 4 - 12 entry Instruction TLB (ITLB) with 4 KB or 1 MB pages
- 8-entry Data TLB (DTLB) with 4 KB or 1 MB pages
- Fully associative Joint TLB (JTLB) with 16, 32, 48, or 64 dual-entries

### 1.2.6.1 Instruction TLB (ITLB)

The ITLB is dedicated to performing translations for the instruction stream. The ITLB is a hybrid structure contains between 4 and 12 entries depending on the number of threads. There are 3 entries that are shared by all TCs, plus an additional entry dedicated to each TC up to nine. Refer to the MMU chapter for more information on the ITLB structure.

The ITLB maps 4 KB or 1 MB pages/subpages. For 4 KB or 1 MB pages, the entire page is mapped in the ITLB. If the main TLB page size is between 4 KB and 1 MB, only the current 4 KB subpage is mapped. Similarly, for page sizes larger than 1 MB, the current 1 MB subpage is mapped.

The ITLB is managed by hardware and is transparent to software. The larger JTLB structure is used as a backup structure for the ITLB. If a fetch address cannot be translated by the ITLB, the JTLB attempts to translate it in the following clock cycle or when available. If successful, the translation information is copied into the ITLB for future use.

### 1.2.6.2 Data TLB (DTLB)

The DTLB is an 8-entry, fully associative TLB dedicated to performing translations for loads and stores. All entries are shared by all TCs. Similar to the ITLB, the DTLB maps either 4 KB or 1 MB pages/subpages.

The DTLB is managed by hardware and is transparent to software. The larger JTLB structure is used as a backup structure for the DTLB. If a fetch address cannot be translated by the DTLB, the JTLB attempts to translate it in the following clock cycle or when available. If successful, the translation information is copied into the DTLB for future use.

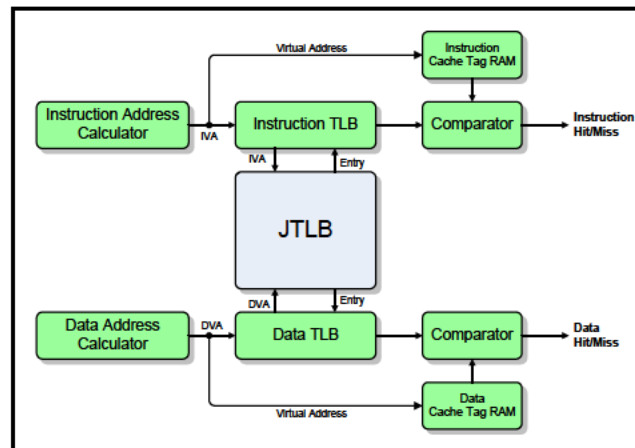
### 1.2.6.3 Translation Lookaside Buffer (TLB)

A TLB provides mapping and protection capability with per-page granularity. The interAptiv core implementation allows a wide range of page sizes to be simultaneously present.

The TLB contains a fully associative Joint TLB (JTLB). To enable higher clock speeds, two smaller micro-TLBs are also implemented: the Instruction TLB (ITLB) and the Data TLB (DTLB). When an instruction or data address is calculated, the virtual address is compared to the contents of the appropriate TLB. If the address is not found in either the ITLB or DTLB, the JTLB is accessed. If the entry is found in the JTLB, that entry is then written into the appropriate TLB. If the address is not found in the JTLB, a TLB exception is taken.

Figure 1.4 shows how the ITLB, DTLB, and JTLB are implemented in the interAptiv core.

**Figure 1.4 Address Translation During a Cache Access**



### 1.2.6.4 Joint TLB (JTLB)

The JTLB is a fully associative TLB cache containing 16, 32, 48, or 64-dual-entries mapping up to 128 virtual pages to their corresponding physical addresses. The address translation is performed by comparing the upper bits of the virtual address (along with the ASID) against each of the entries in the *tag* portion of the joint TLB structure.

The JTLB is organized as pairs of even and odd entries containing pages that range in size from 4 KB to 256 MB, in factors of four, into the 4 GB physical address space. The JTLB is organized in page pairs to minimize the overall size. Each *tag* entry corresponds to two data entries: an even page entry and an odd page entry. The highest order virtual address bit not participating in the tag comparison is used to determine which of the data entries is used. Since page size can vary on a page-pair basis, the determination of which address bits participate in the comparison and which bit is used to make the even-odd determination is decided dynamically during the TLB look-up.

## 1.2.7 Memory Protection Unit

The interAptiv core includes an optional Memory Protection Unit (MPU) for applications that do not require the functionality of a full MMU. The MPU controls whether read, write, or execute access is permitted for a given address and causes an exception if a unauthorized access is attempted. For more information on the MPU, refer to the Memory Protection Unit chapter.

## 1.2.8 Execution Pipelines

The execution unit implements a load/store architecture with single-cycle ALU operations (logical, shift, add, subtract) and an autonomous multiply/divide unit. Each TC contains thirty-one 32-bit general-purpose registers used for integer operations and address calculation. Additional sets of shadow register files can be added to be dedicated for interrupt and exception processing. The general-purpose register file consists of two read ports and one write port and is fully bypassed to minimize operation latency in the pipeline.

The execution unit includes:

- 32-bit adder used for calculating the data address
- Logic for verifying branch prediction
- Load aligner
- Bypass multiplexers used to avoid stalls when executing instructions streams where data producing instructions are followed closely by consumers of their results
- Leading Zero/One detect unit for implementing the CLZ and CLO instructions
- Arithmetic Logic Unit (ALU) for performing bit-wise logical operations
- Shifter and store aligner

The optional Floating Point Unit (FPU) contains two pipelines; one for arithmetic operations and one for data transfer operations. The arithmetic pipeline executes operations such as multiply, divide, and square root.

The data transfer pipeline executes floating point loads, stores, move operations, and register-to-register transfers between the FPU and the integer unit.

## 1.2.9 Instruction and Data Scratch Pad RAM

The interAptiv core allows for optional blocks of scratchpad RAM to be attached to the load/store and/or instruction units. These allow low-latency access to a fixed block of memory. The size of both the instruction scratch pad RAM (ISPRAM) and data scratch pad RAM (DSPRAM) can be configured from a range of 4 KB to 1 MB. These RAM's are used for the temporary storage of information and can be modified by the user at any time.

## 1.2.10 Inter-thread Communication Unit (ITU)

This block provides a mechanism for efficient communication between TCs. This block has a number of locations that can be used as mailboxes, FIFO mailboxes, and semaphores.



## 1.2.11 Bus Interface (BIU)

The Bus Interface Unit (BIU) controls the programmable 64-bit interface. The interface implements the Open Core Protocol (OCP). This implementation features 64-bit read and write data buses to efficiently transfer data to and from the L1 caches.

### 1.2.11.1 Write Buffer

The BIU contains a merging write buffer. The purpose of this buffer is to store and combine write transactions before issuing them to the external interface. The write buffer is organized as eight, 32-byte buffers. Each buffer can contain data from a single 32-byte aligned block of memory.

When performing uncached accelerated writes, the write buffer significantly reduces the number of write transactions on the external interface and reduces the amount of stalling in the core caused by the issuance of multiple writes in a short period of time.

The write buffer also holds eviction data for write-back lines. The load-store unit extracts dirty data from the cache and sends it to the BIU. In the BIU, the dirty data is gathered in the write buffer and sent out as a bursted write.

For uncached accelerated writes, the write buffer can gather multiple writes together and then perform a bursted write in order to increase the efficiency of the bus. Uncached accelerated gathering is supported for any size less than a doubleword.

Gathering of uncached accelerated stores can start on any arbitrary address and can be combined in any order within a cache line. Uncached accelerated stores that do not meet the conditions required to start gathering are treated like regular uncached stores.

### 1.2.11.2 SimpleBE Mode

To aid in attaching the interAptiv core to structures that cannot easily handle arbitrary byte-enable patterns, there is a mode that generates only “simple” byte enables. In this mode, only byte enables representing naturally aligned byte, halfword, word, and doubleword transactions will be generated.

In SimpleBE mode, the *SI\_SimpleBE* input pin only controls the byte enables generated by the interAptiv core(s). It has no effect on byte enables produced by the IOCU. To achieve the effect of setting *SI\_SimpleBE* to ‘one’ in systems with an IOCU, the I/O sub-system must only issue requests to the IOCU with naturally aligned byte enables.

When the *SI\_SimpleBE* input signal to the interAptiv core is asserted, hardware sets bit 21 of the *Config* register (*Config.SB*) to indicate the device is in simple byte enable mode.

## 1.2.12 System Control Coprocessor (CP0)

In the MIPS architecture, CP0 is responsible for the virtual-to-physical address translation and cache protocols, the exception control system, the processor’s diagnostic capability, the operating modes (kernel, user, supervisor, and debug), and whether interrupts are enabled or disabled. Configuration information, such as cache size and associativity, and the presence of features like MIPS16e or a floating point unit, are also available by accessing the CP0 registers.

CP0 also contains the logic used for identifying and managing exceptions. Exceptions can be caused by a variety of sources, including boundary cases in data, external events, or program errors. Most of CP0 is replicated per VPE. A small amount of state is replicated per TC and some are shared between the VPE’s (global registers).

For more information, refer to the CP0 chapter.

## 1.2.13 Interrupt Handling

The interAptiv core supports six hardware interrupts, two software interrupts, a timer interrupt, and a performance counter interrupt. These interrupts can be used in any of three interrupt modes, as defined by Release 3 of the MIPS32 Architecture:

- Interrupt compatibility mode, which acts identically to that in an implementation of Release 1 of the Architecture.
- Vectored Interrupt (VI) mode, which adds the ability to prioritize and vector interrupts to a handler dedicated to that interrupt. The presence of this mode is denoted by the *VInt* bit in the *Config3* register. This mode is architecturally optional. As it is always present on the interAptiv core, the *VInt* bit will always read 1.
- External Interrupt Controller (EIC) mode, which provides support for an external interrupt controller that handles prioritization and vectoring of interrupts. This mode is optional in the Release 3 architecture. The presence of this mode is denoted by the *VEIC* bit in the *Config3* register.

## 1.2.14 Modes of Operation

The interAptiv core supports four modes of operation:

- User mode, most often used for application programs.
- Supervisor mode provides an intermediate privilege level with access to the *ksseg* (kernel supervisor segment) address space.
- Kernel mode, typically used for handling exceptions and operating system kernel functions, including CP0 management and I/O device accesses.
- Debug mode is used during system bring-up and software development. Refer to [Section 1.2.17 “EJTAG Debug Support”](#) for more information on debug mode.

## 1.2.15 Floating Point Unit (FPU)

The optional Floating Point Unit (FPU) implements the MIPS64 ISA (Instruction Set Architecture) for floating-point computation. The FPU contains thirty-two 64-bit registers used for floating point operations. The implementation supports the ANSI/IEEE Standard 754 (IEEE Standard for Binary Floating-Point Arithmetic) for single and double precision data formats.

The FPU can be configured at build time to run at either the same or one-half the clock rate of the integer CPU. The FPU is not as deeply pipelined as the integer CPU so the maximum CPU frequency will only be attained with the FPU running at one-half the CPU frequency.

### 1.2.15.1 FPU Performance

The performance of the FPU is optimized for double-precision formats. Most instructions have a one cycle throughput. The FPU implements the MIPS64 multiply-add (MADD) and multiply-sub (MSUB) instructions with intermediate rounding after the multiply function. The result is guaranteed to be the same as executing a MUL and an ADD instruction separately, but the instruction latency, instruction fetch, dispatch bandwidth, and the total number of register accesses required are greatly improved.

IEEE denormalized input operands and results are supported by hardware for many instructions. IEEE denormalized output results are not supported by hardware in general, but a fast flush-to-zero mode is provided to optimize performance. The fast flush-to-zero mode is enabled through the *FCSR* register, and use of this non-standard mode is rec-

ommended for best performance when denormalized results are generated. This situation occurs most often in GPU driver code or multimedia CODECS handling real-time data streams.

Arithmetic instructions are always dispatched and graduated in order, but loads and stores can complete out-of-order. The integer core performs the data access for load/store operations and transfer data to and from the FPU using the CIU. Load data may arrive in the FPU out-of-order relative to program order. The exception model is ‘precise’ at all times.

The FPU implements a bypass mechanism that allows the result of an operation to be forwarded directly to the instruction that needs it without having to write the result to the FPU register and then read it back.

For more information, refer to the FPU chapter.

## 1.2.16 interAptiv Core Power Management

The interAptiv core offers several power management features, supporting low-power design, such as active power management and power-down modes of operation. The interAptiv core is a static design that supports slowing or halting the clocks to reduce system power consumption during idle periods.

For more information, refer to the Power Management chapter.

## 1.2.17 EJTAG Debug Support

The interAptiv core includes an Enhanced JTAG (EJTAG) block for use in software debugging of application and kernel code. For this purpose, in addition to standard user/supervisor/kernel modes of operation, the interAptiv core provides a Debug mode.

Debug mode is entered when a debug exception occurs (resulting from a hardware breakpoint, single-step exception, etc.) and continues until a debug exception return (DERET) instruction is executed. During this time, the processor executes the debug exception handler routine.

The EJTAG interface operates through the Test Access Port (TAP), a serial communication port used for transferring test data in and out of the interAptiv core. In addition to the standard JTAG instructions, special instructions defined in the EJTAG specification define which registers are selected and how they are used.

There are several types of simple hardware breakpoints defined in the EJTAG specification. These breakpoints stop the normal operation of the CPU and force the system into debug mode.

During synthesis, the interAptiv core can be configured to support the following breakpoint options:

- Zero instruction, zero data breakpoints
- Two instruction, one data breakpoints
- Four instruction, two data breakpoints

Instruction breaks occur on instruction fetch operations, and the break is set on the virtual address. Instruction breaks can also be made on the ASID value used by the MMU. A mask can be applied to the virtual address to set breakpoints on a range of instructions.

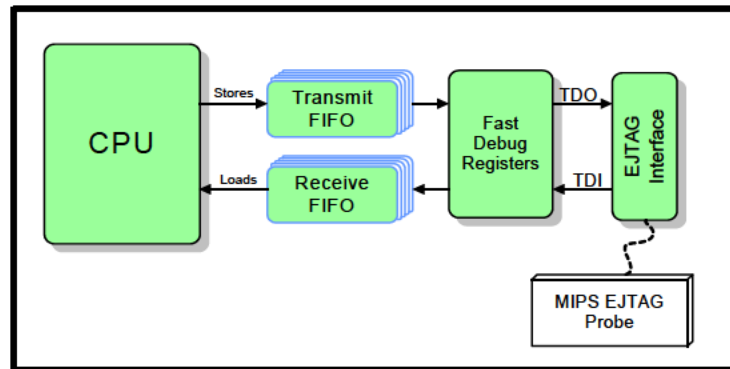
Data breakpoints occur on load and/or store transactions. Breakpoints are set on virtual address and address space identifier (ASID) values, similar to the Instruction breakpoint. Data breakpoints can also be set based on the value of the load/store operation. Finally, masks can be applied to the virtual address, ASID value, and the load/store value.

In debug mode, EJTAG can request that a ‘soft’ reset be masked. This request is signalled via the *EJ\_SRstE* pin. When this pin is deasserted, the system can choose to block some sources of soft reset. Hard resets, such as power-on reset or a reset switch, should not be blocked by this signal. This reset pin has no effect inside the core.

### 1.2.18 Fast Debug Channel

The interAptiv core includes the EJTAG Fast Debug Channel (FDC) as a mechanism for efficient bi-directional data transfer between the CPU and the debug probe. Data is transferred serially via the TAP interface. A pair of memory-mapped FIFOs buffer the data, isolating software running on the CPU from the actual data transfer. Software can configure the FDC block to generate an interrupt based on the FIFO occupancy or can poll the status.

Figure 1.5 Fast Debug Channel



### 1.2.19 PDtrace™

The interAptiv core includes optional PDtrace support for real-time tracing of instruction addresses, data addresses, data values, performance counters, and processor pipeline inefficiencies. The trace information is collected in an on-chip or off-chip memory, for post-capture processing by trace regeneration software. Software-only control of trace is possible in addition to probe-based control.

An optional on-chip trace memory may be configured in size from 256B to 8 MB; it is accessed either through load instructions or the existing EJTAG TAP interface, which requires no additional chip pins. Off-chip trace memory is accessed through a special trace probe and can be configured to use 4, 8, 16, or 64 data pins plus a clock.

### 1.2.20 MIPS16e2 Application Specific Extension

The interAptiv core includes support for the MIPS16e2 ASE. This ASE improves code density through the use of 16-bit encodings of many MIPS32 instructions plus some MIPS16e2 specific instructions. The MIPS16e2 ASE includes additional instructions not accessible in the previous generation to enhance code performance. Refer to the document entitled *MD01172, MIPS16e2 Application Specific Extensions Technical Reference Manual*, for a listing of new MIPS16e2 instructions.

In addition to the architecturally defined instructions in the MIPS16e2 manual, there are four new implementation specific instructions. Two of these instructions, COPYW and UCOPYW, are MIPS16e2 Application Specific Macro (ASMACRO) instructions and are described at the end of Chapter 24 of this manual. The other two instructions, SAVE and RESTORE, are MIPS32 CorExtend instructions that are described at the end of Chapter 25 of this manual.

## 1.2.21 CorExtend® Unit

The CorExtend unit is a custom block that allows the user to connect to the interAptiv core pipeline with access to all programmer-visible general purpose registers and accumulator state.

MIPS provides a template to define the operand format and latency for the new instruction(s) to be added. Each instruction may select up to 2 source GPRs and/or 1 Accumulator from a set of 32 GPRs and 4 accumulators. The instruction may have a destination of either a GPR, an accumulator, or a private state.

## 1.3 Multiprocessing System

The interAptiv Multiprocessing consists of the logic modules shown in [Figure 1.1](#). Each of these blocks is described throughout this section.

### 1.3.1 Cluster Power Controller (CPC)

Individual cores within the cluster can have their clock and/or power gated off when they are not in use. This gating is managed by the Cluster Power Controller (CPC). The CPC controller handles the power shutdown and ramp-up of all CPUs in the cluster. Any interAptiv CPU that supports power-gating features is managed by the CPC.

The CPC also organizes power-cycling of the CM2 dependent on the individual core status and shutdown policy. Reset and root-level clock gating of individual CPUs are considered part of this sequencing.

#### 1.3.1.1 Cluster Power Controller Reset Control

The reset input of the system resets the CPC. Reset sideband signals are required to qualify a reset as system cold, or warm start. Register setting determine the course of action:

- Remain in powered-down
- Go into clock-off mode
- Power-up and start execution

This prevents random power up of power domains before the CPC is properly initialized. In case of a system cold start, after reset is released, the CPC powers up the cores as directed in the CPC cold start configuration. If at least one CPU has been chosen to be powered up on system cold start, the CM2 is also powered up.

When supply rail conditions of power gated cores have reached a nominal level, the CPC will enable clocks and schedule reset sequences for those cores and the coherence manager.

At a warm start reset, the CPC brings all power domains into their cold start configuration. However, to ensure power integrity for all domains, the CPC ensures that domain isolation is raised before power is gated off. Domains that were previously powered and are configured to power up at cold start remain powered and go through a reset sequence.

Within a warm start reset, sideband signals are also used to qualify if coherence manager status registers and GIC watch dog timers are to be reset or remain unchanged. The CPC, after power up of any CPU, provides a test logic reset sequence per domain to initialize TAP and PDTrace logic.

Note that unused CPUs are not held in reset until released by writing into the configuration registers. Rather, unused CPUs remain powered down and are held isolated towards the rest of the cluster. If power gating is not selected for a

given implementation, unused CPUs are powered but receive no clock and remain isolated until activated by the CPC.

In addition to controlling the deassertion of the CPC reset signal, there are memory-mapped registers that can set the value for each CPU's *SI\_ExceptionBase* pins. This allows different boot vectors to be specified for each of the cores so they can execute unique code if required. Each of the cores will have a unique CPU number, so it is also possible to use the same boot vector and branch based on that.

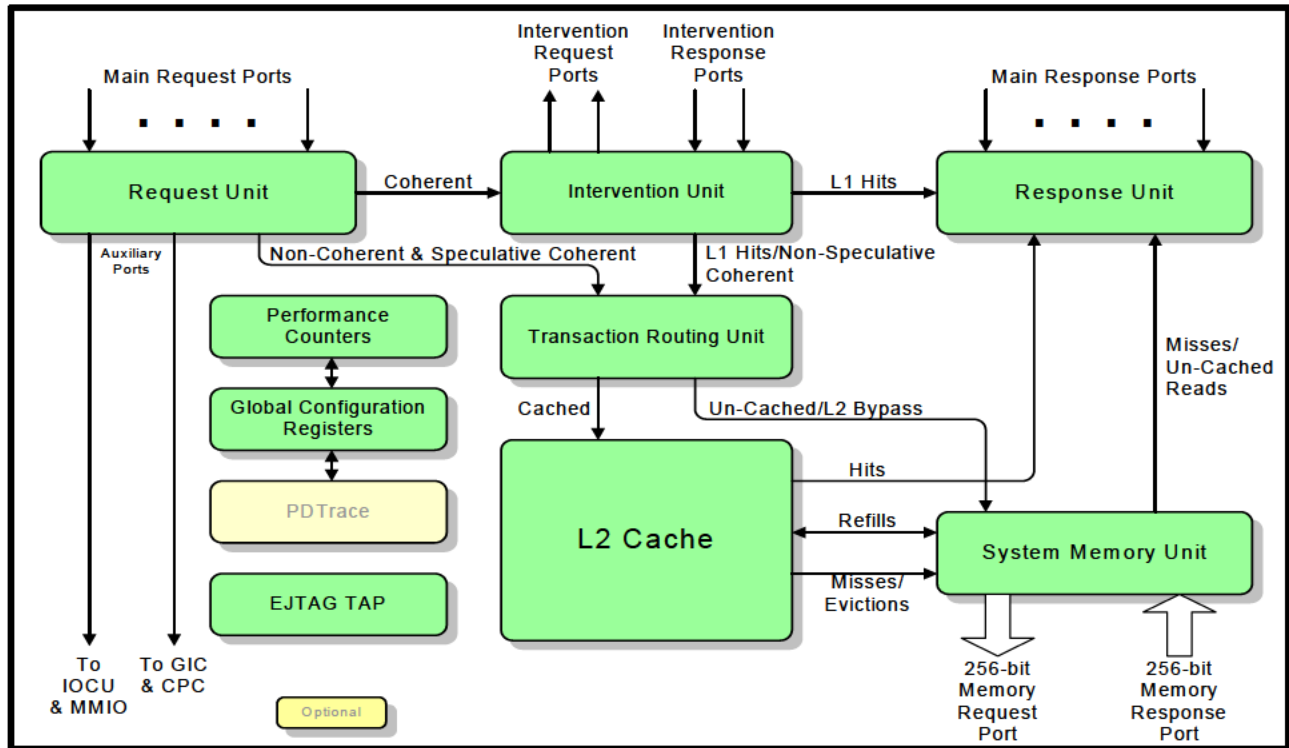
### 1.3.1.2 Optional CPC-basic Power Controller

The interAptiv includes an option for a basic version of the CPC. The type of CPC is selected via the GUI at IP configuration time. The pinout for the CPC-basic is different from the standard CPC, but both versions instantiate all signals and each design only uses the signals that it requires. In the CPC-basic, all functions are controlled by hardware. There is no software control via the CPC-basic registers. Refer to the CPC chapter for more information.

## 1.3.2 Coherence Manager (CM2)

The Coherence Manager with integrated L2 cache (CM2) is responsible for establishing the global ordering of requests and for collecting the intervention responses and sending the correct data back to the requester. A high-level view of the request/response flow through the CM2 is shown in Figure 1.6. Each of the blocks is described in more detail in the following subsections.

Figure 1.6 Coherence Manager with Integrated L2 Cache (CM2) Block Diagram



### 1.3.2.1 Request Unit (RQU)

The Request Unit (RQU) receives OCP bus transactions from multiple CPU cores and/or I/O ports, serializes the transactions and routes them to the Intervention Unit (IVU), Transaction Routing Unit (TRU), or an auxiliary port

used to access a configuration registers or memory-mapped IO. The routing is based on the transaction type, the transaction address, and the CM2's programmable address map.

### 1.3.2.2 Intervention Unit (IVU)

The Intervention Unit (IVU) interrogates the L1 caches by placing requests on the intervention OCP interfaces. Each processor responds with the state of the corresponding cache line. For most transactions, if a CPU core has the line in the MODIFIED or EXCLUSIVE states, it provides the data with its response. If the original request was a read, the IVU routes the data to the original requestor via the Response Unit (RSU). For the MESI protocol, intervention data may also be routed to the L2/Memory via the TRU (implicit writeback).

The IVU gathers the responses from each of the agents and manages the following actions:

- Speculative reads are resolved (confirmed or cancelled).
- Memory reads that are required because they were not speculative are issued to the Memory Interface Unit (MIU).
- Modified data returned from the CPU is sent to the MIU to be written back to memory.
- Data returned from the CPU is forwarded to the Response Unit (RSU) to be sent to the requester.
- The MESI state in which the line is installed by the requesting CPU is determined (the “install state”). If there are no other CPUs with the data, a Shared request is upgraded to Exclusive.

Each device updates its cache state for the intervention and responds when the state transition has completed. The previous state of the line is indicated in the response. If a read type intervention hits on a line that the CPU has in a Modified or Exclusive state, the CPU returns the cache line with its response. A cacheless device, such as the IOCU, does not require an intervention port. Note that the IVU is not included in non-coherent configurations, such as a single core without an IOCU.

### 1.3.2.3 System Memory Unit (SMU)

The System Memory Unit (SMU) provides the interface to the memory OCP port. For an L2 refill, the SMU reads the data from an internal buffer and issues the refill request to the L2 pipeline.

Note that the external interface may operate at a lower frequency than the Coherence Manager (CM2), and the external block may not be able to accept as many requests as multiple CPUs can generate, so some buffering of requests may be required.

### 1.3.2.4 Response Unit (RSU)

The RSU takes responses from the SMU, L2, IVU, or auxiliary port and places them on the appropriate OCP interface. Data from the L2 or SMU is buffered inside a buffer associated with each RSU port, which is an enhancement over the previous generation Coherence Manager.

When a coherent read receives an intervention hit in the MODIFIED or EXCLUSIVE state, the Intervention Unit (IVU) provides the data to the RSU. The RSU then returns the data to the requesting core.

### 1.3.2.5 Transaction Routing Unit

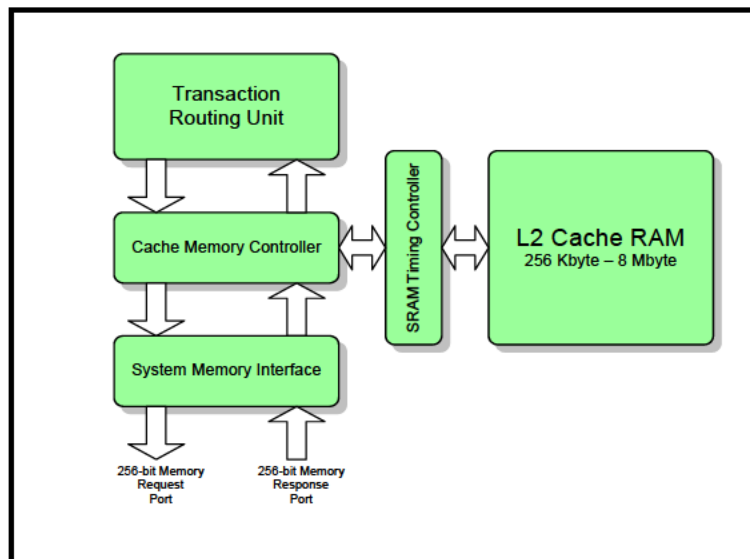
The Transaction Routing Unit (TRU) arbitrates between requests from the RQU and IVU, and routes requests to either the L2 or the SMU. The TRU also contains the request and intervention data buffers which are written directly from the RQU and IVU, respectively. The TRU reads the appropriate write buffer when it processes the corresponding write request.

### 1.3.2.6 Level 2 Cache

The unified L2 cache holds both instruction and data references and achieves high frequencies with low power while using commercially available SRAM generators.

Cache read misses are non-blocking; that is, the L2 can continue to process cache accesses while up to 15 misses are outstanding. The cache is physically indexed and physical tagged. Figure 1.7 shows a block diagram of the L2 cache.

Figure 1.7 L2 Cache Block Diagram



#### L2 Cache Features

- Supports write-back operation.
- Pseudo-LRU replacement algorithm
- Programmable wait state generator to accommodate a wide variety of SRAMs.
- Operates at same clock frequency as CPU.
- Cache line locking support
- Optional ECC support for resilience to soft errors
- Single bit error correction and 2 bit error detection support for Tag and Data arrays
- Single bit detection only for WS array
- Bypass mode
- Fully static design: minimum frequency is 0MHz
- Sleep mode
- Support for extensive use of fine-grained clock gating
- Optional memory BIST for internal SRAM arrays, with support for integrated (March C+, IFA-13) or custom BIST controller



## **L2 Cache Configuration**

The L2 cache in the CM2 can be configured as follows:

- 32 KBytes to 8 MBytes
- Supports 0 KB L2 cache option
- 32- or 64-byte line size
- 8 ways
- 512 to 32768 sets per way (in powers of two)

## **L2 Pipeline Tasks**

The L2 pipeline manages the flow of data to and from the L2 cache. The L2 pipeline performs the following tasks:

- Accesses the tags and data RAMs located in the memory block (MEM).
- Returns data to the RSU for cache hits.
- Issues L2 miss requests.
- Issues L2 write and eviction requests.
- Returns L2 write data to the SMU. The SMU issues refill requests to the L2 for installation of data for L2 allocations

### **1.3.2.7 CM2 Configuration Registers**

The Registers block (GCR) contains the control and status registers for the CM2. It also contains the Trace Funnel, EJTAG TAP state machine, and other multi-core features.

The configuration registers in the CM2 allow software to configure and control various aspects of the operation of the CM2. Some of the control options include:

- *Address map*: the base address for the GCR and GIC address ranges can be specified. An additional four address ranges can be defined as well. These control whether non-coherent requests go to memory or to memory-mapped I/O. A default can also be selected for addresses that do not fall within any range.
- *Error reporting and control*: Logs information about errors detected by the CM2 and controls how errors are handled (ignored, interrupt, etc.).
- *Control Options*: Various features of the CM2 can be disabled or configured. Examples of this are disabling speculative reads and preventing read/shared requests from being upgraded to exclusive.

### **1.3.2.8 PDTrace Unit**

The CM2 PDTrace Unit (PDT) is an optional unit used to collect, pack and send out CM2 debug information.

### **1.3.2.9 Performance Counter Unit**

The CM Performance Counter Unit (PERF) implements the performance counter logic.

### **1.3.2.10 Coherence Manager Performance**

The CM2 has a number of high performance features:

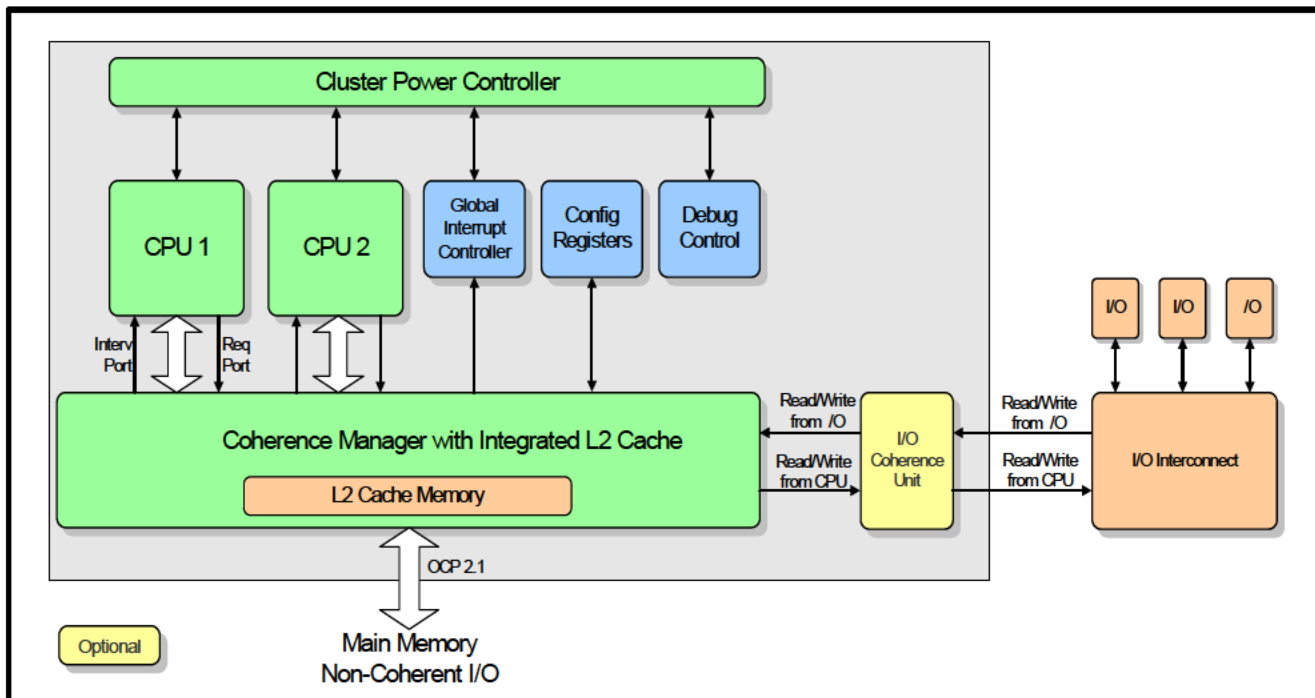
- 256-bit wide internal data paths throughout the CM2
- 64, 128, or 256-bit wide system OCP interface
- Cache to Cache transfers: If a read request hits in another L1 cache in the EXCLUSIVE or MODIFIED state, it will return the data to the CM and it will be forwarded to the requesting CPU, thus reducing latency on the miss.
- Speculative Reads: Coherent read requests are forwarded to the memory interface before they are looked up in the other caches. This is speculating that the cache line will not be found in another CPU's L1 cache. If another cache was able to provide the data, the memory request is not needed, and the CM2 cancels the speculative request—dropping the request if it has not been issued, or dropping the memory response if it has.

### 1.3.3 I/O Coherence Unit (IOCU)

Optional support for hardware I/O coherence is provided by the I/O Coherence Unit (IOCU), which maintains I/O coherence of the caches in all coherent CPUs in the cluster.

The IOCU acts as an interface block between the Coherence Manager (CM2) and I/O devices. Reads and writes from I/O devices may access the L1 and L2 caches by passing through the IOCU and the CM2. Each request from an I/O device may be marked as coherent, non-coherent cached, or uncached. Coherent requests access the L1 and L2 caches. Non-coherent cached requests access only the L2 cache. Uncached requests bypass both the L1 and L2 caches and are routed to main memory. An example system topology is shown in [Figure 1.8](#).

**Figure 1.8 Role of the IOCU in a Two-Core Multiprocessing System**



The IOCU also provides a legacy (without coherent extensions) OCP slave interface to the I/O interconnect for I/O devices to read and write system memory. The reference design also includes an OCP Master port to the I/O interconnect that allows the CPUs to access registers and memory on the I/O devices.

The reference IOCU design provides several features for easier integration:

- A user-defined mapping unit can define cache attributes for each request—coherent or not, cacheable (in L2) or not, and L2 allocation policy.
- Supports incremental bursts up to 16 beats (128 bits) on I/O side. These requests are split into cache-line- sized requests on the CM2 side.
- Ensures proper ordering of responses for the split requests and tagged requests.

In addition, the IOCU contains the following features used to enforce transaction ordering.

- Set-aside buffer: This buffer can delay read responses from the I/O device until previous writes have completed.
- Writes are issued to the CM2 in the order they were received.
- The CM2 provides an acknowledge (ACK) signal to the IOCU when writes are "visible" (guaranteed that a subsequent CPU read will receive that data).
  - Non-coherent write is acknowledged after serialization
  - Coherent write is acknowledged after intervention complete on all CPUs
- The IOCU can be configured to treat incoming writes as non-posted and provide a write ACK when they become visible.

When I/O devices access the same memory that is accessed by the processor cores, care must be taken to account for the caches. When an I/O device is reading memory, dirty data in the caches means that main memory may not contain the latest data, and a read directly from memory can receive stale data. When writing main memory, data in the caches becomes stale—the cores can read the stale value and potentially write it back to memory, overwriting the more recent I/O data.

Taking care of these problems can be handled by hardware, software, or a combination of both.

### 1.3.3.1 Software I/O Coherence

For cases where system redesign to accommodate hardware I/O coherence is not feasible, the CPUs and Coherence Manager provide support for an efficient software-managed I/O coherence. This support is through the globalization of hit-type CACHE instructions.

When a coherent address is used for the CACHE operations, the CPU makes a corresponding coherent request. The CM2 sends interventions for the request to all of the CPUs, allowing all of the L1 caches to be maintained together. The basic software coherence routines developed for single CPU systems can be reused with minimal modifications.

In software managed I/O coherence, software running on the CPU performs any cache operations that are required in accordance with I/O memory accesses. This may include pushing dirty data out of the caches before an I/O read and invalidating stale data after an I/O write. The software I/O coherence code can run on one of the cores and ensure that the appropriate action is taken in all of the caches in the Cluster.

Previous uniprocessor cores from MIPS Technologies have not included support for hardware I/O coherence, and systems based on those cores have relied on software coherence. Generally, the same coherence routines will work on the multi-CPU system.

### 1.3.3.2 Hardware I/O Coherence

For hardware I/O coherence, the coherence features on the CPU are used to ensure that I/O requests are handled properly. Requests from I/O devices go to the I/O Coherence Unit (IOCU) and then to a request port of the Coherence Manager. Requests that are marked coherent will generate interventions to the cores. I/O read requests can obtain any

dirty data directly from a data cache that has it in the M state. I/O write requests will invalidate the line in any data caches that have copies of it. Coherent requests access the in-line L2 cache from the CM, so they will automatically be coherent with the L2 cache.

Note that I/O interventions do not affect the instruction cache. The instruction cache cannot contain dirty data, so I/O reads are not a problem. However, if the I/O device is writing addresses that may reside in the instruction cache, software coherence must be used to invalidate the stale cache data.

In addition, even if hardware I/O coherence is present, there may be a need for software to explicitly maintain coherence. Examples of this are for systems configured without I/O coherence, devices that are not connected to the coherent port, or devices that directly access memory non-coherently. Initially, with the non-coherent I-Cache, this will also be needed to maintain I-Cache coherence with I/O traffic and data operations.

### 1.3.4 Global Interrupt Controller

An optional Global Interrupt Controller handles the distribution of interrupts between and among the VPE's in the system. The GIC is selected at IP configuration time as a GUI menu choice. The customer can opt to include the GIC, or to not include it in the build.

This block has the following features:

- Software interface through relocatable memory-mapped address range.
- Configurable number of system interrupts - from 8 to 256 in multiples of 8.
- Support for different interrupt types:
  - Level-sensitive: active high or low.
  - Edge-sensitive: positive, negative, or double-edge-sensitive.
- Ability to mask and control routing of interrupts to a particular CPU.
- Support for NMI routing.
- Standardized mechanism for sending inter-processor interrupts.

### 1.3.5 Global Configuration Registers (GCR)

The Global Configuration Registers (GCR) are a set of memory-mapped registers that are used to configure and control various aspects of the Coherence Manager and the coherence scheme.

#### 1.3.5.1 Reset Control

The reset input of the system resets the Cluster Power Controller (CPC). Reset sideband signals are required to qualify a reset as system cold, or warm start. Register setting determine the course of action:

- Remain in powered-down
- Go into clock-off mode
- Power-up and start execution

This prevents random power up of power domains before the CPC is properly initialized. In case of a system cold start, after reset is released, the CPC powers up the cores as directed in the CPC cold start configuration. If at least one core has been chosen to be powered up on system cold start, the CM2 is also powered up.

When supply rail conditions of power gated CPUs have reached a nominal level, the CPC will enable clocks and schedule reset sequences for those CPUs and the coherence manager.

At a warm start reset, the CPC brings all power domains into their cold start configuration. However, to ensure power integrity for all domains, the CPC ensures that domain isolation is raised before power is gated off. Domains that were previously powered and are configured to power up at cold start remain powered and go through a reset sequence.

Within a warm start reset, sideband signals are also used to qualify if coherence manager status registers and GIC watch dog timers are to be reset or remain unchanged. The CPC, after power up of any CPU, provides a test logic reset sequence per domain to initialize TAP and PDTrace logic.

Note that unused CPUs are not held in reset until released by writing into the configuration registers. Rather, unused CPUs remain powered down and are held isolated towards the rest of the cluster. If power gating is not selected for a given implementation, unused CPUs are powered but receive no clock and remain isolated until activated by the CPC.

In addition to controlling the deassertion of the CPC reset signal, there are memory-mapped registers that can set the value for each CPU's *SI\_ExceptionBase* pins. This allows different boot vectors to be specified for each of the cores so they can execute unique code if required. Each of the cores will have a unique CPU number, so it is also possible to use the same boot vector and branch based on that.

### 1.3.5.2 Inter-CPU Debug Breaks

The CPS includes registers that enable cooperative debugging across all CPUs. Each core features an *EJ\_DebugM* output that indicates it has entered debug mode (possibly through a debug breakpoint). Registers are defined that allow CPUs to be placed into debug groups such that whenever one CPU within the group enters debug mode, a debug interrupt is sent to all CPUs within the group, causing them to also enter debug mode and stop executing non-debug mode instructions.

### 1.3.5.3 CM2 Control Registers

Control registers in the CM2 allow software to configure and control various aspects of the operation of the CM2. Some of the control options include:

- *Address map*: the base address for the GCR and GIC address ranges can be specified. An additional four address ranges can be defined as well. These control whether non-coherent requests go to memory or to memory-mapped I/O. A default can also be selected for addresses that do not fall within any range.
- *Error reporting and control*: Logs information about errors detected by the CM2 and controls how errors are handled (ignored, interrupt, etc.).
- *Control Options*: Various features of the CM2 can be disabled or configured. Examples of this are disabling speculative reads and preventing ReadShared requests from being upgraded to Exclusive.

## 1.3.6 Clocking Options

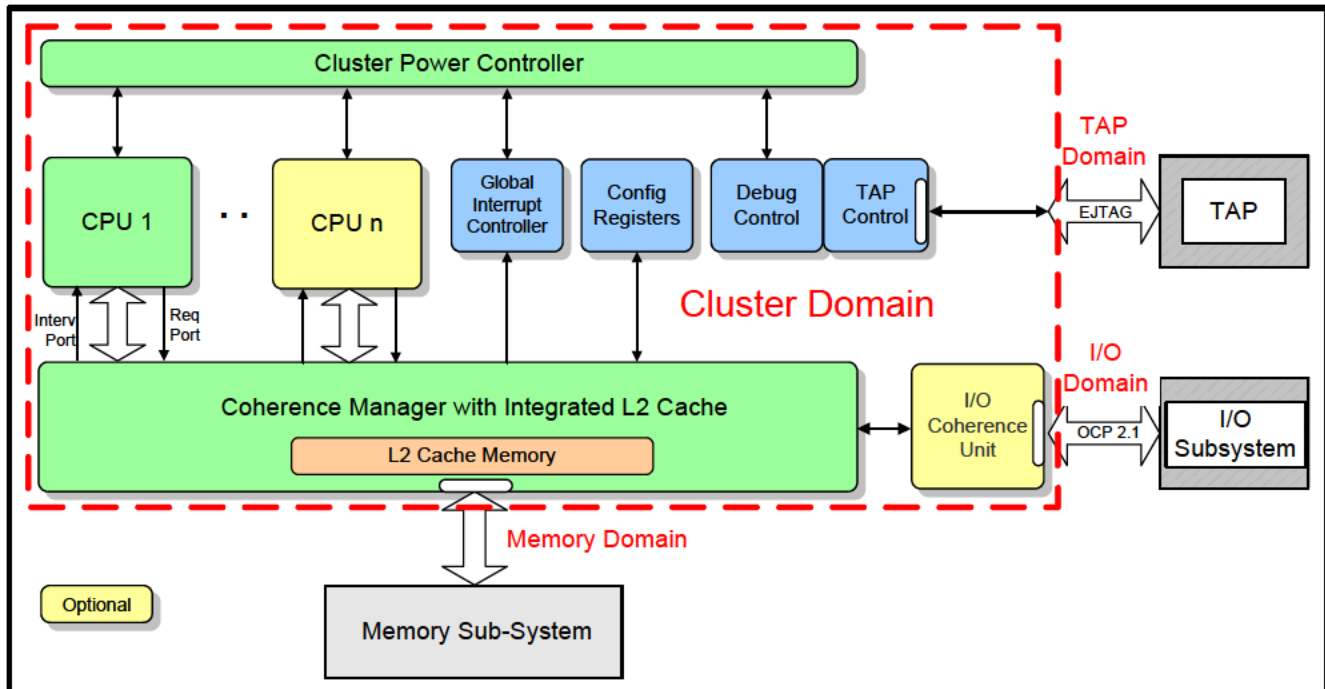
The interAptiv core has the following clock domains:

- Cluster domain — This is the main clock domain, and includes all interAptiv cores (including optional FP2) and the CM2 (including Coherence Manager, Global Interrupt Controller, Cluster Power Controller, trace funnel, IOCU, and L2 cache).
- System Domain - The OCP port connecting to the SOC and the rest of the memory subsystem may operate at a ratio of the cluster domain. Supported ratios are 1:1, 1:1.5, 1:2, 1:2.5, 1:3, 1:3.5, 1:4, 1:5, and 1:10.

- TAP domain - This is a low-speed clock domain for the EJTAG TAP controller, controlled by the *EJ\_TCK* pin. It is asynchronous to *SI\_ClkIn*.
- IO Domain - This is the OCP port connecting the IOCU to the I/O Subsystem. This clock may operate at a ratio of the CM2 domain. Supported ratios are the same as the system domain.

Figure 1.9 shows a diagram with the four clock domains.

**Figure 1.9 interAptiv Multiprocessing System Clocking Domains**



### 1.3.7 Design For Test (DFT) Features

The interAptiv core provides the following tests for determining the integrity of the core. For more information, refer to the Test and Debug chapter.

#### 1.3.7.1 Internal Scan

The interAptiv core supports full mux-based scan for maximum test coverage, with a configurable number of scan chains. ATPG test coverage can exceed 99%, depending on standard cell libraries and configuration options.

#### 1.3.7.2 Memory BIST

The interAptiv core provides an integrated memory BIST solution for testing of all internal SRAMs. These BIST controllers can be configured to utilize the March C+ or IFA-13 algorithms.

Memory BIST can also be inserted with a CAD tool or other user-specified method. Wrapper modules and signal buses of configurable width are provided within the core to facilitate this approach.

## 1.4 Build-Time Configuration Options

The interAptiv Multiprocessing System allows a number of features to be customized based on the intended application. Refer to the interAptiv Data Sheet for more information on the configuration options.





## CP0 Registers

The interAptiv Multiprocessing System Control Coprocessor (CP0) provides the register interface to the interAptiv core and supports memory management, address translation, exception handling, and other privileged operations. Each CP0 register has a unique number that identifies it, referred to as its *register number*. A separate *select number* is used to differentiate additional registers within the *register number*. For example, as shown in the table below, there are eight configuration registers with register number 16. If the *select number* is omitted, it is zero.

This chapter contains the following sections:

- [Section 2.1 “CP0 Register Summary”](#)
- [Section 2.2 “CP0 Register Descriptions”](#)

### 2.1 CP0 Register Summary

The following two subsections show the CP0 register set grouped by function and grouped by number.

#### 2.1.1 CP0 Registers Grouped by Function

The following table lists the CP0 registers in functional order and gives a brief description. Additionally, the table shows whether the register is implemented once per processor, once per VPE, once per TC, or once per system. The individual registers are described throughout this chapter.

**Table 2.1 CP0 Registers Grouped by Function**

Category	Register Name	Register Number	Register Select	Register is per:			Location in Document
				VPE	TC	Core	
Configuration and Status	Config	16	0	X			<a href="#">Section 2.2.1.1 on page 61</a>
	Config1	16	1	X			<a href="#">Section 2.2.1.2 on page 64</a>
	Config2	16	2	X			<a href="#">Section 2.2.1.3 on page 66</a>
	Config3	16	3	X			<a href="#">Section 2.2.1.4 on page 69</a>
	Config4	16	4	X			<a href="#">Section 2.2.1.5 on page 70</a>
	Config5	16	5	X			<a href="#">Section 2.2.1.6 on page 71</a>
	Config7	16	7	X			<a href="#">Section 2.2.1.7 on page 73</a>
	PRId	15	0	X			<a href="#">Section 2.2.1.8 on page 75</a>
	EBase	15	1	X			<a href="#">Section 2.2.1.9 on page 76</a>
	Status	12	0	X		X <sup>1</sup>	<a href="#">Section 2.2.1.10 on page 77</a>
	IntCtl	12	1	X			<a href="#">Section 2.2.1.11 on page 82</a>

**Table 2.1 CP0 Registers Grouped by Function (continued)**

Category	Register Name	Register Number	Register Select	Register is per:			Location in Document
				VPE	TC	Core	
TLB Management	Index	0	0	X			<a href="#">Section 2.2.2.1 on page 85</a>
	Random	1	0	X			<a href="#">Section 2.2.2.2 on page 86</a>
	EntryLo0	2	0	X			<a href="#">Section 2.2.2.3 on page 87</a>
	EntryLo1	3	0	X			
	EntryHi	10	0	X	X <sup>2</sup>		<a href="#">Section 2.2.2.4 on page 88</a>
	Context	4	0	X			<a href="#">Section 2.2.2.5 on page 90</a>
	ContextConfig	4	1	X			<a href="#">Section 2.2.2.6 on page 90</a>
	PageMask	5	0	X			<a href="#">Section 2.2.2.7 on page 91</a>
	PageGrain	5	1	X			<a href="#">Section 2.2.2.8 on page 92</a>
	Wired	6	0	X			<a href="#">Section 2.2.2.9 on page 93</a>
	BadVAddr	8	0	X			<a href="#">Section 2.2.2.10 on page 94</a>
Memory Segmentation	SegCtl0	5	2	X			<a href="#">Section 2.2.3.1 on page 95</a>
	SegCtl1	5	3	X			<a href="#">Section 2.2.3.2 on page 96</a>
	SegCtl2	5	4	X			<a href="#">Section 2.2.3.3 on page 97</a>
Exception Control	Cause	13	0	X			<a href="#">Section 2.2.4.1 on page 100</a>
	EPC	14	0	X			<a href="#">Section 2.2.4.2 on page 103</a>
	ErrorEPC	30	0	X			<a href="#">Section 2.2.4.3 on page 104</a>
Timer	Count	9	0	X			<a href="#">Section 2.2.5.1 on page 105</a>
	Compare	11	0	X			<a href="#">Section 2.2.5.2 on page 106</a>
Cache Management	ITagLo	28	0	X			<a href="#">Section 2.2.6.1 on page 107</a>
	IDataLo	28	1	X			<a href="#">Section 2.2.6.2 on page 109</a>
	IDataHi	29	1	X			<a href="#">Section 2.2.6.3 on page 109</a>
	DTagLo	28	2	X			<a href="#">Section 2.2.6.4 on page 110</a>
	DTagHi	29	2	X			<a href="#">Section 2.2.6.5 on page 113</a>
	DDataLo	28	3	X			<a href="#">Section 2.2.6.6 on page 114</a>
	L23TagLo	28	4	X			<a href="#">Section 2.2.6.7 on page 114</a>
	L23DataLo	28	5	X			<a href="#">Section 2.2.6.8 on page 115</a>
	L23DataHi	29	5	X			<a href="#">Section 2.2.6.9 on page 116</a>
	ErrCtl	26	0	X			<a href="#">Section 2.2.6.10 on page 116</a>
	CacheErr	27	0	X			<a href="#">Section 2.2.6.11 on page 119</a>

**Table 2.1 CP0 Registers Grouped by Function (continued)**

Category	Register Name	Register Number	Register Select	Register is per:			Location in Document
				VPE	TC	Core	
Thread Context	TCStatus	2	1		X		<a href="#">Section 2.2.7.1 on page 127</a>
	TCBind	2	2		X		<a href="#">Section 2.2.7.2 on page 129</a>
	TCRestart	2	3		X		<a href="#">Section 2.2.7.3 on page 130</a>
	TCHalt	2	4		X		<a href="#">Section 2.2.7.4 on page 130</a>
	TCContext	2	5		X		<a href="#">Section 2.2.7.5 on page 131</a>
	TCSchedule	2	6		X		<a href="#">Section 2.2.7.6 on page 131</a>
	TCscheFBack	2	7		X		<a href="#">Section 2.2.7.7 on page 133</a>
	TCOpt	3	7		X		<a href="#">Section 2.2.7.8 on page 134</a>
	SRSCnf0	6	1	X			<a href="#">Section 2.2.7.9 on page 134</a>
	SRSCnf1	6	2	X			<a href="#">Section 2.2.7.10 on page 135</a>
	SRSCnf2	6	3	X			
	SRSCnf3	6	4	X			
	SRSCnf4	6	5	X			
	SRSCtl	12	2	X			<a href="#">Section 2.2.7.11 on page 135</a>
	SRSMAP	12	3	X			<a href="#">Section 2.2.7.12 on page 137</a>
VPE Management	VPEControl	1	1	X			<a href="#">Section 2.2.8.1 on page 138</a>
	VPEConf0	1	2	X			<a href="#">Section 2.2.8.2 on page 139</a>
	VPEConf1	1	3	X			<a href="#">Section 2.2.8.3 on page 140</a>
	VPESchedule	1	5	X			<a href="#">Section 2.2.8.4 on page 141</a>
	VPEScheFBack	1	6	X			<a href="#">Section 2.2.8.5 on page 142</a>
	VPEOpt	1	7	X			<a href="#">Section 2.2.8.6 on page 143</a>
	MVPCControl	0	1			X	<a href="#">Section 2.2.8.7 on page 144</a>
	MVPCnf0	0	2			X	<a href="#">Section 2.2.8.8 on page 144</a>
	MVPCnf1	0	3			X	<a href="#">Section 2.2.8.9 on page 145</a>
Performance Monitoring	PerfCtl0	25	0		X		<a href="#">Section 2.2.9.1 on page 146</a>
	PerfCtl1	25	2		X		
	PerfCnt0	25	1		X		<a href="#">Section 2.2.9.2 on page 156</a>
	PerfCnt1	25	3		X		

**Table 2.1 CP0 Registers Grouped by Function (continued)**

Category	Register Name	Register Number	Register Select	Register is per:			Location in Document
				VPE	TC	Core	
Debug and Trace Registers	Debug	23	0	X	X		<a href="#">Section 2.2.10.1 on page 157</a>
	DEPC	24	0	X			<a href="#">Section 2.2.10.2 on page 160</a>
	DESAVE	31	0	X			<a href="#">Section 2.2.10.3 on page 161</a>
	WatchLo0	18	0	X			<a href="#">Section 2.2.10.4 on page 162</a>
	WatchLo1	18	1	X			
	WatchLo2	18	2	X			
	WatchLo3	18	3	X			
	WatchHi0	19	0	X			<a href="#">Section 2.2.10.5 on page 162</a>
	WatchHi1	19	1	X			
	WatchHi2	19	2	X			
	WatchHi3	19	3	X			
	TraceControl	23	1			X	<a href="#">Section 2.2.11.1 on page 164</a>
	TraceControl2	23	2			X	<a href="#">Section 2.2.11.2 on page 166</a>
	TraceControl3	24	2			X	<a href="#">Section 2.2.11.3 on page 167</a>
	UserTraceData1	23	3			X	<a href="#">Section 2.2.11.4 on page 168</a>
	UserTraceData2	24	3			X	<a href="#">Section 2.2.11.5 on page 169</a>
	TraceIPBC	23	4	X			<a href="#">Section 2.2.11.6 on page 169</a>
TraceDBPC	23	5	X			<a href="#">Section 2.2.11.7 on page 170</a>	
User Mode Support	YQMask	1	4	X			<a href="#">Section 2.2.12.1 on page 172</a>
	HWREna	7	0	X			<a href="#">Section 2.2.12.2 on page 172</a>
	UserLocal	4	2		X		<a href="#">Section 2.2.12.3 on page 174</a>
	LLAddr	17	0		X		<a href="#">Section 2.2.12.4 on page 174</a>
Kernel Mode Support	KScratch1	31	2	X			<a href="#">Section 2.2.13.1 on page 175</a>
	KScratch2	31	3	X			<a href="#">Section 2.2.13.2 on page 175</a>
	KScratch3	31	4	X			<a href="#">Section 2.2.13.3 on page 175</a>
Memory Mapped	CDMMBase	15	2			X	<a href="#">Section 2.2.14.1 on page 176</a>
	CMGCRBase	15	3	See note <sup>3</sup>			<a href="#">Section 2.2.14.2 on page 177</a>

1. *KSU*, *FR*, and *CU0-3* per-TC. See [Section 2.2.1.10 “Status \(CP0 Register 12, Select 0\)”](#).

2. *ASID* per-TC. See [Section 2.2.2.4 “EntryHi \(CP0 Register 10, Select 0\)”](#).

3. This register is per system. It is not instantiated per VPE, per TC, or per Core. Rather, there is only one of these registers in the entire multiprocessing system.

## 2.1.2 CP0 Registers Grouped by Number

The following table provides a numerical listing of the interAptiv CP0 registers. Click on a Name column entry to provide a link to the desired register.

**Table 2.2 CP0 Registers Grouped by Number**

Register			Function	Per			Location
Num	Sel	Name		VPE	TC	Proc	
0	0	Index	Index into the TLB array	X			<a href="#">Section 2.2.2.1</a>
0	1	MVPControl	Processor-wide multithreading control			X	<a href="#">Section 2.2.8.7</a>
0	2	MVPConf0	Processor's multithreading resources			X	<a href="#">Section 2.2.8.8</a>
0	3	MVPConf1	Processor's multithreading resources			X	<a href="#">Section 2.2.8.9</a>
1	0	Random	Randomly generated index into the TLB array	X			<a href="#">Section 2.2.2.2</a>
1	1	VPEControl	VPE control and status	X			<a href="#">Section 2.2.8.1</a>
1	2	VPEConf0	Initializable per-VPE resource lists	X			<a href="#">Section 2.2.8.2</a>
1	3	VPEConf1	Initializable per-VPE resource lists	X			<a href="#">Section 2.2.8.3</a>
1	4	YQMask	Defines valid inputs for yield instruction	X			<a href="#">Section 2.2.12.1</a>
1	5	VPESchedule	Per-VPE thread policy hints	X			<a href="#">Section 2.2.8.4</a>
1	6	VPEScheFBack	Per-VPE information from policy manager	X			<a href="#">Section 2.2.8.5</a>
1	7	VPEOpt	Per-VPE cache-way inhibition	X			<a href="#">Section 2.2.8.6</a>
2	0	EntryLo0	Low-order portion of the TLB entry for even-numbered virtual pages. This register is reserved if the TLB is not implemented.	X			<a href="#">Section 2.2.2.3</a>
2	1	TCStatus	Status and control for each TC		X		<a href="#">Section 2.2.7.1</a>
2	2	TCBind	VPE affiliation and own TC number of this TC		X		<a href="#">Section 2.2.7.2</a>
2	3	TCRestart	Where this TC will next fetch code from		X		<a href="#">Section 2.2.7.3</a>
2	4	TCHalt	Set 1 to freeze the TC for inspection/modification		X		<a href="#">Section 2.2.7.4</a>
2	5	TCContext	Read/write scratch register for OS to maintain thread ID		X		<a href="#">Section 2.2.7.5</a>
2	6	TCSchedule	Per-TC thread scheduling hints		X		<a href="#">Section 2.2.7.6</a>
2	7	TCScheFBack	Per-TC information from policy manager		X		<a href="#">Section 2.2.7.7</a>
3	0	EntryLo1	Low-order portion of the TLB entry for odd-numbered virtual pages. This register is reserved if the TLB is not implemented.	X			<a href="#">Section 2.2.2.3</a>
3	7	TCOpt	Per-TC cache-way inhibition		X		<a href="#">Section 2.2.7.8</a>
4	0	Context	Pointer to page table entry in memory. This register is reserved if the TLB is not implemented.	X			<a href="#">Section 2.2.2.5</a>
4	1	ContextConfig	Defines the bits of the Context register into which the high order bits of the virtual address causing a TLB exception will be written, and how many bits of that virtual address will be extracted.	X			<a href="#">Section 2.2.2.5</a>
4	2	UserLocal	User information that can be written by privileged software and read via RDHWR register 29		X		<a href="#">Section 2.2.12.3</a>

**Table 2.2 CP0 Registers Grouped by Number (continued)**

Register			Function	Per			Location
Num	Sel	Name		VPE	TC	Proc	
5	0	PageMask	PageMask controls the variable page sizes in TLB entries. This register is reserved if the TLB is not implemented.	X			<a href="#">Section 2.2.2.7</a>
5	2	SegCtl0	Segmentation control register 0. Used for enhanced virtual addressing (EVA).	X			<a href="#">Section 2.2.3.1</a>
5	3	SegCtl1	Segmentation control register 1. Used for enhanced virtual addressing (EVA).	X			<a href="#">Section 2.2.3.2</a>
5	4	SegCtl2	Segmentation control register 2. Used for enhanced virtual addressing (EVA).	X			<a href="#">Section 2.2.3.3</a>
6	0	Wired	Controls the number of fixed (“wired”) TLB entries. This register is reserved if the TLB is not implemented.	X			<a href="#">Section 2.2.2.9</a>
6	1	SRSConf0	Write these to use TCs as shadow registers.	X			<a href="#">Section 2.2.7.9</a>
6	2	SRSConf1	Write these to use TCs as shadow registers.	X			<a href="#">Section 2.2.7.10</a>
6	3	SRSConf2	Write these to use TCs as shadow registers.	X			
6	4	SRSConf3	Write these to use TCs as shadow registers.	X			
6	5	SRSConf4	Write these to use TCs as shadow registers.	X			
7	0	HWREna	Enables access via the RDHWR instruction to selected hardware registers in non-privileged mode.	X			<a href="#">Section 2.2.12.2</a>
8	0	BadVAddr	Reports the address for the most recent address-related exception.	X			<a href="#">Section 2.2.2.10</a>
9	0	Count	Processor cycle count.	X			<a href="#">Section 2.2.5.1</a>
10	0	EntryHi	High-order portion of the TLB entry. This register is reserved if the TLB is not implemented.	X	X <sup>1</sup>		<a href="#">Section 2.2.2.4</a>
11	0	Compare	Timer interrupt control.	X			<a href="#">Section 2.2.5.2</a>
12	0	Status	Processor status and control.	X	X <sup>2</sup>		<a href="#">Section 2.2.1.10</a>
12	1	IntCtl	Setup for interrupt vector and interrupt priority features.	X			<a href="#">Section 2.2.1.11</a>
12	2	SRSCtl	Shadow register set selectors	X			<a href="#">Section 2.2.7.11</a>
12	3	SRSMap	In vectored interrupt mode, determines which shadow set is used for each interrupt source.	X			<a href="#">Section 2.2.7.12</a>
13	0	Cause	Cause of last exception.	X			<a href="#">Section 2.2.4.1</a>
14	0	EPC	Program counter at last exception.	X			<a href="#">Section 2.2.4.2</a>
15	0	PRId	Processor identification and revision.	X			<a href="#">Section 2.2.1.8</a>
15	1	EBase	Exception base address.	X			<a href="#">Section 2.2.1.9</a>
15	2	CDMMBase	Common Device Memory Map Base Address			X	<a href="#">Section 2.2.14.1</a>
15	3	CMGCR	Global Configuration Register Base Address	See note <sup>3</sup>			<a href="#">Section 2.2.14.2</a>
16	0	Config	Configuration register.	X			<a href="#">Section 2.2.1.1</a>
16	1	Config1	Configuration for MMU, caches etc.	X			<a href="#">Section 2.2.1.2</a>
16	2	Config2	Configuration for MMU, caches etc.	X			<a href="#">Section 2.2.1.3</a>
16	3	Config3	Interrupt and ASE capabilities	X			<a href="#">Section 2.2.1.4</a>

**Table 2.2 CP0 Registers Grouped by Number (continued)**

Register			Function	Per			Location
Num	Sel	Name		VPE	TC	Proc	
16	4	Config4	Indicates presence of Config5 register	X			<a href="#">Section 2.2.1.5</a>
16	5	Config5	Provides information on EVA and cache error exception vector.	X			<a href="#">Section 2.2.1.6</a>
16	7	Config7	interAptiv family-specific configuration register.	X			<a href="#">Section 2.2.1.7</a>
17	0	LLAddr	Address associated with last LL instruction of a “load-linked/store-conditional” instruction pair.		X		<a href="#">Section 2.2.12.4</a>
18	0	WatchLo0	Watchpoint address associated with instruction watchpoint 0.	X			<a href="#">Section 2.2.10.4</a>
18	1	WatchLo1	Watchpoint address associated with instruction watchpoint 1.	X			
18	2	WatchLo2	Watchpoint address associated with data watchpoint 0.	X			
18	3	WatchLo3	Watchpoint address associated with data watchpoint 1.	X			
19	0	WatchHi0	Watchpoint ASID and Mask associated with instruction watchpoint 0.	X			<a href="#">Section 2.2.10.5</a>
19	1	WatchHi1	Watchpoint ASID and Mask associated with instruction watchpoint 1.	X			
19	2	WatchHi2	Watchpoint ASID and Mask associated with data watchpoint 0.	X			
19	3	WatchHi3	Watchpoint ASID and Mask associated with data watchpoint 1.	X			
23	0	Debug	EJTAG Debug register.	X	X		<a href="#">Section 2.2.10.1</a>
23	1	TraceControl	EJTAG Trace Control register			X	<a href="#">Section 2.2.11.1</a>
23	2	TraceControl2	EJTAG Trace Control2 register			X	<a href="#">Section 2.2.11.2</a>
23	3	UserTraceData1	EJTAG User Trace Data1 register			X	<a href="#">Section 2.2.11.4</a>
23	4	TraceIBPC	EJTAG Trace Instruction breakpoint control register	X			<a href="#">Section 2.2.11.6</a>
23	5	TraceDBPC	EJTAG Trace Debug breakpoint control register	X			<a href="#">Section 2.2.11.7</a>
24	0	DEPC	Restart address from last EJTAG debug exception.	X			<a href="#">Section 2.2.10.2</a>
24	2	TraceControl3	EJTAG Trace Control3 register			X	<a href="#">Section 2.2.11.3</a>
24	3	UserTraceData2	EJTAG User Trace Data2 register			X	<a href="#">Section 2.2.11.5</a>
25	0	PerfCtl0	Performance counter 0 control.		X		<a href="#">Section 2.2.9.1</a>
25	1	PerfCnt0	Performance counter 0.		X		
25	2	PerfCtl1	Performance counter 1 control.		X		<a href="#">Section 2.2.9.2</a>
25	3	PerfCnt1	Performance counter 1.		X		
26	0	ErrCtl	Software test enable of way-select and Data RAM arrays for I-Cache and D-Cache.	X			<a href="#">Section 2.2.6.10</a>
27	0	CacheErr	Records information about cache parity errors	X			<a href="#">Section 2.2.6.11</a>
28	0	ITagLo	Cache tag read/write interface for I-cache.	X			<a href="#">Section 2.2.6.1</a>
28	1	IDataLo	Low-order data read/write interface for I-cache.	X			<a href="#">Section 2.2.6.2</a>
28	2	DTagLo	Cache tag read/write interface for D-cache.	X			<a href="#">Section 2.2.6.4</a>

**Table 2.2 CP0 Registers Grouped by Number (continued)**

Register			Function	Per			Location
Num	Sel	Name		VPE	TC	Proc	
28	3	DDataLo	Low-order data read/write interface for D-cache.	X			<a href="#">Section 2.2.6.6</a>
28	4	L23TagLo	Cache tag read/write interface for L2-cache.	X			<a href="#">Section 2.2.6.7</a>
28	5	L23DataLo	Low-order data read/write interface for L2-cache.	X			<a href="#">Section 2.2.6.8</a>
29	1	IDataHi	High-order data read/write interface for I-cache.	X			<a href="#">Section 2.2.6.3</a>
29	2	DTagHi	High-order portion of data cache tag.	X			
29	5	L23DataHi	High-order data read/write interface for L2-cache.	X			<a href="#">Section 2.2.6.9</a>
30	0	ErrorEPC	Program counter at last error.	X			<a href="#">Section 2.2.4.3</a>
31	0	DESAVE	Debug handler scratchpad register.	X			<a href="#">Section 2.2.10.3</a>
31	2	KScratch1	Kernel scratch pad register 1.	X			<a href="#">Section 2.2.13.1</a>
31	3	KScratch2	Kernel scratch pad register 2.	X			<a href="#">Section 2.2.13.2</a>
31	4	KScratch3	Kernel scratch pad register 3.	X			<a href="#">Section 2.2.13.3</a>

1. *ASID* per-TC. See [Section 2.2.2.4 “EntryHi \(CP0 Register 10, Select 0\)”](#).

2. *KSU,FR*, and *CU0-3* per-TC. See [Section 2.2.1.10 “Status \(CP0 Register 12, Select 0\)”](#).

3. This register is per system. This register is not instantiated per VPE, per TC, or per Core. Rather, there is only one of these registers in the entire multiprocessing system.



## 2.2 CP0 Register Descriptions

The following subsections describe the CP0 registers listed in the tables above.

### 2.2.1 CPU Configuration and Status Registers


This section contains the following CPU Configuration and Status registers.

- [Section 2.2.1.1, "Device Configuration — Config \(CP0 Register 16, Select 0\)" on page 61](#)
- [Section 2.2.1.2, "Device Configuration 1 — Config1 \(CP0 Register 16, Select 1\)" on page 64](#)
- [Section 2.2.1.3, "Device Configuration 2 — Config2 \(CP0 Register 16, Select 2\)" on page 66](#)
- [Section 2.2.1.4, "Device Configuration 3 — Config3 \(CP0 Register 16, Select 3\)" on page 69](#)
- [Section 2.2.1.5, "Device Configuration 4 — Config4 \(CP0 Register 16, Select 4\)" on page 70](#)
- [Section 2.2.1.6, "Device Configuration 5 — Config5 \(CP0 Register 16, Select 5\)" on page 71](#)
- [Section 2.2.1.7, "Device Configuration 7 — Config7 \(CP0 Register 16, Select 7\)" on page 73](#)
- [Section 2.2.1.8, "Processor ID — PRId \(CP0 Register 15, Select 0\)" on page 75](#)
- [Section 2.2.1.9, "Exception Base Address — EBase \(CP0 Register 15, Select 1\)" on page 76](#)
- [Section 2.2.1.10, "Status \(CP0 Register 12, Select 0\)" on page 77](#)
- [Section 2.2.1.11, "Interrupt Control — IntCtl \(CP0 Register 12, Select 1\)" on page 82](#)

#### 2.2.1.1 Device Configuration — Config (CP0 Register 16, Select 0)

The main role of the Config register is to be a read-only repository of information about the interAptiv core resources, encoded so as to be useful to operating system initialization code.

**Figure 2.1 Config Register Format**



**Table 2.3 Field Descriptions for Config Register**

Name	Bit(s)	Description	Read/Write	Reset State
M	31	This bit is hardwired to '1' to indicate the presence of the Config1 register.	R	1
K23	30:28	CCA for kseg2 and kseg3 when the core is configured with MPU and MPU segmentation is not enabled.	R	0 (TLB) 2 (MPU)
KU	27:25	CCA for kuseg when the core is configured with MPU and MPU segmentation is not enabled.	R	0 (TLB) 2 (MPU)
ISP	24	Instruction Scratch Pad RAM present. 0: Instruction scratch pad RAM (ISPRAM) is not implemented. 1: Instruction scratch pad RAM (ISPRAM) is implemented.	R	Preset

**Table 2.3 Field Descriptions for Config Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
DSP	23	Data Scratch Pad RAM present. 0: Data scratch pad RAM (DSPRAM) is not implemented. 1: Data scratch pad RAM (DSPRAM) is implemented. This bit should not be confused with the MIPS DSP ASE, whose presence is indicated by <code>Config3<sub>DSP</sub></code> .	R	Preset
UDI	22	User-Defined Instruction. 0: The interAptiv core does not contain user-defined "CorExtend" instructions. 1: The interAptiv core contains user-defined "CorExtend" instructions. This bit is automatically updated by hardware when <code>VPEconf0<sub>NCX</sub></code> is written.	R	Preset
SB	21	Read-only "SimpleBE" bus mode indicator, which reflects the interAptiv input signal <code>SI_SimpleBE</code> . 0: No reserved byte enabled on the OCP interface. 1: Only simple byte enables allows on the OCP interface. If set by hardware, the interAptiv core will only do simple partial-word transfers on its OCP interface; that is, the only partial-word transfers will be byte, aligned half-word, and aligned word. If zero, it may generate partial-word transfers with an arbitrary set of bytes enabled. This generates less requests, but may not be supported by all downstream devices.	R	Externally Set
MDU	20	MDU Implementation. This bit is encoded as follows:  0: High performance MDU 1: Reduced-area MDU	R	Preset
0	19	Must be written as zeros; returns zeros on reads.	0	0
MM	18	Write Merge. This bit indicates whether write-through merging is enabled in the 32-byte collapsing write buffer. 0: No merging allowed 1: Merging allowed Setting this bit allows writes resulting from separate store instructions in write-through mode to be merged into a single transaction at the interface. The state of this bit does not affect cache writebacks (which are always whole blocks together) or uncached writes (which are never merged). Note that write-through caching is not supported in the interAptiv core, so this bit has no meaning. This bit is implemented per-processor and not per-VPE as are other writable fields of this register.	R/W	1
0	17	Must be written as zero; returns zero on read.	R	0

**Table 2.3 Field Descriptions for Config Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
BM	16	<p>Burst Mode.</p> <p>0: Sequential burst mode</p> <p>1: SubBlock burst mode</p> <p>This bit reads 0 when the bus uses sequential burst ordering and reads 1 when it uses sub-block burst ordering. This bit is set by the input signal <i>SI_SBlock</i> signal to match the system controller.</p> <p>Note that the interAptiv core only supports sequential burst ordering. Hence this bit is always zero.</p>	R	0
BE	15	<p>Endian mode.</p> <p>0: Little endian</p> <p>1: Big endian</p> <p>This bit is written by hardware based on the state of the <i>SI_Endian</i> input pin.</p>	R	Externally Set
AT	14:13	<p>Architecture type implemented by the processor.</p> <p>This field is always 00 to indicate the MIPS32 architecture.</p>	R	0
AR	12:10	<p>Architecture release.</p> <p>0x0 = Release 1</p> <p>0x1 = Release 2 or Release 3</p> <p>This bit always reads 1 to reflect Release 3 of the MIPS32 architecture.</p>	R	1
MT	9:7	<p>MMU Type.</p> <p>000: Reserved</p> <p>001: Standard TLB</p> <p>010: Reserved</p> <p>011: FMT. If MPU, but MPU_Config.EN = 0</p> <p>100 - 101: Reserved</p> <p>110: CDMM. If MPU and MPU_Config_EN = 1</p> <p>111: Reserved</p> <p>If the MPU option was selected at build time and the MPU_Config,EN bit is 0, then the MPU address translation and default attributes are based on a Fixed Mapping Translation.</p> <p>If the MPU_Config,EN bit is 1, then segment control in the MPU is enabled and used for address translation and default attributes. For more information on segment control, refer to Section 4.2, MPU registers.</p>	R	Preset
0	6:3	Must be written as zero; returns zero on read.	R	0
K0	2:0	Kseg0 coherency attribute of the page. See <a href="#">Table 2.19</a> for the field encoding.	R/W	2

### 2.2.1.2 Device Configuration 1 — Config1 (CP0 Register 16, Select 1)

The Config1 register provides information such as the size of the TLB and the L1 instruction and data cache parameters. It also contains a series of single bits that indicate the presence of selected logic units on the interAptiv core.

**Figure 2.2 Config1 Register Format**

31	30	25	24	22	21	19	18	16	15	13	12	10	9	7	6	5	4	3	2	1	0
M	MMUSize	IS	IL	IA	DS	DL	DA	C2	MD	PC	WR	CA	EP	FP							

**Table 2.4 Field Descriptions for Config1 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
M	31	Continuation bit, set to 1 to indicate that the Config2 register is implemented.	R	1
MMU Size	30:25	This field contains the number of entries in the TLB minus one. This field reads as 0 when the MPU option is selected at build time.	R	Preset
IS	24:22	L1 Instruction cache number of sets per way. This field indicates the number of sets per way in the L1 instruction cache. The number of sets is multiplied by the number of ways and the line size to derive the cache size. In this case, the number of sets defines the cache size since the line size and number of ways in the interAptiv core are fixed. This field is encoded as follows:  000: 64 sets per way (equates to 8 KByte instruction cache) 001: 128 sets per way (equates to 16 KByte instruction cache) 010: 256 sets per way (equates to 32 KByte instruction cache) 011: 512 sets per way (equates to 64 KByte instruction cache) 100 - 110: Reserved 111: 32 sets per way (equates to 4 KByte instruction cache)  Because the line size and associativity are fixed for the interAptiv instruction cache as defined in the IL and IA fields below, the IS field is used to determine the overall cache size as follows:  If this field is set to 2, the instruction cache size would be: 256 sets/way x 32 bytes/line x 4 sets per way = 32 KBytes.  If this field is set to 3, the instruction cache size would be: 512 sets/way x 32 bytes/line x 4 sets per way = 64 KBytes.	R	Preset
IL	21:19	L1 Instruction cache line size. In the interAptiv core, the instruction cache line size is fixed at 32 bytes. As such, this field is encoded as follows:  000 - 011: Reserved 100: 32 byte line size 101 - 111: Reserved	R	Preset

**Table 2.4 Field Descriptions for Config1 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
IA	18:16	L1 Instruction cache associativity. In the interAptiv core, the instruction cache associativity is fixed at 4 ways. As such, this field is encoded as follows: 000 - 010: Reserved 011: 4-ways 100 - 111: Reserved  A default value of 3 indicates a 4-way set associative instruction cache. Refer to the IS field above to determine how to calculate the size of the L1 instruction cache.	R	3
DS	15:13	L1 Data cache number of sets per way. This field indicates the number of sets per way in the L1 data cache and is encoded as follows: The number of sets is multiplied by the number of ways and the line size to derive the cache size. In this case, the number of sets defines the cache size since the line size and number of ways in the interAptiv core are fixed. This field is encoded as follows: 000: 64 sets per way (equates to 8 KByte instruction cache) 001: 128 sets per way (equates to 16 KByte instruction cache) 010: 256 sets per way (equates to 32 KByte instruction cache) 011: 512 sets per way (equates to 64 KByte instruction cache) 100 - 110: Reserved 111: 32 sets per way (equates to 4 KByte instruction cache)  Because the line size and associativity are fixed for the interAptiv data cache as defined in the DL and DA fields below, the DS field is used to determine the overall cache size as follows:  If this field is set to 2, the data cache size would be: $256 \text{ sets/way} \times 32 \text{ bytes/line} \times 4 \text{ sets per way} = 32 \text{ KBytes.}$  If this field is set to 3, the data cache size would be: $512 \text{ sets/way} \times 32 \text{ bytes/line} \times 4 \text{ sets per way} = 64 \text{ KBytes.}$	R	Preset
DL	12:10	L1 data cache line size. In the interAptiv core, the data cache line size is fixed at 32 bytes. As such, this field is encoded as follows: 000 - 011: Reserved 100: 32 byte line size 101 - 111: Reserved	R	Preset
DA	9:7	L1 data cache associativity. In the interAptiv core, the data cache associativity is fixed at 4 ways. As such, this field is encoded as follows: 000 - 010: Reserved 011: 4-ways 100 - 111: Reserved  A default value of 3 indicates a 4-way set associative data cache.	R	3
C2	6	This bit is cleared to indicate that a coprocessor 2 does not exist in the system. This bit is automatically updated by hardware when VPEConf1 <sub>NCP2</sub> is written.	R	Preset

**Table 2.4 Field Descriptions for Config1 Register**


Name	Bit(s)	Description	Read/ Write	Reset State
MD	5	MDMX Application Specific Extension (ASE). A logic '0' indicates that the MDMX ASE is not implemented in the floating point unit (FPU) of the interAptiv core. Note that if the FPU is not implemented, this bit has no meaning.	R	0
PC	4	Performance counter present. There is at least one performance counter implemented in the interAptiv core. Hardware sets this bit if performance counters are present. Refer to the PerfCnt0-3 and PerfCnt0-3 registers for more information.	R	Preset
WR	3	Watchpoint registers present. This bit reads as '1' if the Watchpoint registers are present. Refer to the WatchLo 0-3/WatchHi 0-3 registers in <a href="#">Section 2.2.10.4 "Watch Low 0-3 — WatchLo0-3 (CP0 Register 18, Select 0-3)"</a> .	R	Preset
CA	2	MIPS16e present. This bit reads 1 to indicate the MIPS16e compressed-code instruction set is available.	R	Preset
EP	1	EJTAG unit present. This bit reads 1 if the EJTAG debug unit is provided on the interAptiv core.	R	1
FP	0	Floating Point Unit present. This bit is set to indicate that a floating point unit is present. The floating point unit is optional on the interAptiv core. <ul style="list-style-type: none"> <li>• When no FPU is present, this will be 0</li> <li>• When the multi-threaded FPU is present, this will be 1</li> </ul> When the single-threaded FPU is present, hardware automatically updates this field when VPEConf1 <sub>NCP1</sub> is written.	R	Preset

### 2.2.1.3 Device Configuration 2 — Config2 (CP0 Register 16, Select 2)

The Config2 register provides information about the size and organization of L2 and L3 caches. The Config2 register also has fields that indicate the presence of some extensions to the base MIPS32 architecture.

An L3 cache can be used with the interAptiv core. However, the core does not support passing of the L3 configuration information via the Config2 register. As such, the TU, TS, TL and TA bits of this register, which handle L3 operations, are not used and are all tied to 0. Information on L3 transfers may be available in an implementation specific register elsewhere in the system.

**Figure 2.3 Config2 Register Format**



**Table 2.5 Field Descriptions for Config2 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
M	31	This bit is hardwired to '1' to indicate the presence of the Config3 register.	R	1

**Table 2.5 Field Descriptions for Config2 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
TU	30:28	An L3 cache can be used with the interAptiv core. However, the core does not support passing of the L3 configuration data via the Config2 register. As such, the TU, TS, TL and TA bits of this register, which report L3 information, are not used and are all tied to 0. Details of the L3 configuration may be available in an implementation specific register elsewhere in the system.	R	0
TS	27:24		R	0
TL	23:20		R	0
TA	19:16		R	0
SU	15:13	These bits are reserved in the interAptiv core and is always 0.	R	0
L2B	12	<p>L2 cache bypass. Setting this bit disables or bypasses the L2 cache. Setting this bit also forces Config2<sub>SL</sub> to 0. Based on this information, most operating system code will conclude that there is no L2 cache on the system.</p> <p>The L2 cache receives the value of this bit from core 0. Setting this bit forces hardware to drive a series of internal handshake signals between the core to the CM2, placing the L2 cache into bypass mode.</p> <p>When this bit is set through a write operation, a subsequent read of this bit will not indicate a logic 1 until the L2 has asserted its internal handshake signal, indicating that it has been bypassed. Note that the L2 Bypass is a feature available on the MIPS L2 cache controller, and may not be available with non-MIPS L2 cache implementations.</p>	R/W	0

**Table 2.5 Field Descriptions for Config2 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
SS	11:8	<p>L2 cache number of sets per way. This field indicates the number of sets per way in the L2 cache of the Coherent Processing System (CPS) and is written by hardware at reset based on the state of the <i>L2_Sets[3:0]</i> signals.</p> <p>At IP configuration time, the user selects the cache size and the line size. Hardware then takes this information and selects the appropriate number of sets. See the example formulas below for determining the number of sets based on cache and line size. Note that a value of 0x0 in the SL field of this register indicates the 0K L2 cache option.</p> <p>This field is encoded as follows:</p> <p>0x0: 64 sets per way            0x1: 128 sets per way            0x2: 256 sets per way            0x3: 512 sets per way            0x4: 1024 sets per way            0x5: 2048 sets per way            0x6: 4096 sets per way            0x7: 8192 sets per way            0x8: 16384 sets per way            0x9: 32768 sets per way            0xA- 0xF: Reserved</p> <p>For example:</p> <p>If this field is set to 0x5, the SL field is set to 0x4, and the SA field is set to 0x4, the L2 cache size would be:</p> <p>2048 sets/way x 32 bytes/line x 8 ways = 512 KBytes</p> <p>Similarly, if this field is set to 0x9, the SL field is set to 0x4, and the SA field is set to 0x4, the L2 cache size would be:</p> <p>32768 sets/way x 32 bytes/line x 8 ways = 8 MBytes</p> <p>Note that the setting for 32768 sets/way cannot be used with the 64-byte line size because the interAptiv core does not support a 16 MB L2 cache size.</p>	R	Preset
SL	7:4	<p>L2 data cache line size. This field is written by hardware at reset based on the state of the <i>L2_LineSize[3:0]</i> signals. These signals are driven based on the customer's line size choice during IP configuration. As such, this field is encoded as follows:</p> <p>0x0: No L2 cache present            0x1 - 0x3: Reserved            0x4: 32 byte line size            0x5: 64 byte line size            0x6 - 0xF: Reserved</p>	R	Preset




**Table 2.5 Field Descriptions for Config2 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
SA	3:0	L2 cache associativity. In the interAptiv core, the L2 cache associativity is fixed at 8 ways. This field is written by hardware at reset based on the state of the <i>L2_Assoc[3:0]</i> signals. As such, this field is encoded as follows:  0x0 - 0x6: Reserved 0x7: 8-way set associative 0x8 - 0xF: Reserved	R	Preset

**2.2.1.4 Device Configuration 3 — Config3 (CP0 Register 16, Select 3)**

Config3 provides information about the presence of optional extensions to the base MIPS32 architecture in addition to those specified in Config2. All fields in the Config3 register are read-only.

**Figure 2.4 Config3 Register Format**



**Table 2.6 Field Descriptions for Config3 Register**

Name	Bit(s)	Description	Read/ Write	Reset State
M	31	Configuration continuation bit. This bit is always one to indicate the presence of Config4.	R	1
0	30	Must be written as zeros; returns zeros on read	0	0
CMGCR	29	Reads 1 to indicate that the Coherence Manager has a Global Configuration Register Space and the CMGCRBase cop0 register is implemented.	R	1
0	28:26	Must be written as zeros; returns zero on read.	R	0
SC	25	Segment Control implemented. This bit indicates whether the Segment Control registers SegCtl0, SegCtl1 and SegCtl2 are present.	R	1 (TLB) 0 (MPU)
0	24:14	Must be written as zero; returns zero on read.	R	0
ULRI	13	Reads 1 if the UserLocal Register is implemented.	R	Preset
RXI	12	Reads 1 if the <i>RIE</i> and <i>XIE</i> fields exist in the PageGrain register.	R	1 (TLB) 0 (MPU)
DSP2P	11	Indicates the MIPS DSP ASE revision.  0: Revision 1 (DSP R1) 1: Revision 2 (DSP R2)  If a DSP is implemented, it will be DSPR2 compliant and this bit will be set.	R	Preset
DSPP	10	Reads 1 if the MIPS DSP ASE extension is implemented.	R	Preset
CTXTC	9	Reads 1 if the ContextConfig register is implemented. The width of the BadVPN2 field in the Context register depends on the contents of the ContextConfig register.	R	1 (TLB) 0 (MPU)
0	8:7	Must be written as zero; returns zero on read.	R	0

**Table 2.6 Field Descriptions for Config3 Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
VEIC	6	Support for an external interrupt controller. This bit is set or cleared by hardware depending on whether the EIC option was selected at build time.  0: Support for EIC mode not supported. 1: Support of EIC mode supported.  The value of this bit is set by the static input, <i>SI_EICPresent</i> . This allows external logic to communicate whether an external interrupt controller is attached to the processor.	R	Externally Set
VInt	5	Vectored interrupts implemented. This bit indicates whether vectored interrupts are implemented. On the interAptiv core, this bit reads 1 to indicate the CPU can handle vectored interrupts.	R	1
SP	4	Reads 0 to indicate the CPU does not support 1 Kbyte TLB pages.	R	0
CDMM	3	Reads 1 to indicate the Common Device Memory Map (CDMM) feature is implemented, as well as the CDMMBase register is present.	R	1
MT	2	This bit reads 1 to indicate that the MIPS MT (multi-threading) ASE is implemented. This bit is always set in the interAptiv core to indicate that multithreading is present.	R	1
SM	1	Reads 0 to indicate the interAptiv does not include the instructions of the SmartMIPS ASE.	R	0
TL	0	Reads 1 to indicate PDTrace is supported.	R	Preset

### 2.2.1.5 Device Configuration 4 — Config4 (CP0 Register 16, Select 4)

The Config4 register encodes additional capabilities such as TLBINV instruction support and the number of kernel scratch registers.

**Figure 2.5 Config4 Register Format**



**Table 2.7 Field Descriptions for Config4 Register**

Name	Bit(s)	Description	Read/Write	Reset State
M	31	Configuration continuation bit. This bit is one to indicate the presence of Config5.	R	1
IE	30:29	TLBINV instruction support. This field is encoded as follows:  00: MPU 10: TLB based MMU  When configured with an MPU, the TLBINV and TLBINVF instructions are not supported. When configured with an MMU, the TLBINV and TLBINVF instructions are supported and operate on a single TLB entry.	R	0: MPU 2: TLB based MMU
0	28:24	Reserved. Must be written as zero. Ignored on reads.	R	0


**Table 2.7 Field Descriptions for Config4 Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
KScrExist	23:16	Indicates how many scratch registers are available to kernel-mode software within CP0 Register 31. In the interAptiv architecture, three kernel scratch registers are included at register selects 2, 3, and 4.  Each bit represents a select for CP0 Register 31. Bit 16 represents Select 0, Bit 23 represents Select 7. If the bit is set, the associated scratch register is implemented and available for kernel-mode software. Therefore, this field contains a value of 0x1C (8'b00011100). This indicates that bits 18 - 20 are set, corresponding to selects 2, 3, and 4.  These registers are used by the kernel for temporary storage of information. Refer to <a href="#">Section 2.2.13, "Kernel Mode Support Registers" on page 174</a> for more information.	R	0x1C
0	15:0	Reserved. Must be written as zero. Ignored on reads.	R	0

**2.2.1.6 Device Configuration 5 — Config5 (CP0 Register 16, Select 5)**

The Config5 register encodes additional capabilities for the address mode programming and cache error exceptions.

**Config5 Register Format**



**Table 2.8 Field Descriptions for Config5 Register**

Name	Bit(s)	Description	Read/Write	Reset State
R	31	Reserved. Must be written as zero. Ignored on reads.	R	0


**Table 2.8 Field Descriptions for Config5 Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
K	30	<p>This bit effects the cache coherency attributes, the boot exception vector overlay, and the location of the exception vector as follows:</p> <p>When this bit is cleared, the following events occur:</p> <ol style="list-style-type: none"> <li>1. For a TLB-based MMU, the <code>Config<sub>K0</sub></code> field is used to set the cache coherency attributes for the <code>kseg0</code> region (0x8000_0000 - 0x9FFF_FFFF). For a MPU, the use of <code>Config.K0</code> is controlled by <code>MPU_Config.En=0</code>.</li> <li>2. Hardware creates two boot overlay segments, one for <code>kseg0</code> and one for <code>kseg1</code>.</li> <li>3. The exception vectors are forced to reside in <code>kseg0/kseg1</code> by ignoring the state of bits 31:30 of the <code>EBase</code> register as well as the <code>SI_ExceptionBase[31:30]</code> pins and forcing them to a value of 2'b10.</li> </ol> <p>When this bit is set, the following events occur:</p> <ol style="list-style-type: none"> <li>1: For a TLB-based MMU, the <code>Config<sub>K0</sub></code> field is ignored and the cache coherency attributes are derived from the C fields of the various segmentation control registers (<code>SegCtl0</code> - <code>SegCtl2</code>).</li> <li>2. Hardware creates one boot overlay segment that can reside anywhere in virtual address space.</li> <li>3. The exception vectors are not forced to reside in <code>kseg0/kseg1</code>. Rather, bits 31:30 of the <code>EBase</code> register, as well as the <code>SI_ExceptionBase[31:30]</code> signals and used to place the exception vectors anywhere within virtual address space.</li> </ol>	R/W	0
CV	29	<p>Cache error exception vector control. Disables logic forcing use of <code>kseg1</code> region in the event of a Cache Error exception when <code>Status<sub>BEV</sub> = 0</code>.</p> <p>When the CV bit is cleared, bits 31:30 of the <code>EBase</code> Register are fixed with the value 2'b10 to force the exception base address to be in the <code>kseg0</code> or <code>kseg1</code> unmapped virtual address segments. Bit 29 of exception base address will be forced to 1 on Cache Error exceptions so the exception handler will be executed from the uncached <code>kseg1</code> segment.</p> <p>When the CV bit is set, the <code>ExcBase</code> field is expanded to include bits 31:30 to facilitate programmable memory segmentation.</p>	R/W	0
EVA	28	This bit is a logic one to indicate support for enhanced virtual address (EVA).	R	1 (TLB) 0 (MPU)
R	27:15	Reserved. Must be written as zero. Ignored on reads.	R	0
CA2	14	<p>Extension of <code>Config<sub>1CA</sub></code>. Indicates implementation of version 2.0 of MIPS16e, MIPS16e2. For CA2 to be 1, <code>Config<sub>1CA</sub></code> must also be 1. This bit is encoded as follows:</p> <p>0: Version 2.0 is not implemented. 1: Version 2.0 is implemented.</p>	R	Preset
R	13:0	Reserved. Must be written as zero. Ignored on reads.	R	0

### 2.2.1.7 Device Configuration 7 — Config7 (CP0 Register 16, Select 7)

This register controls machine-specific features of the interAptiv core. A few of them are for hardware interface adaptation, but most are for chip or system test only. They default to a "safe" value. Most software, including boot-strap software, can and should ignore these registers unless specifically advised to use them.

**Figure 2.6 Config7 Register Format**



**Table 2.9 Field Descriptions for Config7 Register**

Name	Bit(s)	Description	Read/Write	Reset State						
WII	31	Wait IE/IXMT Ignore: Indicates that this processor will allow an interrupt to unblock a WAIT instruction even if IE or IXMT is preventing the interrupt from being taken. This avoids problems using the WAIT instruction for 'bottom half' interrupt servicing.	R	1						
0	30:24	These bits are unused and should be written as 0.	R	0						
MCRO	23	Indicates that SAVE/RESTORE macro instructions are available as MIPS32 CorExtend opcodes.	R	Preset						
CP	22	Indicates that COPYW/UCOPYW ASMACRO instructions are supported.	R	Preset						
RAMSlp	21	RAM sleep. Setting to 1 disables RAM sleep modes by preventing the sleep input to the L1 cache RAMs from being asserted.	R/W	0						
NCWB	20	Non-Coherent Writeback. When set, HitWB cacheops to a non-coherent address are written using a non-coherent CCA.	R/W	0						
PCT	19	Performance Counters per TC: This bit indicates to software that the performance counter registers are instantiated per TC rather than per processor. This bit is implemented per-processor.	R	1						
HCI	18	Hardware Cache Initialization: Indicates that a cache does not require initialization by software. This bit will most likely only be set on simulation-only cache models and not on real hardware.	R	Preset						
FPR	17	Floating Point Ratio: Indicates clock ratio between integer CPU and floating point unit on interAptiv CPUs. This bit is implemented per-processor. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>FP clock frequency is the same as the integer clock</td> </tr> <tr> <td>1</td> <td>FP clock frequency is one-half the integer clock</td> </tr> </tbody> </table>	Encoding	Description	0	FP clock frequency is the same as the integer clock	1	FP clock frequency is one-half the integer clock	R	Preset
Encoding	Description									
0	FP clock frequency is the same as the integer clock									
1	FP clock frequency is one-half the integer clock									
AR	16	Alias removed. Hardware sets this bit to indicate that the L1 data cache is configured to avoid cache aliases.	R	Preset						
0	15:13	Reserved. Write as zero. Ignored on reads.	R	0						

**Table 2.9 Field Descriptions for Config7 Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
ICWP	12	Instruction cache way prediction. Setting this bit disables instruction cache way prediction.	R/W	0
BusSlp	11	Bus sleep mode. Setting to 1 prevents the core from going to sleep while bus requests are pending.	R/W	0
IAR	10	Instruction Alias Removed. Indicates that the interAptiv core has hardware support to remove instruction cache aliasing. The virtual aliasing hardware can be disabled via the IVAD bit described below.	R	Preset
IVAD	9	Instruction Virtual Aliasing disabled. The hardware required to resolve instruction cache virtual aliasing is always present in the interAptiv core as noted by the default state of the IAR bit shown above. However, software can toggle the IVAD bit to enable or disable the virtual aliasing hardware for the instruction cache. Setting this bit disables the hardware alias removal on the instruction cache. If this bit is cleared, the <b>CACHE Hit Invalidate</b> and <b>SYNCI</b> instructions look up all possible aliased locations and invalidate the given cache line in all of them. This bit is Read-only if IAR = 0. This bit is implemented per-VPE.	R/W	0 (hardware aliasing removal enabled)
ES	8	Externalize <b>sync</b> . If this bit is set, and if the downstream device (toward memory) is capable of accepting SYNCs (indicated by the pin <i>Sl_SyncTxEn</i> ), the <b>sync</b> instruction causes a SYNC-specific transaction to go out on the external bus. If this bit is cleared or if <i>Sl_SyncTxEn</i> is deasserted, no transaction will go out, but all SYNC handling internal to the CPU will nevertheless be performed. The <b>sync</b> instruction is signalled on the interAptiv's OCP interface as an "ordering barrier" transaction. The transaction is an extension to the OCP standards, and system controllers which don't support it typically under-decode it as a read from the boot ROM area. But that's going to be quite slow, so set this bit only if your system understands the synchronizing transaction. When this bit is read, the value returned depends on the state of the <i>Sl_SyncTxEn</i> pin. If <i>Sl_SyncTxEn</i> is 0, a value of 0 is returned. If <i>Sl_SyncTxEn</i> is 1, the value returned is the last value that was written to this bit.	R	1
BTLM	7	Block TC on Load Miss: Setting this bit will cause a TC to be suspended once a load miss has been detected, rather than waiting for a dependent instruction to try to access the load data. This can increase pipeline utilization and provide fairer allocation of miss resources, but does limit the parallel servicing of cache misses from a single TC. This bit is implemented per-processor.	R/W	0
CPOOO	6	Out-of-order data return on the Coprocessor interfaces: Writing 1 to this bit disables the out-of-order data return for the FPU and COP2.	R/W	0
NBSLU	5	Non-Blocking LSU: Writing 1 to this field will lock the LSU and ALU pipelines together. This forces LSU pipeline stalls to also stall the ALU pipeline. This bit is implemented per-processor.	R/W	0


**Table 2.9 Field Descriptions for Config7 Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
ULB	4	Uncached load blocking. Set to 1 to make all uncached loads blocking (a program usually only blocks when it uses the data which is loaded).	R/W	0
BP	3	Branch prediction. When set, no branch prediction is done, and all branches and jumps stall. This bit is implemented per-VPE.	R/W	0
RPS	2	Return prediction stack. When set, the return address branch predictor is disabled, so <code>jr \$31</code> is treated just like any other jump register. An instruction fetch stalls after the branch delay slot, until the jump instruction reaches the Address Generation pipeline and can provide the right address.	R/W	0
BHT	1	Branch history table. When set, the branch history table is disabled and all branches are predicted taken. This bit is don't care if <code>Config7_BP</code> is set.	R/W	0
SL	0	Scheduled loads. When set, non-blocking loads are disabled. Normally the interAptiv core continues after a load instruction, even if it misses in the D-cache, until the data is used. When this bit is set, the CPU stalls on any D-cache load miss.	R/W	0

### 2.2.1.8 Processor ID — PRId (CP0 Register 15, Select 0)

The Processor Identification (PRId) register is a 32 bit read-only register that contains information identifying the manufacturer, manufacturing options, processor identification, and revision level of the processor.

**Figure 2.7 PRId Register Format**



**Table 2.10 Field Descriptions for PRId Register**

Name	Bit(s)	Description	Read/Write	Reset State
CoOpt	31:24	Company Option. Should be a number between 0 and 127— higher values are reserved by MIPS Technologies.	R	Preset
CoID	23:16	Company ID. Identifies the company that designed or manufactured the processor. In the interAptiv, this field contains a value of 1 to indicate MIPS Technologies, Inc.	R	1
ProcType	15:8	Processor ID. Identifies the type of processor. This field allows software to distinguish between the various types of processors from MIPS Technologies. The value of this field is 0xA1 for the interAptiv core.	R	A1

**Table 2.10 Field Descriptions for PRId Register**

Name	Bit(s)	Description	Read/Write	Reset State												
Rev	7:0	<p>The revision number of the interAptiv design. This field allows software to distinguish between one revision and another of the same processor type.</p> <p>This field is broken up into the following three subfields:</p> <table border="1"> <thead> <tr> <th>Bit(s)</th> <th>Name</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>7:5</td> <td>Major Revision</td> <td>This number is increased on major revisions of the interAptiv core.</td> </tr> <tr> <td>4:2</td> <td>Minor Revision</td> <td>This number is increased on each incremental revision of the processor and reset on each new major revision.</td> </tr> <tr> <td>1:0</td> <td>Patch Level</td> <td>If a patch is made to modify an older revision of the processor, this field will be incremented.</td> </tr> </tbody> </table>	Bit(s)	Name	Meaning	7:5	Major Revision	This number is increased on major revisions of the interAptiv core.	4:2	Minor Revision	This number is increased on each incremental revision of the processor and reset on each new major revision.	1:0	Patch Level	If a patch is made to modify an older revision of the processor, this field will be incremented.	R	Preset
Bit(s)	Name	Meaning														
7:5	Major Revision	This number is increased on major revisions of the interAptiv core.														
4:2	Minor Revision	This number is increased on each incremental revision of the processor and reset on each new major revision.														
1:0	Patch Level	If a patch is made to modify an older revision of the processor, this field will be incremented.														

### 2.2.1.9 Exception Base Address — EBase (CP0 Register 15, Select 1)

The EBase register is a read/write register containing the base address of the exception vectors used when StatusBEV equals 0, and a read-only CPU number value that may be used by software to distinguish different processors in a multi-processor system.

The EBase register provides the ability for software to identify the specific processor within a multi-processor system, and allows the exception vectors for each processor to be different. Bits 31:12 of the EBase register are concatenated with zeros to form the base of the exception vectors when StatusBEV is 0. The exception vector base address comes from the fixed defaults when StatusBEV is 1, or for any EJTAG Debug exception. The reset state of bits 31:12 of the EBase register initialize the exception base register to 0x8000.0000, providing backward compatibility with Release 1 implementations.

The size of the ExcBase field depends on the state of the WG bit. At reset, the WG bit is cleared by default. In this case, the ExcBase field is comprised of bits 29:12. Bits 31:30 of the EBase Register are not writeable and are forced to a value of 2'b10 by hardware so that the exception handler will be executed from the *kseg0/kseg1* segments. This is shown in [Figure 2.8](#).

When the WG bit is set, bits 31:30 of the ExcBase field become writeable and are used to relocate the ExcBase field to other segments after they have been setup using the SegCtl0 through SegCtl2 registers. This is shown in [Figure 2.9](#). Note that if the WG bit is set by software (allowing bits 31:30 to become part of the ExcBase field) and then cleared, bits 31:30 can no longer be written by software and the state of these bits remains unchanged for any writes after WG was cleared. Therefore, it is the responsibility of software to write a value of 2'b10 to bits 31:30 of the EBase register prior to clearing the WG bit if it wants to ensure that future exceptions will be executed from the *kseg0* or *kseg1* segments.

Note that the WG bit is different from the CV bit in the SegCtl0 register located in [Section 2.2.3.1, "Segmentation Control 0 — SegCtl0 \(CP0 Register 5, Select 2\)"](#). Although their functions are similar, the CV bit applies only to cache error exceptions, whereas the WG bit applies to all exceptions.

If the value of the exception base register is to be changed, this must be done with StatusBEV equal to 1. The operation of the processor is **UNDEFINED** if the exception base field is written with a different value when StatusBEV is 0.



Combining bits 31:12 with the Exception Base field allows the base address of the exception vectors to be placed at any 4K page boundary.

**Figure 2.8 EBase Register Format — WG = 0**



**Figure 2.9 EBase Register Format — WG = 1**



**Table 2.11 Field Descriptions for EBase Register**

Name	Bit(s)	Description	Read/ Write	Reset State
ExcBase	31:12	Exception Base Address. The size and behavior of this field depends on the state of the WG bit. When the WG bit is set, the ExcBase field includes bits 31:30 to facilitate programmable memory segmentation. This field specifies the base address of the exception vectors when Status <sub>BEV</sub> is zero. Bits 31:30 can be written only when WG is set. When WG is zero, these bits are unchanged on a write.  When the WG bit is cleared, bits 31:30 of this field must be 2'b10 to make sure the exception vector maps to kseg0 or kseg1, conventionally used for OS code.  In a multi-core environment, setting EBase in any CPU to a unique value allows that CPU can have its own unique exception handlers.  This field should be written only when Status <sub>BEV</sub> is set so that any exception will be handled through the ROM entry points.	R/W	0x8000.0
WG	11	Write gate.  When the WG bit is set, the ExcBase field is expanded to include bits 31:30 of the EBase register to facilitate programmable memory segmentation controlled by the SegCtl0 through SegCtl2 registers.  When the WG bit is cleared, bits 31:30 of the EBase register are not writeable and remain unchanged from the last time that WG was cleared.	R/W	Externally Set
0	10	Reserved. Write as zero. Ignored on reads.	R	0
CPUNum	9:0	This field contains an identifier that will be unique among the CPU's in a multi-processor system. The value in this field is set by the <i>SI_CPUNum[9:0]</i> static input pins to the interAptiv core.  Note that in 2 VPE configurations, CPUNum[0] contains the VPE number.	R	Externally Set

### 2.2.1.10 Status (CP0 Register 12, Select 0)


The Status register is a read/write register that contains the operating mode, interrupt enabling, and diagnostic states of the processor. Fields in this register and the CP0 Debug register combine to create operating modes for the processor. Selected bits are encoded as follows to place the processor into one of the operating modes. Refer to the MMU chapter for more information on the various operating modes. A brief summary is provided below.

**Table 2.12 Operating Mode Encoding**

Status <sub>IE</sub>	Status <sub>ERL</sub>	Status <sub>EXL</sub>	Status <sub>KSU</sub>	Debug <sub>DM</sub>	Mode of Operation
1	0	0	x	0	Individual interrupts can be disabled/enabled using the Status <sub>IM7-0</sub> mask bits.
x	0	0	2'b2	0	<i>User Mode</i> . In user mode, the CPU has access only to the mapped kuseg address region.
x	0	0	2'b1	0	<i>Supervisor Mode</i> . In supervisor mode, the CPU has access to the top half of the kseg2 region (sometimes known as kseg3), but no access to CP0 registers or most kernel memory.  Note that Supervisor mode is not supported with the core is configured with an MPU.
x	x	x	2'b0	0	<i>Kernel addressing mode</i> . In this mode, a TLB miss goes to the TLB Refill Handler.
x	x	1	x	0	<i>Kernel addressing mode</i> . In this mode, a TLB miss goes to the TLB Refill Handler.
x	1	x	x	0	<i>Kernel addressing mode</i> . In this mode, a TLB miss goes to the general exception handler as opposed to the TLB Refill handler.
x	x	x	x	1	<i>Debug Mode</i> . In debug mode, the processor has full access to all resources that are available in Kernel Mode operation, in addition to those provided by EJTAG.

Figure 2.10 shows the format of the Status Register; Table 2.13 describes the Status register fields.

**Figure 2.10 Status Register Format**



**Table 2.13 Field Descriptions for Status Register**

Name	Bit(s)	Description	Read/Write	Reset State
CU3	31	Coprocessor 3 Usable. Reserved. This is a per-VPE view of the TCStatus <sub>TCU3</sub> per-TC field.	R	0
CU2	30	Coprocessor 2 Usable. Controls access to coprocessor 2.  0: Access not allowed. 1: Access allowed.  CU2 is reserved for a customer's coprocessor. This is a per-VPE view of the TCStatus <sub>TCU2</sub> per-TC field.	R	0

**Table 2.13 Field Descriptions for Status Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
CU1	29	<p>Coprocessor 1 Usable. Controls access to coprocessor 1.</p> <p>0: Access not allowed. 1: Access allowed.</p> <p>CU1 is most often used for a floating-point unit. When no coprocessor 1 is present, this bit is read-only and reads zero. This is a per-VPE view of the TCStatus<sub>TCU1</sub> per-TC field.</p>	R/W	Undefined
CU0	28	<p>Coprocessor 0 accessible in User Mode. This bit controls user mode access to coprocessor 0.</p> <p>0: Access not allowed. 1: Access allowed.</p> <p>Coprocessor 0 is always usable when the processor is running in Kernel or Debug Mode, regardless of the state of the CU0 bit.</p> <p>Setting Status<sub>CU0</sub> to 1 has the effect of allowing privileged instructions to execute in user mode, although this is not something a secure OS is likely to allow.</p>	R/W	Undefined
RP	27	<p>Reduced Power. Enable/disable reduced power mode.</p> <p>0: Disable reduced power mode. 1: Enable reduced power mode.</p> <p>The state of the RP bit is visible on the core's external interface signal <i>SI_RP</i>.</p>	R/W	0
FR	26	<p>Floating Register. This bit is used to control the floating-point register mode for 64-bit floating point units:</p> <p>0: Floating point registers can contain any 32-bit data type. 64-bit data types are stored in even-odd pairs of registers. 1: Floating point registers can contain any datatype.</p> <p>This bit must be ignored on writes and read as zero under the following conditions</p> <ul style="list-style-type: none"> <li>• No floating point unit is implemented</li> <li>• 64-bit floating point unit is not implemented</li> </ul> <p>If the interAptiv core is equipped with an optional FPU, set this bit to 0 for MIPS I compatibility mode, which allows for 16 real FP registers, with 16 odd FP register numbers reserved for access to the high-order bits of double-precision values. This is a per-VPE view of the TCStatus<sub>TFR</sub> per-TC field.</p>	R/W	0
RE	25	<p>Reverse Endian. Enables Reverse endianness for instructions that execute in User mode. This bit is always 0 as this feature is not supported in the interAptiv core.</p>	R	0

**Table 2.13 Field Descriptions for Status Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
MX	24	<p>MIPS DSP Extension. Enables access to DSP ASE resources.</p> <p>0: Access not allowed. 1: Access allowed.</p> <p>An attempt to execute any DSP ASE instruction before when this bit is 0 will cause a <i>DSP State Disabled</i> exception (see section 14.3 Software Detection of the DSP ASE Revision 2).</p> <p>This bit is read-only and defaults to 0 if the DSP ASE option is not selected during IP configuration. The bit is read-write and defaults to an Undefined state if the DSP ASE option is selected during IP configuration.</p> <p>This is a per-VPE view of the TCStatus.TMX per-TC field.</p>	<p>R (no DSP)</p> <p>R/W (DSP)</p>	<p>0 (no DSP)</p> <p>Undefined (DSP)</p>
0	23	Reserved. Write as zero. Ignored on reads.	R	0
BEV	22	<p>Boot Exception Vector. Controls the location of exception vectors:</p> <p>0: Normal. Refer to the EBase register for more information. 1: Bootstrap</p> <p>When set to 1, all exception entry points are relocated to near the reset start address.</p>	R/W	1
TS	21	<p>TLB shutdown. Normally this bit would indicate that a machine check exception was taken due to a TLBWI or TLBWR that would have created conflicting TLB entries. However, with the MIPS MT ASE, multiple writes are not an error condition and the conflicting TLB write instruction is silently dropped without a machine check exception. This bit will always be 0.</p>	R	0
SR	20	Soft Reset. The interAptiv core only supports a full external reset, so this bit is not used and always reads zero.	R	0
NMI	19	<p>Indicates that the entry through the reset exception vector was due to an NMI.</p> <p>0: Not NMI (reset) 1: NMI</p> <p>Software can only write a 0 to this bit to clear it and cannot force a 0 to 1 transition.</p>	R/W0	1 for NMI 0 otherwise
0	18	Reserved. Write as zero. Ignored on reads.	R	0
CEE	17	<p>CorExtend Enable. Enable/disable CorExtend User Defined Instructions (UDIs).</p> <p>0: Disable CorExtend block 1: Enable CorExtend block</p> <p>The presence of the CorExtend extension is indicated in ConfigUDI, which is set when the core is configured. This bit is reserved if CorExtend is not present.</p> <p>This is a per-VPE view of the TCStatus<sub>TCEE</sub> per-TC field.</p>	R/W	Undefined
0	16	Reserved. Write as zero. Ignored on reads.	R	0

**Table 2.13 Field Descriptions for Status Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State						
IM7:IM2	15:10	<p>Interrupt Mask: Controls the enabling of each of the hardware interrupts.</p> <p>An interrupt is taken if interrupts are enabled and the corresponding bits are set in both the Interrupt Mask field of the <b>Status</b> register and the Interrupt Pending field of the <b>Cause</b> register and the <i>IE</i> bit is set in the <b>Status</b> register.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt request disabled</td> </tr> <tr> <td>1</td> <td>Interrupt request enabled</td> </tr> </tbody> </table> <p>In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled (<math>\text{Config3}_{VEIC} = 1</math>), these bits take on a different meaning and are interpreted as the <b>IPL</b> field, described below.</p>	Encoding	Meaning	0	Interrupt request disabled	1	Interrupt request enabled	R/W	Undefined
Encoding	Meaning									
0	Interrupt request disabled									
1	Interrupt request enabled									
IPL	15:10	<p>Interrupt Priority Level: When EIC is enabled, this field is the encoded (0:63) value of the current IPL. An interrupt will be signaled only if the requested IPL is higher than this value.</p> <p>If EIC interrupt mode is not enabled (<math>\text{Config3}_{VEIC} = 0</math>), these bits take on a different meaning and are interpreted as the <b>IM7:IM2</b> bits, described above.</p>	R/W	Undefined						
IM1:IM0	9:8	<p>Interrupt Mask: Controls the enabling of each of the software interrupts.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt request disabled</td> </tr> <tr> <td>1</td> <td>Interrupt request enabled</td> </tr> </tbody> </table> <p>In EIC interrupt mode, the processor sends the state of the software interrupt requests (<b>Cause</b>IP1:IP0), to the external interrupt controller, where it prioritizes these interrupts in a system-dependent way with other hardware interrupts. This allows the interrupt controller to be more specific or more general as a function of the system environment and needs.</p> <p>When EIC interrupt mode is enabled (<math>\text{Config3}_{VEIC} = 1</math>), these bits are writable, but have no effect on the interrupt system.</p>	Encoding	Meaning	0	Interrupt request disabled	1	Interrupt request enabled	R/W	Undefined
Encoding	Meaning									
0	Interrupt request disabled									
1	Interrupt request enabled									
0	7:5	Reserved. Write as zero. Ignored on reads.	R	0						
KSU	4:3	<p>These bits denote the processor's operating mode.</p> <p>2'b00: Kernel Mode</p> <p>2'b01: Supervisor Mode</p> <p>2'b10: User Mode.</p> <p>Note that the processor can also be in Kernel mode if <b>ERL</b> or <b>EXL</b> is set, regardless of the state of these bits.</p> <p>This is a per-VPE view of the <b>TCStatus</b><sub>TKSU</sub> per-TC field. Note that supervisor mode is not supported when the core is configured with an MPU.</p>	R/W	Undefined						


**Table 2.13 Field Descriptions for Status Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
ERL	2	<p>Error Level; Set by the processor when a Reset, NMI, or Cache Error exception is taken.</p> <p>0: Normal level 1: Error level</p> <p>When ERL is set:</p> <ul style="list-style-type: none"> <li>• The processor is running in kernel mode</li> <li>• Interrupts are disabled</li> <li>• The ERET instruction will use the return address held in ErrorEPC instead of EPC</li> <li>• When ERL = 1 in the Status register, the segment kuseg (legacy) or xksego (EVA) is treated as an unmapped and uncached address space. While in this setting, the kuseg virtual address maps directly to the same physical address, and does not include the ASID field.</li> <li>• The lower 2<sup>29</sup> bytes of kuseg are treated as an unmapped and uncached region. This allows main memory to be accessed in the presence of cache errors. The operation of the processor is <b>UNDEFINED</b> if the ERL bit is set while the processor is executing instructions from kuseg.</li> </ul>	R/W	1
EXL	1	<p>Exception Level; Set by the processor when any exception other than Reset, Cache Error, or NMI exception is taken.</p> <p>0: Normal level 1: Exception level</p> <p>When EXL is set:</p> <ul style="list-style-type: none"> <li>• The processor is running in Kernel Mode.</li> <li>• Hardware and software interrupts are disabled.</li> <li>• TLB Refill exceptions use the general exception vector instead of the TLB Refill vector.</li> <li>• EPC, Cause<sub>BD</sub> and SRSCtl are not be updated if another exception is taken.</li> </ul> <p>When an exception occurs and EXL is set, a nested TLB Refill exception is sent to the general exception handler (rather than to its dedicated handler) and the values in EPC, Cause<sub>BD</sub> and SRSCtl are not overwritten. The result is that, after returning from the second exception, the processor jumps back to the code that was executing before the first exception occurred.</p>	R/W	Undefined
IE	0	<p>Interrupt Enable. Acts as the master enable for software and hardware interrupts.</p> <p>0: Interrupts are disabled 1: Interrupts are enabled</p> <p>This bit can be written using the <b>di/ei</b> instructions.</p>	R/W	Undefined

### 2.2.1.11 Interrupt Control — IntCtl (CP0 Register 12, Select 1)

The IntCtl register controls the interrupt capabilities of the *interAptiv* core, including vectored interrupts and support for an external interrupt controller.

**Figure 2.11 IntCtl Register Format**



**Table 2.14 Field Descriptions for IntCtl Register**

Name	Bit(s)	Description	Read/Write	Reset State
IPTI	31:29	<p>For <i>Interrupt Compatibility</i> and <i>Vectored Interrupt</i> modes, this field specifies the IP number to which the Timer Interrupt request is merged, and allows software to determine whether to consider <b>Cause<sub>T</sub></b> for a potential interrupt. This field is encoded as shown in <a href="#">Table 2.15, "Encoding of IPTI, IPPCI, and IPFDCI Fields"</a>.</p> <p>The value of this bit is set by the static input, <i>SI_IPTI</i>[2:0]. This allows external logic to communicate the specific <i>SI_Int</i> hardware interrupt pin to which the <i>SI_TimerInt</i> signal is attached.</p> <p>The value of this field is not meaningful if External Interrupt Controller Mode is enabled. The external interrupt controller is expected to provide this information for that interrupt mode.</p>	R	Externally Set
IPPCI	28:26	<p>For <i>Interrupt Compatibility</i> and <i>Vectored Interrupt</i> modes, this field specifies the IP number to which the Performance Counter Interrupt request is merged, and allows software to determine whether to consider <b>Cause<sub>PC</sub></b> for a potential interrupt. This field is encoded as shown in <a href="#">Table 2.15, "Encoding of IPTI, IPPCI, and IPFDCI Fields"</a>.</p> <p>The value of this bit is set by the static input <i>SI_IPPCI</i>[2:0]. This allows external logic to communicate the specific <i>SI_Int</i> hardware interrupt pin to which the <i>SI_PCInt</i> signal is attached.</p> <p>The value of this field is not meaningful if External Interrupt Controller Mode is enabled. The external interrupt controller is expected to provide this information for that interrupt mode.</p>	R	Externally Set
IPFDCI	25:23	<p>For <i>Interrupt Compatibility</i> and <i>Vectored Interrupt</i> modes, this field specifies the IP number to which the Fast Debug Channel Interrupt request is merged, and allows software to determine whether to consider <b>Cause<sub>FDC</sub></b> for a potential interrupt. This field is encoded as shown in <a href="#">Table 2.15, "Encoding of IPTI, IPPCI, and IPFDCI Fields"</a>.</p> <p>The value of this bit is set by the static input, <i>SI_IPFDCI</i>[2:0]. This allows external logic to communicate the specific <i>SI_Int</i> hardware interrupt pin to which the <i>SI_FDCInt</i> signal is attached.</p> <p>The value of this field is not meaningful if External Interrupt Controller Mode is enabled. The external interrupt controller is expected to provide this information for that interrupt mode.</p>	R	Externally Set
0	22:10	Reserved. Write as zero. Ignored on reads.	R	0

**Table 2.14 Field Descriptions for IntCtl Register**

Name	Bit(s)	Description	Read/Write	Reset State																					
VS	9:5	<p>Vector Spacing. If vectored interrupts are implemented (as denoted by Config3<sub>VIInt</sub> or Config3<sub>VEIC</sub>), this field specifies the spacing between vectored interrupts.</p> <table border="1"> <thead> <tr> <th>VS Field Encoding</th> <th>Spacing Between Vectors (hex)</th> <th>Spacing Between Vectors (decimal)</th> </tr> </thead> <tbody> <tr> <td>0x00</td> <td>0x000</td> <td>0</td> </tr> <tr> <td>0x01</td> <td>0x020</td> <td>32</td> </tr> <tr> <td>0x02</td> <td>0x040</td> <td>64</td> </tr> <tr> <td>0x04</td> <td>0x080</td> <td>128</td> </tr> <tr> <td>0x08</td> <td>0x100</td> <td>256</td> </tr> <tr> <td>0x10</td> <td>0x200</td> <td>512</td> </tr> </tbody> </table> <p>All other values are reserved. The operation of the processor is <b>UNDEFINED</b> if a reserved value is written to this field.</p>	VS Field Encoding	Spacing Between Vectors (hex)	Spacing Between Vectors (decimal)	0x00	0x000	0	0x01	0x020	32	0x02	0x040	64	0x04	0x080	128	0x08	0x100	256	0x10	0x200	512	R/W	0
VS Field Encoding	Spacing Between Vectors (hex)	Spacing Between Vectors (decimal)																							
0x00	0x000	0																							
0x01	0x020	32																							
0x02	0x040	64																							
0x04	0x080	128																							
0x08	0x100	256																							
0x10	0x200	512																							
0	4:0	Reserved. Write as zero. Ignored on reads.	R	0																					

**Table 2.15 Encoding of IPTI, IPPCI, and IPFDCI Fields**

Encoding	IP bit	Hardware Interrupt Source
0	0	Reserved
1	1	Reserved
2	2	HW0
3	3	HW1
4	4	HW2
5	5	HW3
6	6	HW4
7	7	HW5



## 2.2.2 TLB Management Registers

This section contains the following TLB management registers. Note that all of the registers in section, except BadVAddr, are only implemented if a TLB-based MMU is supported. If an MPU is implemented, only the BadVAddr register is used.

- Section 2.2.2.1, "Index (CP0 Register 0, Select 0)" on page 85
- Section 2.2.2.2, "Random (CP0 Register 1, Select 0)" on page 86
- Section 2.2.2.3, "EntryLo0 - EntryLo1 (CP0 Registers 2 and 3, Select 0)" on page 87
- Section 2.2.2.4, "EntryHi (CP0 Register 10, Select 0)" on page 88
- Section 2.2.2.5, "Context (CP0 Register 4, Select 0)" on page 90
- Section 2.2.2.6, "Context Configuration — ContextConfig (CP0 Register 4, Select 1)" on page 90
- Section 2.2.2.7, "PageMask (CP0 Register 5, Select 0)" on page 91
- Section 2.2.2.8, "Page Granularity — PageGrain (CP0 Register 5, Select 1)" on page 92
- Section 2.2.2.9, "Wired (CP0 Register 6, Select 0)" on page 93
- Section 2.2.2.10, "Bad Virtual Address — BadVAddr (CP0 Register 8, Select 0)" on page 94

### 2.2.2.1 Index (CP0 Register 0, Select 0)

Index is used as the TLB index when reading or writing the TLB with **TLBR**/**TLBWI**/**TLBINV**/**TLBINVF** respectively. It is also set by a TLB probe (**TLBP**) instruction to return the location of an address match in the TLB.

During execution of a **TLBR** instruction, the Index field that was previously written by software or by a **TLBP** instruction is used to indicate the TLB entry to be read. Hardware then uses this information to perform the read operation.

During execution of a **TLBWI**, **TLBINV**, or **TLBINVF** instruction, the Index field that was previously written by software or by a **TLBP** instruction is used to indicate the TLB entry to be written or invalidated. Hardware then uses this information to perform the respective write or invalidate operation.

Prior to executing a **TLBP** instruction, the VPN to be searched should have been written to the VPN2 field in the EntryHi register. During the **TLBP** instruction, hardware searches the TLB array for a match to the VPN stored in the EntryHi register. If a match is found, hardware writes the index into the *Index* field of this register.

The P bit of this register is set by hardware to indicate that a match was not found. If this bit is not set, software can then read the corresponding index from this register.

The operation of the processor is **UNDEFINED** if a value greater than or equal to the number of TLB entries is written to the Index register.

**Figure 2.12 Index Register Format**



**Table 2.16 Field Descriptions for Index Register**

Name	Bit(s)	Description	Read/Write	Reset State
P	31	Probe Failure. This bit is automatically set when a <b>TLBP</b> search of the TLB fails to find a matching entry.  Software can set to 1 to avoid locking up an entry when the TLB is shared. If 0, TLBWR on the other VPE will skip the selected Index value to allow refills on the other VPE to occur at the same time as TLB maintenance on this one.	R	Undefined
0	30:6	Must be written as zero; returns zero on reads.	0	0
Index	5:0	Index to the TLB entry affected by the <b>TLBR</b> and <b>TLBWI</b> instructions.  For 16 or 32 entry TLBs, behavior is undefined if index points to a non-existent entry.	R/W	Undefined

**2.2.2.2 Random (CP0 Register 1, Select 0)**

The Random register is a read-only register whose value is used to index the TLB during a **TLBWR** instruction. It provides a quick way of replacing a TLB entry at random. As a result, it will not take values less than the value programmed in the *Wired* register. The *Random* register employs a pseudo-random least-recently-used (LRU) algorithm that ensures that no wired entries are selected, Only those LRU entries that are not in the *Wired* register are targeted for replacement. The contents of the *Random* register are modified after a TLB write, or on a write to the *Wired* register.

The value of the register varies between an upper and lower bound as follow:

- A lower bound is set by the number of TLB entries reserved for exclusive use by the operating system (the contents of the *Wired* register). The entry indexed by the *Wired* register is the first entry available to be written by a TLB Write Random operation.
- An upper bound is set by the total number of TLB entries minus 1.

The Random register is decremented by one almost every clock, wrapping after the value in the *Wired* register is reached. To enhance the level of randomness and reduce the possibility of a live lock condition, an LFSR register is used which prevents the decrement pseudo-randomly.

**Figure 2.13 Random Register Format**



**Table 2.17 Field Descriptions for Random Register**

Name	Bit(s)	Description	Read/Write	Reset State
0	31:6	Must be written as zero; returns zero on reads.	0	0
Random	5:0	This field cycles "randomly" through the potential indices of the TLB, so its length varies with the TLB size. It is a pseudo-least-recently-used TLB index.	R	TLB Entries — 1


### 2.2.2.3 EntryLo0 - EntryLo1 (CP0 Registers 2 and 3, Select 0)

The pair of *EntryLo* registers act as the interface between the TLB and the **TLBP**, **TLBR**, **TLBWI**, and **TLBWR** instructions. These registers store the contents of a TLB entry. Each entry maps a pair of pages. The EntryLo0 and EntryLo1 register store even and odd numbered virtual pages respectively. These registers are read during a **TLBWR** or **TBLWI** instruction, and written by a **TLBR** instruction. They are not used for any other purpose.

Software may determine the value of PABITS by writing all ones to the EntryLo0 or EntryLo1 registers and reading the value back. Bits read as “1” from the PFN field allow software to determine the boundary between the PFN and Fill fields to calculate the value of PABITS.

The contents of the EntryLo0 and EntryLo1 registers are not defined after an address error exception and some fields may be modified by hardware during the address error exception sequence. Software writes of the EntryHi register (via MTC0) do not cause the implicit update of address-related fields in the BadVAddr or Context registers.

**Figure 2.14 EntryLo0 and EntryLo1 Register Format**



**Table 2.18 Field Descriptions for EntryLo0 and EntryLo1 Registers**

Name	Bit(s)	Description	Read/Write	Reset State
RI	31	Read Inhibit. If this bit is set in a TLB entry, any attempt (other than a MIPS16 PC-relative load) to read data on the virtual page causes either a TLB Invalid or a TLBRI exception, even if the V (Valid) bit is set. The RI bit is writable only if the RIE bit of the PageGrain register is set. For more information, refer to <a href="#">Section 2.2.2.8, "Page Granularity — PageGrain (CP0 Register 5, Select 1)"</a> .  If the RIE bit of the PageGrain register is not set, the RI bit of Entry 0 and Entry 1 are set to zero on any write to the register, regardless of the value written.	R/W	Undefined
XI	30	Execute Inhibit. If this bit is set in a TLB entry, any attempt to fetch an instruction or to load MIPS16 PC-relative data from the virtual page causes a TLB Invalid or a TLBXI exception, even if the V (Valid) bit is set. The XI bit is writable only if the XIE bit of the PageGrain register is set. For more information, refer to <a href="#">Section 2.2.2.8, "Page Granularity — PageGrain (CP0 Register 5, Select 1)"</a> .  If the XIE bit of the PageGrain register not set, the XI bit of TLB Entry 0 - 1 is set to zero on any write to the register, regardless of the value written.	R/W	Undefined
U	29:26	The upper 4 bits of the PFN cannot be written by software and will return 0 on reads.	R/W	Undefined
PFN	25:6	The "Physical Frame Number" represents bits 31:12 of the physical address. The 20 bits of PFN, together with 12 bits of in-page address, make up a 32-bit physical address. The MIPS32® Architecture permits the PFN to be as large as 24 bits. The interAptiv core supports a 32-bit physical address bus.	R/W	Undefined
C	5:3	Coherency attribute of the page. See <a href="#">Table 2.19</a> .	R/W	Undefined
D	2	The "Dirty" flag. Indicates that the page has been written, and/or is writable. If this bit is a one, stores to the page are permitted. If this bit is a zero, stores to the page cause a TLB Modified exception.  Software can use this bit to track pages that have been written to. When a page is first mapped, this bit should be cleared. It is set on the first write that causes an exception.	R/W	Undefined

**Table 2.18 Field Descriptions for EntryLo0 and EntryLo1 Registers**

Name	Bit(s)	Description	Read/Write	Reset State
V	1	The “Valid” flag. Indicates that the TLB entry, and thus the virtual page mapping, are valid. If this bit is a set, accesses to the page are permitted. If this bit is a zero, accesses to the page cause a <i>TLB Invalid</i> exception.  This bit can be used to make just one of a pair of pages valid.	R/W	Undefined
G	0	The “Global” bit. On a TLB write, the logical AND of the G bits in both the Entry 0 and Entry 1 registers become the G bit in the TLB entry. If the TLB entry G bit is a one, then the ASID comparisons are ignored during TLB matches. On a read from a TLB entry, the G bits of both Entry 0 and Entry 1 reflect the state of the TLB G bit.	R/W	Undefined

**Table 2.19 Cache Coherency Attributes Encoding of the C Field**

C[5:3] / K0[2:0] <sup>1</sup>	Name	Cache Coherency Attribute
0	—	Reserved
1	—	Reserved
2	UC	Uncached, non-coherent
3	WB	Cacheable, non-coherent, write-back, write allocate
4	CWBE	Cacheable, coherent, write-back, write-allocate, read misses request Exclusive
5	CWB	Cacheable, coherent, write-back, write-allocate, read misses request Shared
6	—	Reserved
7	UCA	Uncached Accelerated, non-coherent

1. State of the K0 field at bits 2:0 of the Config register. See [Section 2.2.1.1 “Device Configuration — Config \(CP0 Register 16, Select 0\)”](#)

#### 2.2.2.4 EntryHi (CP0 Register 10, Select 0)

The EntryHi register contains the upper portion of the virtual address match information used for TLB read, write, and access operations. The remaining information is stored in the EntryLo0 and EntryLo1 registers described in [Section 2.2.2.3 “EntryLo0 - EntryLo1 \(CP0 Registers 2 and 3, Select 0\)”](#).


A TLB exception (TLB Refill, TLB Invalid, or TLB Modified) causes bits VA<sub>31:13</sub> of the virtual address to be written into the VPN2 field of the EntryHi register. A TLBR instruction writes the EntryHi register with the corresponding fields from the selected TLB entry. The ASID field is written by software with the current address space identifier value and is used during the TLB comparison process to determine TLB match.

Because the ASID and EHINV fields are overwritten by a TLBR instruction, software must save and restore the value of ASID around use of the TLBR. This is especially important in TLB Invalid and TLB Modified exceptions, and in other memory management software.

The VPN2 field of the EntryHi register is not defined after an address error exception and this field may be modified by hardware during the address error exception sequence. Software writes of the EntryHi register (via MTC0) do not cause the implicit write of address-related fields in the BadVAddr and Context registers.

The EntryHiEHINV field has been added to support explicit invalidation of TLB entries via the **TLBWI** instruction. When EntryHiEHINV = 1, the **TLBWI** instruction acts as a TLB invalidate operation, setting the hardware valid bit associated with a TLB entry to the invalid state. When EntryHiEHINV = 1, only the *Index* register is required to be valid. Behavior of the **TLBWR** instruction is unmodified by EntryHiEHINV. The **TLBR** instruction copies the EHINV bit from the TLB Entry to EntryHiEHINV. Note that execution of the **TLBP** instruction does not change this value.

**Figure 2.15 EntryHi Register Format**



**Table 2.20 Field Descriptions for EntryHi Register**

Name	Bit(s)	Description	Read/Write	Reset State
VPN2	31:13	EntryHi <sub>VPN2</sub> is the virtual address to be matched on a <b>TLBP</b> . This field consists of VA <sub>31:13</sub> of the virtual address (virtual page number / 2). It is also the virtual address to be written into the TLB on a <b>TLBWI</b> and <b>TLBWR</b> , and the destination of the virtual address on a <b>TLBR</b> .  On a TLB-related exception, the VPN2 field is automatically set to the virtual address that was being translated when the exception occurred.  This field is written by software before a <b>TLBP</b> or <b>TLBWI</b> and written by hardware in all other cases.	R/W	Undefined
0	12:11	Reserved. Write as zero. Ignored on reads.	R	0
EHINV	10	TLBWI invalidate enable. When this bit is set, the TLBWI instruction acts as a TLB invalidate operation, setting the hardware valid bit associated with the TLB entry to the invalid state. When this bit is set, the PageMask and EntryLo0/EntryLo1 registers do not need to be valid. Only the <i>Index</i> register is required to be valid.  This bit is ignored on a <b>TLBWR</b> instruction.	R/W	0
0	9:8	Reserved. Write as zero. Ignored on reads.	R	0
ASID	7:0	Address space identifier. This field is used to stage data to and from the TLB, but in normal running software it's also the source of the current "ASID" value, used to extend the virtual address and help to map address translations for the current process.  This field is written by hardware on a TLB read and by software to establish the current ASID value for TLB write and against which TLB references match each entry's TLB ASID field.  This field supports up to 256 unique ASID values, consisting of a virtual tag that is in addition to the 32-bit address.  This field is per-TC field visible in TCStatus <sub>TASID</sub> .	R/W	0

### 2.2.2.5 Context (CP0 Register 4, Select 0)


The Context register is a read/write register containing a pointer to an entry in the page table entry (PTE) array. This array is an operating system data structure that stores virtual-to-physical translations. During a TLB miss, the operating system loads the TLB with the missing translation from the PTE array. The Context register duplicates some of the information provided in the BadVAddr register but is organized in such a way that the operating system can directly reference an 8-byte page table entry (PTE) in memory.

The BadVPN2 field of the Context register is not defined after an address error exception, and this field may be modified by hardware during the address error exception sequence.

The pointer implemented by the Context register can point to any power-of-two-sized PTE structure within memory. This allows the TLB refill handler to use the pointer without additional shifting and masking steps. For example, if the low-order bit of the PTEBase field is 20, the page table entry (PTE) structure occurs on a 1M boundary. If the low-order bit is 21, PTE structure occurs on a 2M boundary, etc.

Figure 2.16 shows the format of the Context Register; Table 2.21 describes the Context register fields.

**Figure 2.16 Context Register Format**



**Table 2.21 Context Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
PTEBase	31:23	This field is for use by the operating system and is normally written with a value that allows the operating system to use the Context Register as a pointer to an array of data structures in memory corresponding to the address region that contains the virtual address which caused the exception.	R/W	Undefined
BadVPN2	22:4	This field is written by hardware on a TLB exception. It contains bits VA <sub>31:13</sub> of the virtual address that caused the exception.	R	Undefined
0	3:0	Ignored on write; returns zero on read.	R	0


### 2.2.2.6 Context Configuration — ContextConfig (CP0 Register 4, Select 1)

The ContextConfig register defines the bits of the Context register into which the high order bits of the virtual address causing a TLB exception will be written, and how many bits of that virtual address will be extracted. Bits above the selected BadVPN2 field of the Context register are read/write to software and serve as the PTEBase field. Bits below the selected BadVPN2 field of the Context register serve as the PTEBaseLow field.

Software writes a set of contiguous ones to the VirtualIndex field of the ContextConfig register. Hardware then determines which bits of this register are high and low. The highest order bit that is a logic ‘1’ serves as the MSB of the BadVPN2 field of the Context register. The lowest order bit that is a logic ‘1’ serves as the LSB of the BadVPN2 field of the Context register. A value of all zero’s in the VirtualIndex field means that the full 32 bits of the Context register are R/W for software and are unaffected by TLB exceptions.

Figure 2.17 shows the formats of the ContextConfig register; Table 2.22 describes the ContextConfig register fields.

**Figure 2.17 ContextConfig Register Format**



**Table 2.22 ContextConfig Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:23	Ignored on write; returns zero on read.	R	0x00
VirtualIndex	22:2	A mask of 0 to 21 contiguous 1 bits in this field causes the corresponding bits of the Context register to be written with the high-order bits of the virtual address causing a TLB exception.  Behavior of the processor is <b>UNDEFINED</b> if non-contiguous 1 bits are written into the register field. Note that it is the responsibility of software to ensure that this field is written with contiguous ones because if non-contiguous 1 bits are written, no exception will be taken.	R/W	0x1F_FFFC
0	1:0	Ignored on write; returns zero on read.	R	0


**2.2.2.7 PageMask (CP0 Register 5, Select 0)**

Every TLB entry has an independent virtual-address mask that allows it to ignore some address bits when deciding to match. By selectively ignoring lower page addresses, the entry can be made to match all the addresses in a "page" larger than 4KB.

Software can determine the maximum page size supported by writing all ones to the PageMask register, then reading the value back. If a pair of bits reads back as ones, the processor implements that page size. Note that the bits are read in pairs, so bits 14:13 are read first and can have only a value of 00 or 11. If they are both 11, bits 16:15 are read, and so on.

The operation of the processor is **UNDEFINED** if software loads the Mask field with a value other than one of those listed in Table 2.24, even if the hardware returns a different value on read. Hardware may depend on this requirement in implementing hardware structures.

**Figure 2.18 PageMask Register Format**



**Table 2.23 Field Descriptions for PageMask Register**

Name	Bit(s)	Description	Read/ Write	Reset State
0	31:29	Ignored on write; returns zero on read.	R	0
Mask	28:13	Acts as a kind of backward mask, in that a 1 bit means "don't compare this address bit when matching this address". However, only a restricted range of PageMask values are legal (i.e., with "1"s filling the PageMask <sub>Mask</sub> field from low bits upward, two at a time)	R/W	Undefined

**Table 2.23 Field Descriptions for PageMask Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
0	12:0	Ignored on write; returns zero on read.	R	0

**Table 2.24 PageMask Register Values**

PageMask Register Value	Size of Each Output Page
0x0000.E000	4 Kbytes
0x0000.6000	16 Kbytes
0x0001.E000	64 Kbytes
0x0007.E000	256 Kbytes
0x001F.E000	1 Mbyte
0x007F.E000	4 Mbytes
0x01FF.E000	16 Mbytes
0x07FF.E000	64 Mbytes
0x1FFF.E000	256 Mbytes

Software may determine which page sizes are supported by writing all ones to the PageMask register, then reading the value back. If a pair of bits reads back as ones, the processor implements that page size. The operation of the processor is **UNDEFINED** if software loads the *Mask* field with a value other than one of those listed in Table 2.24, even if the hardware returns a different value on read. Hardware may depend on this requirement in implementing hardware structures.

### 2.2.2.8 Page Granularity — PageGrain (CP0 Register 5, Select 1)

The PageGrain register is a read/write register used for XI/RI TLB protection bits. The PageGrain register is present in Release 3 (and subsequent releases) of the architecture.

Figure 2.19 shows the format of the PageGrain register; Table 2.25 describes the PageGrain register fields.

**Figure 2.19 PageGrain Register Format**



**Table 2.25 Field Descriptions for PageGrain Register**

Name	Bit(s)	Description	Read/Write	Reset State
RIE	31	Read inhibit enable. 0: RI bit of the Entry0 and Entry1 registers is disabled and not writeable by software. 1: RI bit of the Entry0 and Entry1 registers is enabled.	R/W	0



**Table 2.25 Field Descriptions for PageGrain Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
XIE	30	Execute inhibit enable. 0: XI bit of the Entry0 and EntryI registers is disabled and not writeable by software. 1: XI bit of the Entry0 and EntryI registers is enabled.	R/W	0
0	29	Reserved. Ignored on write; returns zero on read.	R	0
ESP	28	This bit is always 0 as 1K pages are not supported. This bit must be written with 0.	R	0
IEC	27	Enables unique exception codes for the Read-Inhibit and Execute-Inhibit exceptions. 0: Read-Inhibit and Execute-Inhibit exceptions both use the TLBL exception code. 1: Read-Inhibit exceptions use the TLBRI exception code. Execute-Inhibit exceptions use the TLBXI exception code.	R/W	0
0	26:13	Reserved. Ignored on write; returns zero on read.	R	0
ASE	12:8	Ignored on write; returns zero on read.	R	0
0	7:0	Reserved. Ignored on write; returns zero on read.	R	0

### 2.2.2.9 Wired (CP0 Register 6, Select 0)

The Wired register is a read/write register that specifies the boundary between the wired and random entries in the TLB as shown in Figure 2.26. Wired entries are fixed, non-replaceable entries that cannot be overwritten by a TLBWR instruction. Wired entries can be overwritten by a TLBWI instruction.

Note that wired entries in the TLB must be contiguous and start from 0. For example, if the Wired field of this register contains a value of 5, this indicates that entries 4, 3, 2, 1, and 0 of the TLB are wired.

The Wired register is reset to zero by a Reset exception. Writing the Wired register may cause the Random register to change state.

The operation of the processor is undefined if a value greater than or equal to the number of TLB entries is written to the Wired register. Wired can be set to a non-zero value to prevent the random replacement of up to 63 TLB pages.

**Figure 2.20 Wired Register Format**



**Table 2.26 Field Descriptions for Wired Register**

Name	Bit(s)	Description	Read/Write	Reset State
0	31:6	Ignored on write; returns zero on read.	R	0

**Table 2.26 Field Descriptions for Wired Register**

Name	Bit(s)	Description	Read/Write	Reset State
Wired	5:0	<p>TLB wired boundary.</p> <p>For 16 and 32 entry TLBs, behavior is undefined if value is set to a value larger than last TLB entry.</p> <p>This field is encoded as follows:</p> <p>0x00: 0 TLB entries are hardwired            0x01: 1 TLB entry is hardwired            0x02: 2 TLB entries are hardwired            .....            0x3F: 63 TLB entries are hardwired</p> <p>These entries become a good place for an OS to keep translations which must never cause a TLB translation-not-present exception.</p>	R/W	0

**2.2.2.10 Bad Virtual Address — BadVAddr (CP0 Register 8, Select 0)**


The BadVAddr register is a read-only register that captures the most recent virtual address that caused one of the following exceptions:

- Address error (AdEL or AdES)
- TLB Refill
- TLB Invalid (TLBL, TLBS)
- TLB Read Inhibit (TLBRI)
- TLB Execute Inhibit (TLBXI)
- TLB Modified
- MPU Protection

The BadVAddr register does not capture address information for cache or bus errors, since they are not addressing errors.

There is more information about this register in the notes to the Cause<sub>ExcCode</sub> field.

**Figure 2.21 BadVAddr Register Format**



**Table 2.27 BadVAddr Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bits			
BadVAddr	31:0	Bad virtual address. This register stores the virtual address that causes one of the exceptions listed above.	R	Undefined

## 2.2.3 Memory Segmentation Registers

This section contains the following memory segmentation registers. Note that all of the registers in section are only implemented if a TLB-based MMU is supported. If an MPU is implemented, none of these registers are used.

- [Section 2.2.3.1, "Segmentation Control 0 — SegCtl0 \(CP0 Register 5, Select 2\)" on page 95](#)
- [Section 2.2.3.2, "Segmentation Control 1 — SegCtl1 \(CP0 Register 5, Select 3\)" on page 96](#)
- [Section 2.2.3.3, "Segmentation Control 2 — SegCtl2 \(CP0 Register 5, Select 4\)" on page 97](#)

Programmable segmentation allows for the virtual address space segments to be programmed with different access modes and attributes. Control of the 4GB of virtual address space is divided into six segments that are controlled using three CP0 registers; SegCtl0 through SegCtl2. Each register has two 16-bit fields. Each field controls one of the six address segments as shown in [Table 2.28](#). For more information, refer to Section 2.6 of the MMU chapter of this manual.

**Table 2.28 Programmable Segmentation Register Interface**


Register	CP0 Location	Memory Segment	Register Bits	Virtual Address Space Controlled	Virtual Address Range (Hex)
SegCtl2	Register 5 Select 4	CFG5	31:16	0.0 GB to 1.0 GB	0x0000_0000 - 0x3FFF_FFFF
		CFG4	15:0	1.0 GB to 2.0 GB	0x4000_0000 - 0x7FFF_FFFF
SegCtl1	Register 5 Select 3	CFG3	31:16	2.0 GB to 2.5 GB	0x8000_0000 - 0x9FFF_FFFF
		CFG2	15:0	2.5 GB to 3.0 GB	0xA000_0000 - 0xBFFF_FFFF
SegCtl0	Register 5 Select 2	CFG1	31:16	3.0 GB to 3.5 GB	0xC000_0000 - 0xDFFF_FFFF
		CFG0	15:0	3.5 GB to 4.0 GB	0xE000_0000 - 0xFFFF_FFFF

### 2.2.3.1 Segmentation Control 0 — SegCtl0 (CP0 Register 5, Select 2)

The SegCtl0 register works in conjunction with the SegCtl1 and SegCtl2 registers to allow for configuration of the memory segmentation system. The address is split into the six segments defined in [Table 2.28](#).

[Figure 2.22](#) shows the format of the SegCtl0 Register. Note that the Config3<sub>SR</sub> bit must be set to enable this register.

**Figure 2.22 SegCtl0 Register Format (CP0 Register 5, Select 2)**



**Table 2.29 SegCtl0 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
CFG1_PA	31:25	Physical address bits 31:29 for segment 1. For use when unmapped. Bits 27:25 correspond to physical address bits 31:29. Bits 31:28 are reserved for future expansion. For more information, refer to <a href="#">Section 3.5.3.6 “Defining the Physical Address Range for Each Memory Segment”</a>	R/W	Configuration Dependent
0	24:23	Reserved.	RO	0
CFG1_AM	22:20	Configuration 1 access control mode. See <a href="#">Table 2.32</a> for encoding. For more information, refer to <a href="#">Section 3.5.3.5 “Setting the Access Control Mode”</a> .	R/W	Configuration Dependent
CFG1_EU	19	Error condition behavior. Configuration segment 1 becomes unmapped and uncached when $Status_{ERL} = 1$ .	R/W	Configuration Dependent
CFG1_C	18:16	Cache coherency attribute for segment 1. The encoding of the CFG1_C field is the same as the C field of the EntryLo0/EntryLo1 registers described in <a href="#">Section 2.2.2.3</a> . Refer to <a href="#">Table 2.19</a> for the encoding of this field. For more information, refer to <a href="#">Section 3.5.3.5 “Setting the Access Control Mode”</a> .	R/W	Configuration Dependent
CFG0_PA	15:9	Physical address bits 31:29 for segment 0. For use when unmapped. Bits 11:9 correspond to physical address bits 31:29 for segment 0. Bits 15:12 are reserved for future expansion. For more information, refer to <a href="#">Section 3.5.3.6 “Defining the Physical Address Range for Each Memory Segment”</a> .	R/W	Configuration Dependent
0	8:7	Reserved.	RO	0
CFG0_AM	6:4	Configuration 0 access control mode. See <a href="#">Table 2.32</a> for encoding. For more information, refer to <a href="#">Section 3.5.3.5 “Setting the Access Control Mode”</a> .	R/W	Configuration Dependent
CFG0_EU	3	Error condition behavior. Configuration segment 0 becomes unmapped and uncached when $Status_{ERL} = 1$ .	R/W	Configuration Dependent
CFG0_C	2:0	Cache coherency attribute for segment 0. The encoding of the CFG0_C field is the same as the C field of the EntryLo0/EntryLo1 registers described in <a href="#">Section 2.2.2.3</a> . Refer to <a href="#">Table 2.19</a> for the encoding of this field. For more information, refer to <a href="#">Section 3.5.3.5 “Setting the Access Control Mode”</a> .	R/W	Configuration Dependent

### 2.2.3.2 Segmentation Control 1 — SegCtl1 (CP0 Register 5, Select 3)

The SegCtl1 register works in conjunction with the SegCtl0 and SegCtl2 registers to allow for configuration of the memory segmentation system. The address is split into six segments defined in [Table 2.28](#).

Segmentation Control allows address-specific behaviors defined by the Privileged Resource Architecture to be modified or disabled.

[Figure 2.23](#) shows the format of the SegCtl1 Register. Note that the Config3<sub>SR</sub> bit must be set to enable this register. For more information on the reset states of these fields, refer to [Section 3.6.6 “Switching the Addressing Scheme from Legacy to EVA After Boot-up”](#).





Table 2.32 describes the access control modes specifiable in the CFG<sub>AM</sub> fields.

**Table 2.32 Segment Configuration Access Control Modes**

Mode		Action when referenced from Operating Mode			Description
		User mode	Supervisor mode	Kernel mode	
UK	000	Address Error	Address Error	Unmapped	Kernel-only unmapped region e.g. kseg0, kseg1
MK	001	Address Error	Address Error	Mapped	Kernel-only mapped region e.g. kseg3
MSK	010	Address Error	Mapped	Mapped	Supervisor and kernel mapped region e.g. ksseg, sseg
MUSK	011	Mapped	Mapped	Mapped	User, supervisor and kernel mapped region e.g. useg, kuseg, suseg
MUSUK	100	Mapped	Mapped	Unmapped	Used to implement a fully-mapped flat address space in user and supervisor modes, with unmapped regions which appear in kernel mode.
USK	101	Address Error	Unmapped	Unmapped	Supervisor and kernel unmapped region e.g. sseg in a fixed mapping TLB.
-	110	Undefined	Undefined	Undefined	Reserved
UUSK	111	Unmapped	Unmapped	Unmapped	Unrestricted unmapped region

Table 2.33 describes the state of each Segment Configuration register at reset in legacy mode.

**Table 2.33 Segment Configuration Reset States in Legacy Mode**

CFG	Segment	CFG <sub>AM</sub>	CFG <sub>PA</sub>	CFG <sub>C</sub>	CFG <sub>EU</sub>
0	kseg3	MK	Undefined	Undefined	0
1	ksseg, sseg	MSK	Undefined	Undefined	0
2	kseg1	UK	3'b000	2	0
3	kseg0	UK	3'b000	3	0
4	kuseg, suseg, useg	MUSK	3'b010	Undefined	1
5	kuseg, suseg, useg	MUSK	3'b000	Undefined	1

## 2.2.4 Exception Control Registers

This section contains the following exception control registers.


- [Section 2.2.4.1, "Cause \(CP0 Register 13, Select 0\)" on page 100](#)
- [Section 2.2.4.2, "Exception Program Counter — EPC \(CP0 Register 14, Select 0\)" on page 103](#)
- [Section 2.2.4.3, "Error Exception Program Counter — ErrorEPC \(CP0 Register 30, Select 0\)" on page 104](#)

Also refer to the Interrupt Control register in [Section 2.2.1.11, "Interrupt Control — IntCtl \(CP0 Register 12, Select 1\)" on page 82](#).

### 2.2.4.1 Cause (CP0 Register 13, Select 0)

The Cause register describes the cause of the most recent exception and controls software interrupt requests and the vector through which interrupts are dispatched. With the exception of the IP1:0, DC, IV, and WP fields, all fields in the Cause register are read-only. IP7:2 are interpreted as the Requested Interrupt Priority Level (RIPL) in External Interrupt Controller (EIC) interrupt mode.

**Figure 2.25 Cause Register Format**



**Table 2.34 Field Descriptions for Cause Register**

Name	Bit(s)	Description	Read/Write	Reset State
BD	31	Indicates whether the last exception taken occurred in a branch delay slot. 0: Exception taken was not in delay slot 1: Exception taken was in delay slot  The processor updates BD only if the EXL bit in the Status register was zero when the exception occurred.  If the exception occurred in a branch delay slot, the exception program counter (EPC) is set to restart execution at the branch. Software should read this bit to determine if the exception was taken in a delay slot.	R	Undefined
TI	30	Timer Interrupt. Denotes whether a timer interrupt is pending (analogous to the IP bits for other interrupt types) 0: No timer interrupt is pending 1: Timer interrupt is pending  Hardware sets this bit based on the state of the external <i>SI_TimerInt</i> signal. See also the descriptions of the Count and Compare registers.	R	Undefined
CE	29:28	Coprocessor unit number referenced when a Coprocessor Unusable exception is taken. This field is loaded by hardware on every exception, but is <b>UNPRE-DICTABLE</b> for all exceptions except Coprocessor Unusable.  00: Coprocessor 0 01: Coprocessor 1 10: Coprocessor 2 11: Coprocessor 3	R	Undefined



**Table 2.34 Field Descriptions for Cause Register(continued)**

Name	Bit(s)	Description	Read/Write	Reset State
DC	27	Disable Count register. In some power-sensitive applications, the Count register is not used but may still be the source of some noticeable power dissipation. This bit allows the Count register to be stopped in such situations. For example, this can be useful during low-power operation following a <b>wait</b> instruction.  0: Enable counting of Count register 1: Disable counting of Count register	R/W	0
PCI	26	Performance Counter Interrupt. Indicates whether a performance counter interrupt is pending (analogous to the IP bits for other interrupt types).  0: No performance counter interrupt is pending 1: Performance counter interrupt is pending  See also the description of the PerfCnt registers.	R	Undefined
0	25:24	Reserved. Write as zero. Ignored on reads.	R	0
IV	23	Indicates whether an interrupt exception uses the general exception vector or a special interrupt vector:  0: Use the general exception vector (0x180) 1: Use the special interrupt vector (0x200)  When the IV bit in the Cause register is 1 and the BEV bit in the Status register is 0, the special interrupt vector represents the base of the vector interrupt table.	R/W	Undefined
WP	22	Indicates that a watch exception was deferred because either the Status <sub>EXL</sub> bit or the Status <sub>ERL</sub> bit was a logic '1' at the time the watch exception was detected. This bit both indicates that the watch exception was deferred, and causes the exception to be initiated when Status <sub>EXL</sub> and Status <sub>ERL</sub> are both zero. As such, software must clear this bit as part of the watch exception handler to prevent a watch exception loop.  Software should never write a 1 to this bit when its value is a 0, thereby causing a 0-to-1 transition. If such a transition is caused by software, it is <b>UNPRE-DICTABLE</b> whether hardware ignores the write, accepts the write with no side effects, or accepts the write and initiates a watch exception once Status <sub>EXL</sub> and Status <sub>ERL</sub> are both zero. Software should clear this bit, but never set it. It is set by hardware.	R/W	Undefined
FDCI	21	Fast Debug Channel Interrupt: This bit denotes whether an FDC interrupt is pending (analogous to the IP bits for other interrupt types).  0: No FDC interrupt is pending 1: FDC interrupt is pending  This bit is set by hardware based on the state of the external <i>SI_FDCInt</i> signal.	R	Undefined
0	20:16	Reserved. Write as zero. Ignored on reads.	R	0

**Table 2.34 Field Descriptions for Cause Register(continued)**

Name	Bit(s)	Description	Read/Write	Reset State																					
IP7-2 RIPL	15:10	<p>Indicates an interrupt is pending.</p> <p>If External Interrupt Controller (EIC) mode is disabled (<math>\text{Config3}_{VEIC} = 0</math>), timer interrupts are combined in a system-dependent way with any hardware interrupt. Each bit of this field maps to an individual hardware interrupt.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Name</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>15</td> <td>IP7</td> <td>Hardware interrupt 5</td> </tr> <tr> <td>14</td> <td>IP6</td> <td>Hardware interrupt 4</td> </tr> <tr> <td>13</td> <td>IP5</td> <td>Hardware interrupt 3</td> </tr> <tr> <td>12</td> <td>IP4</td> <td>Hardware interrupt 2</td> </tr> <tr> <td>11</td> <td>IP3</td> <td>Hardware interrupt 1</td> </tr> <tr> <td>10</td> <td>IP2</td> <td>Hardware interrupt 0</td> </tr> </tbody> </table> <p>If EIC interrupt mode is enabled (<math>\text{Config3}_{VEIC} = 1</math>), these bits take on a different meaning and are interpreted as the Requested Interrupt Priority Level (RIPL) field.</p> <p>When EIC interrupt mode is enabled, this field (RIPL) contains the encoded (0 - 63) value of the requested interrupt. A value of zero indicates that no interrupt is requested.</p>	Bit	Name	Meaning	15	IP7	Hardware interrupt 5	14	IP6	Hardware interrupt 4	13	IP5	Hardware interrupt 3	12	IP4	Hardware interrupt 2	11	IP3	Hardware interrupt 1	10	IP2	Hardware interrupt 0	R	Undefined
Bit	Name	Meaning																							
15	IP7	Hardware interrupt 5																							
14	IP6	Hardware interrupt 4																							
13	IP5	Hardware interrupt 3																							
12	IP4	Hardware interrupt 2																							
11	IP3	Hardware interrupt 1																							
10	IP2	Hardware interrupt 0																							
IP1-0	9:8	<p>Controls the request for software interrupts:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Name</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>9</td> <td>IP1</td> <td>Request software interrupt 1</td> </tr> <tr> <td>8</td> <td>IP0</td> <td>Request software interrupt 0</td> </tr> </tbody> </table> <p>These bits are exported to an external interrupt controller for prioritization in EIC interrupt mode with other interrupt sources. The state of these bits are driven onto the external <i>SL_SWInt[1:0]</i> bus.</p>	Bit	Name	Meaning	9	IP1	Request software interrupt 1	8	IP0	Request software interrupt 0	R/W	Undefined												
Bit	Name	Meaning																							
9	IP1	Request software interrupt 1																							
8	IP0	Request software interrupt 0																							
0	7	Reserved. Write as zero. Ignored on reads.	R	0																					
ExcCode	6:2	Encodes the cause of the last exception as described in <a href="#">Table 2.35</a> .	R	Undefined																					
0	1:0	Reserved. Write as zero. Ignored on reads.	R	0																					

**Table 2.35 Exception Code Values in ExcCode Field of Cause Register**

Value (decimal)	Value (hex)	Code	Description
0	0x0	Int	Interrupt
1	0x01	Mod	TLB modification exception
2	0x2	TLBL	Load or fetch, but page not present or marked as invalid in the TLB
3	0x3	TLBS	Store, but page not present or marked as invalid in the TLB

**Table 2.35 Exception Code Values in ExcCode Field of Cause Register (continued)**

Value (decimal)	Value (hex)	Code	Description
4	0x4	AdEL	Address error on load/fetch or store respectively. Address is either wrongly aligned, or a privilege violation.
5	0x5	AdES	
6	0x6	IBE	Bus error signaled on instruction fetch
7	0x7	DBE	Bus error signaled on load/store (imprecise)
8	0x8	Sys	System call, i.e. <b>syscall</b> instruction executed.
9	0x9	Bp	Breakpoint, i.e. <b>break</b> instruction executed. If an SDBBP instruction is executed while the processor is running in EJTAG Debug Mode, this value is written to the <b>Debug<sub>DExcCode</sub></b> field to denote an SDBBP in Debug mode.
10	0xA	RI	Reserved instruction. Instruction code not recognized (or not legal)
11	0xB	CpU	Coprocessor Unusable Exception. Instruction code was for a co-processor which is not enabled in <b>Status<sub>CU3-0</sub></b> .
12	0xC	Ov	Overflow exception. Overflow from a trapping variant of integer arithmetic instructions.
13	0xD	Tr	Trap exception. Condition met on one of the conditional trap instructions <b>teq</b> etc.
14	0xE	-	Reserved
15	0xF	FPE	Floating point unit exception — more details in the FPU control/status registers.
16	0x10	IS1	Coprocessor 2 implementation specific exception
17	0x11	CEU	CorExtend Unusable exception
18	0x12	C2E	Precise Coprocessor 2 exception
21-22	0x15 - 0x16	-	Reserved.
23	0x17	WATCH	Instruction or data reference matched a watchpoint. Refer to <b>WatchHi/WatchLo</b> address.
24	0x18	-	Reserved
25	0x19	Thread	Thread exception. <b>VPEControl<sub>EXCPT</sub></b> specifies the type of the thread exception.
26	0x1A	DSPDis	DSP ASE not enabled or not present exception. This exception occurs when trying to run an instruction from the MIPS DSP ASE, but the ASE is either not enabled or not available. If this exception occurs and the DSP ASE is present in the system, check the state of the <b>Status<sub>MX</sub></b> bit to make sure it is set to '1'.
27-28	0x1B - 0x1C	-	Reserved.
29	0x1D	Prot	An attempt to access an address that is not allowed by the MPU programming.
30	0x1D	CacheErr	Cache error. In normal mode, a cache error exception has a dedicated vector and the Cause register is not updated. If a cache error occurs while in Debug Mode, this code is written to the <b>Debug<sub>DExcCode</sub></b> field to indicate that re-entry to Debug Mode was caused by a cache error.
31	0x1F	-	Reserved.

#### 2.2.4.2 Exception Program Counter — EPC (CP0 Register 14, Select 0)

Following an exception other than an error or debug exception, the Exception Program Counter (EPC) contains the address at which processing resumes after the exception has been serviced (the corresponding debug and error exception use DEPC and ErrorEPC respectively).

Unless the *EXL* bit in the Status register is set (indicating, among other things, that interrupts are disabled), the processor writes the EPC register when an exception occurs.

- For synchronous (precise) exceptions, EPC contains either:
  - The virtual address of the instruction that was the direct cause of the exception, or
  - The virtual address of the immediately preceding branch or jump instruction, when the exception causing instruction is in a branch delay slot, and the Branch Delay bit in the Cause register is set.
- For asynchronous (imprecise) exceptions, EPC contains the address of the instruction at which to resume execution.

On a new exception, the processor does not write to the EPC register when the *EXL* bit in the Status register is set. However, the register can still be written via the *MTC0* instruction.

In processors that implement the MIPS16 ASE, a read of the EPC register (via *MFC0*) returns the following value in the destination GPR:

$$\text{GPR}[rt] \leftarrow \text{ExceptionPC}_{31:1} \ || \ \text{ISAMode}_0$$

That is, the upper 31 bits of the exception PC are combined with the lower bit of the *ISAMode* field of *DEPC7* and written to the GPR.


Similarly, a write to the EPC register (via *MTC0*) takes the value from the GPR and distributes that value to the exception PC and the *ISAMode* field, as follows

$$\begin{aligned} \text{ExceptionPC} &\leftarrow \text{GPR}[rt]_{31:1} \ || \ 0 \\ \text{ISAMode} &\leftarrow 2\#0 \ || \ \text{GPR}[rt]_0 \end{aligned}$$

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the exception PC, and the lower bit of the exception PC is cleared. The upper bit of the *ISAMode* field is cleared and the lower bit is loaded from the lower bit of the GPR.

The processor reads the EPC register as the result of execution of the *eret* instruction.

**Figure 2.26 EPC Register Format**



**Table 2.36 EPC Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
EPC	31:0	Exception Program Counter.	R/W	Undefined

### 2.2.4.3 Error Exception Program Counter — ErrorEPC (CP0 Register 30, Select 0)

The ErrorEPC register is a read/write register, similar to the EPC register, except that ErrorEPC is used on error exceptions. All bits of the ErrorEPC register are significant and must be writable. It is also used to store the program counter on Reset, Soft Reset, and nonmaskable interrupt (NMI) exceptions.

This full 32-bit register is filled with the restart address on a cache error exception or any kind of CPU reset — in fact, any exception which sets *Status<sub>ERL</sub>* and leaves the CPU in "error mode".


The ErrorEPC register contains the virtual address at which instruction processing can resume after servicing an error. This address can be:

- The virtual address of the instruction that caused the exception, or
- the virtual address of the immediately preceding branch or jump instruction when the error causing instruction is in a branch delay slot.

On a reset exception, VPE0's ErrorEPC contains the virtual address at which TC0 would have resumed processing after servicing the error. This, in conjunction with TCRestart registers of other TCs, can provide valuable debug information about the state of the various TCs when the error occurred.

Unlike the EPC register, there is no corresponding branch delay slot indication for the ErrorEPC register.

**Figure 2.27 ErrorEPC Register Format**



**Table 2.37 ErrorEPC Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
ErrorEPC	31:0	Error Exception Program Counter.	R/W	Undefined

## 2.2.5 Timer Registers

This section contains the following timer registers.

- [Section 2.2.5.1, "Count \(CP0 Register 9, Select 0\)" on page 105](#)
- [Section 2.2.5.2, "Compare \(CP0 Register 11, Select 0\)" on page 106](#)

### 2.2.5.1 Count (CP0 Register 9, Select 0)

The Count register acts as a timer, incrementing at a constant rate. Incrementing of this register occurs whether or not an instruction is executed, retired, or any forward progress is made through the pipeline. When enabled by clearing the DC bit in the Cause register, the counter increments every other clock (half the clock rate).

The Count may be stopped in either of the following two circumstances.

- Some implementations may stop Count in the low-power mode, for example, through the **wait** instruction, but only if the Cause<sub>DC</sub> flag is set to 1.
- When the device is in debug mode, the Count register can be stopped by setting Debug<sub>CountDM</sub>. By writing the Count<sub>DM</sub> bit, it is possible to control whether the Count register continues incrementing while the processor is in debug mode.


The Count field starts counting from whatever value is loaded into it. However, OS timers are usually implemented by leaving Count free-running and writing Compare as necessary. This counter rolls over when reaching its maximum value.

By writing the Count<sub>DM</sub> bit in the Debug register, it is possible to control whether the Count register continues incrementing while the processor is in debug mode.

**Figure 2.28 Count Register Format**

31

0



**Table 2.38 Count Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bits			
Count	31:0	Interval counter.	R/W	Undefined

**2.2.5.2 Compare (CP0 Register 11, Select 0)**

The Compare register acts in conjunction with the Count register to implement a timer and timer interrupt function. When the value of the Count register equals the value of the Compare register, the *SI\_TimerInt* output pin is asserted. *SI\_TimerInt* remains asserted until the Compare register is written.


The *SI\_TimerInt* output can be fed back into the interAptiv core on one of the interrupt pins to generate an interrupt. Traditionally, this has been done by multiplexing it with hardware interrupt 5 in order to set interrupt bit IP(7) in the Cause register.

For diagnostic purposes, the Compare register is a read/write register. In normal use however, the Compare register is write-only. As a side effect, writing a value to this register clears the timer interrupt.

**Figure 2.29 Compare Register Format**

31

0



**Table 2.39 Compare Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Compare	31:0	Interval count compare value.	R/W	Undefined

**2.2.6 Cache Management Registers**

This section contains the following cache management registers.

- [Section 2.2.6.1, "Level 1 Instruction Cache Tag Low — ITagLo \(CP0 Register 28, Select 0\)" on page 107](#)
- [Section 2.2.6.2, "Level 1 Instruction Cache Data Low — IDataLo \(CP0 Register 28, Select 1\)" on page 109](#)
- [Section 2.2.6.3, "Level 1 Instruction Cache Data High — IDataHi \(CP0 Register 29, Select 1\)" on page 109](#)
- [Section 2.2.6.4, "Level 1 Data Cache Tag Low — DTagLo \(CP0 Register 28, Select 2\)" on page 110](#)
- [Section 2.2.6.5, "DTagHi \(CP0 Register 29, Select 2\): L1 Data Cache and DSPRAM ECC" on page 113](#)
- [Section 2.2.6.6, "Level 1 Data Cache Data Low — DDataLo \(CP0 Register 28, Select 3\)" on page 114](#)
- [Section 2.2.6.7, "Level 2/3 Cache Tag Low — L23TagLo \(CP0 Register 28, Select 4\)" on page 114](#)
- [Section 2.2.6.8, "Level 2/3 Cache Data Low — L23DataLo \(CP0 Register 28, Select 5\)" on page 115](#)

- Section 2.2.6.9, "Level 2/3 Cache Data High — L23DataHi (CP0 Register 29, Select 5)" on page 116
- Section 2.2.6.10, "ErrCtl (CP0 Register 26, Select 0)" on page 116
- Section 2.2.6.11, "Cache Error — CacheErr (CP0 Register 27, Select 0)" on page 119

### 2.2.6.1 Level 1 Instruction Cache Tag Low — ITagLo (CP0 Register 28, Select 0)

The ITagLo register acts as the interface to the instruction cache tag array. The Index Store Tag and Index Load Tag operations of the CACHE instruction use the ITagLo register as the source of tag information. Note that the interAptiv CPU does not implement the ITagHi register

When the WST bit of the ErrCtl register is asserted, this register becomes the interface to the way-selection RAM. In this mode, the fields are redefined to give appropriate access the contents of the WS array instead of the Tag array.

These registers are a staging location for cache tag information being read/written with **cache** load-tag/store-tag operations.

The interpretation of this register changes depending on the settings of ErrCtl<sub>WST</sub> and ErrCtl<sub>SPR</sub>.

- Default cache interface mode (ErrCtl<sub>WST</sub> = 0, ErrCtl<sub>SPR</sub> = 0)
- Diagnostic "way select test mode" (ErrCtl<sub>WST</sub> = 1, ErrCtl<sub>SPR</sub> = 0)
- For scratchpad memory setup (ErrCtl<sub>WST</sub> = 0, ErrCtl<sub>SPR</sub> = 1)

See the diagrams below for a description. Note that the interAptiv core does not implement the ITagHi register.

#### **ITagLo (ErrCtl<sub>WST</sub> = 0, ErrCtl<sub>SPR</sub> = 0)**

In this mode, this register is a staging location for cache tag information being read/written with **cache** load-tag/store-tag operations—routinely used in cache initialization.

**Figure 2.30 ITagLo Register Format (ErrCtl<sub>WST</sub> = 0)**

, ErrCtl<sub>SPR</sub> = 0)



**Table 2.40 Field Descriptions for ITagLo Register**

Name	Bit(s)	Description	Read/Write	Reset State
PTagLo	31:10	This field contains the physical address of the cache line. Bit 31 corresponds to bit 31 of the PA and bit 10 corresponds to bit 10 of the PA.  Bit 10 is only used when 4 KB caches are implemented. For other cache sizes, this bit will not exist in the tag and will be written as a 0 on IndexLoadTag operations.	R/W	Undefined
0	9:8	Must be written as zero; returns zero on read.	R	0
V	7	Set to 1 if this cache entry is valid (set to zero to initialize the cache).	R/W	Undefined
0	6	Must be written as zero; returns zero on read.	R	0


**Table 2.40 Field Descriptions for ITagLo Register(continued)**

Name	Bit(s)	Description	Read/Write	Reset State
L	5	Specifies the lock bit for the cache tag. This bit is set to lock this cache entry, preventing it from being replaced by another line when a cache miss occurs. When this bit is set, and the V bit is set, the corresponding cache line will not be replaced by the cache replacement algorithm.  This can be used for critical data that must not be removed from the cache. However, this can reduce the efficiency of the cache for memory data competing for space at this index.	R/W	Undefined
0	4:1	Must be written as zero; returns zero on read.	R	0
P	0	Parity bit over the cache tag entries (excluding the D bit). This bit is updated with tag array parity on CACHE Index Load Tag operations and used as tag array parity on Index Store Tag operations when the PO bit of the ErrCtl register is set.	R/W	Undefined

**ITagLo-WST (ErrCtlWST = 1, ErrCtlSPR = 0)**

The way-select RAM is an independent slice of the cache memory (distinct from the tag and data arrays). Test software can access the data in these fields either by **cache** load-tag or store-tag operations when ErrCtl<sub>WST</sub> is set.

**Figure 2.31 ITagLo Register Format (ErrCtlWST = 1, ErrCtlSPR = 0)**




**Table 2.41 Field Descriptions for ITagLo-WST Register**

Name	Bit(s)	Description	Read/Write	Reset State
Unused	31:16	Unused field.	R/W	Undefined
LRU	15:10	LRU bits. This field contains the value read from the WS array after a CACHE Index Load WS operation. It is used to store into the WS array during CACHE Index Store WS operations.  When reading or writing the tag in way-select test mode (that is, with ErrCtl <sub>WST</sub> set), this field reads or writes the LRU ("least recently used") state bits, held in the way-select RAM.	R/W	Undefined
0	9:8	Must be written as zero; returns zero on read..	R	0

**ITagLo-WST (ErrCtlWST = 0, ErrCtlSPR = 1)**

In this mode, the ITagLo register becomes the interface to the instruction scratchpad RAM.

**Figure 2.32 ITagLo Register Format (ErrCtlWST = 0, ErrCtlSPR = 1)**





**Table 2.42 Field Descriptions for ITagLo-SPR Register**


Name	Bit(s)	Description	Read/Write	Reset State
ErrCtlWST = 0, ErrCtlSPR = 1 — Tag 0				
BasePA	31:12	When reading pseudo-tag 0 of a scratchpad RAM, this field contains bits [31:12] of the base address of the scratchpad region.	R/W	Undefined
E	7	When reading pseudo-tag 0 of a scratchpad RAM, this bit indicates whether the scratchpad is enabled.	R/W	Undefined
ErrCtlWST = 0, ErrCtlSPR = 1 — Tag 1				
Size	19:12	When reading pseudo-tag 1 of a scratchpad RAM, this field indicates the size of the scratchpad array. This field is the number of 4KB sections it contains. Combined with bits 11:0, the register will contain the number of bytes in the scratchpad region.	R/W	Undefined

### 2.2.6.2 Level 1 Instruction Cache Data Low — IDataLo (CP0 Register 28, Select 1)

The IDataLo register is a register that acts as the interface to the instruction cache data array and is intended for diagnostic operations only. The Index Load Tag operation of the CACHE instruction reads the corresponding data values into the IDataLo register. If the WST bit in the ErrCtl register is set, then the contents of IDataLo can be written to the cache data array by doing an Index Store Data CACHE instruction. If the SPR bit in the ErrCtl register is set, then the contents of IDataLo can be written to the scratchpad RAM data array by doing an Index Store Data CACHE instruction.

Two registers (IDataHi, IDataLo) are needed, because the interAptiv core loads I-cache data at least 64 bits at a time.

**Figure 2.33 IDataLo Register Format**



**Table 2.43 IDataLo Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DATA	31:0	Low-order data read from the cache data array.	R/W	Undefined

### 2.2.6.3 Level 1 Instruction Cache Data High — IDataHi (CP0 Register 29, Select 1)

The IDataHi register is a register that acts as the interface to the cache data array and is intended for diagnostic operations only. The Index Load Tag operation of the CACHE instruction reads the corresponding data values into the IDataHi register. If the WST bit in the ErrCtl register is set, then the contents of IDataHi can be written to the cache data array by doing an Index Store Data CACHE instruction. If the SPR bit in the ErrCtl register is set, then the contents of IDataHi can be written to the scratchpad RAM data array by doing an Index Store Data CACHE instruction.

Because the interface to the I-cache only operates on pairs of instructions, two registers (IDataHi, IDataLo) are needed because the interAptiv core loads I-cache data at least 64-bits at a time. The high instruction is written into the IDataHi register. Note that IDataHi and IDataLo reflect the memory ordering of the instructions. Depending on the

endianness of the system, Instruction0 belongs in either IDataHi (BigEndian) or IDataLo (LittleEndian) and vice versa for Instruction1.

**Figure 2.34 IDataHi Register Format**



**Table 2.44 IDataHi Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DATA	31:0	High-order data read from the cache data array.	R/W	Undefined

**2.2.6.4 Level 1 Data Cache Tag Low — DTagLo (CP0 Register 28, Select 2)**

The DTagLo register acts as the interface to the data cache tag array. The Index Store Tag and Index Load Tag operations of the CACHE instruction use the DTagLo register as the source of tag information.

When the WST bit of the ErrCtl register is asserted, this register becomes the interface to the way-selection RAM. In this mode, the fields are redefined to give appropriate access the contents of the WS array instead of the Tag array.

These registers are a staging location for cache tag information being read/written with **cache** load-tag/store-tag operations.

The D-cache has five logical memory arrays associated with this DTagLo register. The tag RAM stores tags and other state bits with special attention to the needs of the CPU. The duplicate tag RAM also stores tags and state, but is optimized for the needs of interventions. Both of these arrays are set-associative (4-way). The Dirty RAM and duplicate Dirty RAM store the dirty bits (indicating modified data) for intervention uses, and each combine their ways together in a single entry per set. The WS RAM also combines the lock and LRU data in a single entry per set. Accessing these arrays for index cache loads and stores is controlled by using three bits in the *ErrCtl* register to create modes that allow the correct access to these arrays.

**DTagLo (ErrCtl<sub>WST</sub> = 0, , ErrCtl<sub>SPR</sub> = 0)**

The DTagLo register acts as the interface to the data cache tag array. The Index Store Tag and Index Load Tag operations of the CACHE instruction use the DTagLo register as the source of tag information.

When the WST bit of the ErrCtl register is asserted, this register becomes the interface to the way-selection RAM. In this mode, the fields are redefined to give appropriate access the contents of the WS array instead of the Tag array.

**DTagLo Register Format (ErrCtl<sub>WST</sub> = 0, , ErrCtl<sub>SPR</sub> = 0)**




**Table 2.45 Field Descriptions for DTagLo Register**

Name	Bit(s)	Description	Read/Write	Reset State
PTagLo	31:10	This field contains the physical address of the cache line. Bit 31 corresponds to bit 31 of the PA and bit 10 corresponds to bit 10 of the PA. Bit 10 is only used when 4 KB caches are implemented. For other cache sizes, this bit will not exist in the tag and will be written as a 0 on Index-LoadTag operations.	R/W	Undefined
0	9:8	Reserved. Write as zero. Ignored on reads.	R	0
V	7	Valid entry: This bit is set if this cache entry is valid (set zero to initialize the cache). Index Load: load from V field in primary tag RAM Index Store: store to V field in primary and duplicate tag RAM	R/W	Undefined
D	6	This field indicates whether the cache line is dirty. It will only be set if bit 7 (valid) is also set. For L1 I-cache, this field must be written as zero and returns zero on read.	R/W	Undefined
L	5	Specifies the lock bit for the cache tag. When this bit is set, and the valid bit is set, the corresponding cache line will not be replaced by the cache replacement algorithm.	R/W	Undefined
0	4:1	Reserved. Write as zero. Ignored on reads.	R	0
P	0	Parity. Specifies the parity bit for the cache tag. This bit is updated with tag array parity on CACHE Index Load Tag operations and used as tag array parity on Index Store Tag operations when the PO bit of the ErrCtl register is set. This parity does not cover the dirty bit; the dirty bit has a separate parity bit placed in the way selection RAM. Note that this bit is read only when L1 data cache ECC is enabled.	R/W	Undefined

***DTagLo-WST(ErrCtl<sub>WST</sub> = 1, , ErrCtl<sub>SPR</sub> = 0)***

**Figure 2.35 DTagLo Register Format (ErrCtl<sub>WST</sub> = 1, , ErrCtl<sub>SPR</sub> = 0)**



**Table 2.46 Field Descriptions for DTagLo-WST Register**


Name	Bit(s)	Description	Read/Write	Reset State
U	31:24	Undefined.	R	0

**Table 2.46 Field Descriptions for DTagLo-WST Register**

Name	Bit(s)	Description	Read/Write	Reset State
WSDP	23:20	Dirty Parity (Optional). This field contains the value read from the WS array during a CACHE Index Load WS operation. If the PO field of the ErrCtl register is asserted, then this field is used to store the dirty parity bits during a CACHE Index Store WS operation.	R/W	Undefined
WSD	19:16	Dirty bits. This field contains the value read from the WS array after a CACHE Index Load WS operation. It is used to store into the WS array during CACHE Index Store WS operations.	R/W	Undefined
WSLRU	15:10	LRU bits. This field contains the value read from the WS array after a CACHE Index Load WS operation. It is used to store into the WS array during CACHE Index Store WS operations.	R/W	Undefined
0	9:8	Reserved. Write as zero. Ignored on reads.	R	0
U	7:5	Undefined.	R	0
R	4:3	Must be written as zero; returns zero on read.	R	0
U	2:1	Unused bit.	R/W	0

If a *scratchpad* RAM has been implemented, it must be initialized and managed using **cache** load/store operations while ErrCtl<sub>SPR</sub> is set. The tag load/store operations are used to read and write control registers: During these operations, the DTagLo register has the following bit assignments.

**Figure 2.36 DTagLo Register Format (ErrCtl<sub>WST</sub> = 0, ErrCtl<sub>DYT</sub> = 0, ErrCtl<sub>SPR</sub> = 1)**



**Table 2.47 Field Descriptions for DTagLo-SPR Register**

Name	Bit(s)	Description	Read/Write	Reset State
PTAG	31:12	Scratchpad control. Sets base address.	R/W	Undefined
U	11:10	Undefined.	R	0
0	9:8	Reserved. Write as zero. Ignored on reads.	R	0
E	7	Scratchpad control enable.	R/W	Undefined
U	6:5	Undefined.	R	0
0	4:1	Reserved. Write as zero. Ignored on reads.	R	0
U	0	Undefined.	R	0
Tag = 0				
BasePA	31:12	When reading pseudo-tag 0 of a scratchpad RAM, this field will contain bits [31:12] of the base address of the scratchpad region.	R/W	Undefined


**Table 2.47 Field Descriptions for DTagLo-SPR Register**

Name	Bit(s)	Description	Read/Write	Reset State
E	7	When reading pseudo-tag 0 of a scratchpad RAM, this bit will indicate whether the scratchpad is enabled	R/W	Undefined
Tag = 1				
Size	19:12	When reading pseudo-tag 1 of a scratchpad RAM, this field indicates the size of the scratchpad array. This field is the number of 4KB sections it contains. (Combined with the 0's in 11:0, the register will contain the number of bytes in the scratchpad region.)	R/W	Undefined


**2.2.6.5 DTagHi (CP0 Register 29, Select 2): L1 Data Cache and DSPRAM ECC**

This register functions in the same manner as DTagLo and is the interface for CACHE instructions accessing the ECC bits as might be done in an ECC error handler or software cache test routine. The format of this register depends on the type of error as shown in the figures below.

**Figure 2.37 DTagHi Register Format — DSPRAM ECC Error**



**Figure 2.38 DTagHi Register Format — L1 Data RAM and/or L1 Tag RAM ECC Error**



**Table 2.48 Field Descriptions for DTagHi Register**

Name	Bit(s)	Description	Read/Write	Reset State
DSPRAM ECC error format				
0	31:20	Reserved.	R	Undefined
DSPRAM ECC	19:0	DSPRAM ECC value. This field contains the ECC bits read from the DSPRAM array during a CACHE Index Load Tag instruction. This value is written into the ECC bits on a CACHE Index Store Data if the PO bit of the ErrCtl register is set.	R/W	Undefined
L1 Data or Tag RAM ECC Format				
0	31:17	Reserved.	R	Undefined
L1 Data RAM ECC	16:10	L1 Data RAM ECC value. Bits 16:10 are used to store the L1 Data RAM ECC value.	R/W	Undefined
0	9:7	Reserved.	R	Undefined
L1 Tag RAM ECC	6:0	L1 Tag RAM ECC value. Bits 6:0 store the L1 tag RAM ECC value.	R/W	Undefined

### 2.2.6.6 Level 1 Data Cache Data Low — DDataLo (CP0 Register 28, Select 3)

In the interAptiv core, software can read or write cache data using a **cache** index load tag/index store data instruction. Which word of the cache line is transferred depends on the low address fed to the **cache** instruction.

The DDataLo register acts as the interface to the data cache data array and is intended for diagnostic operations only. The Index Load Tag operation of the CACHE instruction reads the corresponding data values into the DDataLo register. If the WST bit in the ErrCtl register is set, then the contents of DDataLo can be written to the cache data array by doing an Index Store Data CACHE instruction. If the SPR bit in the ErrCtl register is set, then the contents of DDataLo can be written to the scratchpad RAM data array by doing an Index Store Data CACHE instruction.

**Figure 2.39 DDataLo Register Format**



**Table 2.49 DDataLo Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DATA	31:0	Low-order data read from the cache data array.	R/W	Undefined


### 2.2.6.7 Level 2/3 Cache Tag Low — L23TagLo (CP0 Register 28, Select 4)

The L23TagLo register acts as the interface to the L2 or L3 cache tag array. The L2 and L3 Index Store Tag and Index Load Tag operations of the CACHE instruction use the L23TagLo register as the source of tag information. Note that the interAptiv CPU does not implement the L23TagHi register.

The core can be configured without L2/L3 cache support. In this case, this register will be a read-only register that reads as 0.

Figure 2.40 and Table 2.50 describe the fields of L23TagLo as interpreted by the L2 during Index Load Tag and Index Store Tag cache-ops. In Figure 2.41, the Tag field is always left justified so system address bit 31 is at L23TagLo[31].

**Figure 2.40 L23TagLo Register (Tag Accesses)**



**Table 2.50 L23TagLo Register (Tag Accesses) Field Descriptions**

Fields		Description	Read/ Write	Reset State
Name	Bits			
Tag	31:14	Tag.	R/W	Undefined
Reserved	13:9	Reserved. This field should be written with 0s and reads should ignore it.	R/W	Undefined
TP	8	Total Parity.	R/W	Undefined
V	7	Valid.	R/W	Undefined
D	6	Dirty.	R/W	Undefined

**Table 2.50 L23TagLo Register (Tag Accesses) Field Descriptions(continued)**

Fields		Description	Read/Write	Reset State
Name	Bits			
L	5	Lock.	R/W	Undefined
Parity	4:0	Parity.	R/W	Undefined

**Figure 2.41 L23TagLo Register (WS Accesses)**



**Table 2.51 L23TagLo Register (WS Accesses) Field Descriptions**

Fields		Description	Read/Write	Reset State
Name	Bits			
DP	31:24	Dirty Parity.	R/W	Undefined
D	23:16	Dirty.	R/W	Undefined
LRU	15:9	LRU algorithm. For Cache-Ops that access the LRU field, the associativity impacts the number of LRU bits present and how they affect line replacement and refill. The interAptiv core supports an 8-way set associative L2 cache.  The 8-way configuration uses all bits of the LRU field (15:9), but since it is a pseudo-LRU algorithm, the value of the LRU field does not directly correspond to the least-to-most order of the 8 ways.	R/W	Undefined
Reserved	8:0	Reserved. This field should be written with 0s and reads should ignore it.	R/W	Undefined


**2.2.6.8 Level 2/3 Cache Data Low — L23DataLo (CP0 Register 28, Select 5)**

The *L23DataLo* register is a register that acts as the interface to the L2 or L3 cache data array and is intended for diagnostic operations only. The Index Load Tag operation of the CACHE instruction reads the corresponding data values into the *L23DataLo* register. If the WST bit in the ErrCtl register is set, then the contents of *L23DataLo* can be written to the cache data array by doing an Index Store Data CACHE instruction.

The core can be configured without L2/L3 cache support. In this case, this register will be a read-only register that reads as 0.

On interAptiv family cores, test software can read or write cache data using a **cache** index load/store data instruction. Which word of the cache line is transferred depends on the low address fed to the **cache** instruction.

**Figure 2.42 L23DataLo Register Format**




**Table 2.52 L23DataLo Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DATA	31:0	Low-order data read from the cache data array.	R/W	Undefined

### 2.2.6.9 Level 2/3 Cache Data High — L23DataHi (CP0 Register 29, Select 5)

On interAptiv family cores, test software can read or write cache data using a **cache** index load/store data instruction. Which word of the cache line is transferred depends on the low address fed to the **cache** instruction.

**Figure 2.43 L23DataHi Register Format**



**Table 2.53 L23DataHi Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DATA	31:0	High-order data read from the cache data array.	R/W	Undefined

### 2.2.6.10 ErrCtl (CP0 Register 26, Select 0)

The ErrCtl register controls parity protection of data and instruction caches and provides for software testing of the way-selection and scratchpad RAMs.

Parity protection can be enabled or disabled using the PE bit. When parity is implemented and the PO bit is deasserted, the CACHE Index Store Tag and Index Store Data operations will internally generate parity to be written into the RAM arrays. However, when the PO bit is asserted, tag array parity is written using the P bit of the TagLo register and data array parity is written using the PI/PD bits of ErrCtl.

ECC protection for the secondary cache is controlled by a combination of PE and the L2P bits.

A CACHE Index Load Tag operation to the instruction cache will update the PCI field with the instruction precode bits from the data array and the PI field with the parity bits from the data array if parity is supported. A CACHE Index Load Tag operation to the data cache will cause the PD bits to be updated with the byte parity for the selected word of the data array if parity is implemented.

The PCO field can be used for testing the precode bits of the instruction cache data array. When the PCO bit is cleared, the CACHE Index Store Data instruction will internally generate the precode bits to be written into the instruction cache data array. However, when the PCO bit is set, the CACHE Index Store Data instruction will write the value in the PCI field to the precode bits in the data array. Setting an illegal value in the precode bits will cause unpredictable behavior. This mechanism should only be used for software testing of the cache arrays. Furthermore, the cache should be flushed after testing.


The WST, SPR, and ITC bits are used to enable CACHE instruction access to different arrays. On previous cores, these bits have been defined as orthogonal - only one of them should be set at any time. On the interAptiv core, these bits are treated as a three-bit field to allow access to additional arrays. The different test modes are listed in [Table 2.54](#).



**Table 2.54 CACHE Test Mode Control**

WST	SPR	ITC	Description
0	0	0	Normal mode.
0	1	0	SPRAM Access - Index Ld/St Tag instructions will access SPRAM tag values.
1	0	0	Way Select Test - Index Ld/St Tag instructions access Way Select RAM.
1	1	0	Duplicate Tag Array - Index Ld/St Tag instructions will read and write only the duplicate cache tag array.
0	0	1	ITC Access - Index Ld/St Tag instructions will access ITC tag values.
Others			Reserved for future use.

**Figure 2.44 Error Control Register Format**



**Table 2.55 Field Descriptions for ErrCtl Register**

Name	Bit(s)	Description	Read/Write	Reset State						
PE	31	Parity or ECC enable. This bit enables or disables the cache parity or ECC protection for both the instruction cache and the data cache. The selection of parity or ECC is done at build time and the result is stored in the L1ECC bit of this register.  <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Parity or ECC disabled</td> </tr> <tr> <td>1</td> <td>Parity or ECC enabled</td> </tr> </tbody> </table> <p>This field is only writable if the cache parity option was implemented when the CPU was built. If cache parity or ECC is not supported, this field is always read as 0. Software can test for cache parity support by attempting to write a 1 to this field, then read back the value.</p>	Encoding	Meaning	0	Parity or ECC disabled	1	Parity or ECC enabled	R/W	0
Encoding	Meaning									
0	Parity or ECC disabled									
1	Parity or ECC enabled									
PO	30	If set, the PD fields of this register overwrites calculated parity for the data array. In addition, the P field of the TagLo register overwrites calculated parity for the tag array. This bit only has significance during CACHE Index Store Tag and CACHE Index Store Data operations.  0 = User calculated parity 1 = Override calculated parity	R/W	0						
WST	29	Write to 1 for test mode for <b>cache IndexLoadTag/</b> <b>cache IndexStoreTag</b> instructions, which then read/write the cache's internal <i>way-selection RAM</i> instead of the cache tags. This bit works in conjunction with the SPR and ITC bits. Refer to <a href="#">Table 2.54</a> for more information.	R/W	0						
SPR	28	Scratchpad RAM.  This bit works in conjunction with the WST and ITC bits. Refer to <a href="#">Table 2.54</a> for more information.	R/W	0						
PCO	27	Precode override. Used for diagnostic/test of the instruction cache.	R/W	0						

**Table 2.55 Field Descriptions for ErrCtl Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State															
ITC	26	InterThread Communication. If set, Index Load Tag and Index Store Tag CACHE instructions operate on the ITC tag.  CACHE instruction behavior is undefined if this bit is set at the same time as WST or SPR. This bit works in conjunction with the WST and SPR bits of this register. See Table 2.54 above.	R/W	0															
LBE	25	Bit indicating that the most recent Data Bus Error was involved a load instruction. A Per-TC BE bit will indicate which TCs were impacted.	R	Undefined															
WABE	24	Bit indicating that the most recent Data Bus Error was due to a write allocate and that store data was lost. There is no indication of which TC(s) the store request came from.  It is possible for both LBE and WABE to be set if the bus error was on a line being used for both loads and stores.	R	Undefined															
L2P	23	L2 ECC enable. This bit can be set only if the L2 cache is ECC-capable. This bit, in conjunction with the PE bit, enables or disables the ECC protection for the L2 cache:  <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>PE</th> <th>L2P</th> <th>L2 check</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>0</td> <td>1</td> <td>1</td> </tr> </tbody> </table>	PE	L2P	L2 check	1	0	1	1	1	0	0	0	0	0	1	1	R/W	0
PE	L2P	L2 check																	
1	0	1																	
1	1	0																	
0	0	0																	
0	1	1																	
L1ECC	22	L1 ECC configured. This bit indicates if the core was configured with ECC on the L1 data cache.	R	Preset															
0	21	Must be written as zeroes; returns zeroes when read.	R	0															
SE	20	Indicates that a second cache or TLB error was detected before the first error was processed. This is an unrecoverable error. This bit is set when a cache error is detected while the FE bit is set. This bit is cleared on reset or when a cache error is detected with FE cleared.	R	0															
FE	19	Indicates that this is the first cache or TLB error and therefore potentially recoverable. Error handling software should clear this bit when the error has been processed. This bit is set by hardware and cleared by software on reset. Refer to the SE bit description for implications of this bit.  Note that software can only write a 0 to this bit. A write value of 1 will not have any effect.	R/W	0															
PCI	18:13	Instruction precode bits read from or written to the instruction cache data RAM.	R/W	Undefined															
PI	12:4	This field contains the instruction cache data RAM parity bits per doubleword being written or read.  <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Bits</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>12</td> <td>Even parity bit for the pre-code bits.</td> </tr> <tr> <td>11:4</td> <td>Per-byte even parity bits for the 64b of data.</td> </tr> </tbody> </table>	Bits	Meaning	12	Even parity bit for the pre-code bits.	11:4	Per-byte even parity bits for the 64b of data.	R/W	Undefined									
Bits	Meaning																		
12	Even parity bit for the pre-code bits.																		
11:4	Per-byte even parity bits for the 64b of data.																		

**Table 2.55 Field Descriptions for ErrCtl Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
PD	3:0	Parity bits being read/written to the data cache when PO is set.  This field is 0 if L1 data cache ECC is implemented as indicated by the L1ECC bit of this register. Refer to <a href="#">Section 2.2.6.5 “DTagHi (CP0 Register 29, Select 2): L1 Data Cache and DSPRAM ECC”</a> for more information.	R/W	0x0

**2.2.6.11 Cache Error — CacheErr (CP0 Register 27, Select 0)**

The CacheErr register provides an interface with the cache error-detection logic. When a Cache Error exception is signaled, the fields of this register are set accordingly. The format of the CacheErr register is different for Primary caches and the Secondary Cache, as well as the type of error. The EREC field [31:30] indicates the format to be used for decoding the contents of the CacheErr register. Each of these register formats is described in the following subsections.

- ['CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error \(ER:EC = 00 or 10\)" on page 119](#)
- ['CacheErr Register Format — Secondary \(L2\) Cache Error \(ER:EC = 01\)" on page 122](#)
- ['CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error \(ER:EC = 11\)" on page 124](#)

The CacheErr register is used by hardware to report the location of a error in a cache or a scratch pad memory. In cores that implement simple parity, L1 cache errors and scratch pad errors are uncorrected. In cores that implement ECC, there are two error classes: uncorrected and corrected. The CacheErr register has been extended to record both uncorrected and corrected L1 cache and scratch pad errors.

Uncorrected ECC errors are recorded in the same manner as parity errors are recorded in implementations that support simple parity. For correctable errors, a new CacheErr register overlay has been defined for when CacheErr register bits 31:30 equal 11<sub>2</sub>.

On a corrected error (ER:EC = 11) the CacheErr register must be cleared after being processed by software so it can record a new error. This register is cleared by writing back the value currently stored in the register. This particular write-current-value-to-clear behavior is employed to avoid any period of uncertainty — and a lost error record — between when software decides to write the CacheErr register and a subsequent error event.

**CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)**

This format is used when the following error types occur:

- Uncorrectable L1 instruction cache error
- Uncorrectable L1 data cache error
- Uncorrectable instruction scratch pad RAM (ISPRAM) error
- Uncorrectable data scratch pad RAM (DSPRAM) error

**Figure 2.45 CacheErr Register — Uncorrectable Primary Cache or Scratch Pad RAM Errors**



**Table 2.56 CacheErr Register Descriptions — Uncorrectable Primary Cache or Scratch Pad RAM Errors**

Fields		Description	Read / Write	Reset State															
Name	Bits																		
ER:EC	31:30	<p>The ER:EC field provides a first-level indication of the nature and location of a cache or scratch pad RAM error. This indication is encoded as follows.</p> <table border="1"> <thead> <tr> <th>ER</th> <th>EC</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>Uncorrected L1 instruction cache or instruction scratch pad error. In this case, bits 29:0 of this register are used to provide additional information about the error.</td> </tr> <tr> <td>0</td> <td>1</td> <td>Secondary Cache Error. Refer to <a href="#">Section “CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)”</a> below for more information.</td> </tr> <tr> <td>1</td> <td>0</td> <td>Uncorrected L1 Data Cache or Data Scratch Pad Error. In this case, bits 29:0 of this register are used to provide additional information about the error.</td> </tr> <tr> <td>1</td> <td>1</td> <td>Corrected L1 Cache or Scratch Pad Error. Refer to <a href="#">Section “CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)”</a> below for more information.</td> </tr> </tbody> </table>	ER	EC	Meaning	0	0	Uncorrected L1 instruction cache or instruction scratch pad error. In this case, bits 29:0 of this register are used to provide additional information about the error.	0	1	Secondary Cache Error. Refer to <a href="#">Section “CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)”</a> below for more information.	1	0	Uncorrected L1 Data Cache or Data Scratch Pad Error. In this case, bits 29:0 of this register are used to provide additional information about the error.	1	1	Corrected L1 Cache or Scratch Pad Error. Refer to <a href="#">Section “CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)”</a> below for more information.	R	Undefined
ER	EC	Meaning																	
0	0	Uncorrected L1 instruction cache or instruction scratch pad error. In this case, bits 29:0 of this register are used to provide additional information about the error.																	
0	1	Secondary Cache Error. Refer to <a href="#">Section “CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)”</a> below for more information.																	
1	0	Uncorrected L1 Data Cache or Data Scratch Pad Error. In this case, bits 29:0 of this register are used to provide additional information about the error.																	
1	1	Corrected L1 Cache or Scratch Pad Error. Refer to <a href="#">Section “CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)”</a> below for more information.																	
ED:ET	29:28	<p>These bits encode the array in which an error was detected.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>No tag or data RAM error detected for either the instruction or data caches</td> </tr> <tr> <td>01</td> <td>Primary tag RAM error in either the instruction or data cache</td> </tr> <tr> <td>10</td> <td>Primary instruction or data cache data RAM error</td> </tr> <tr> <td>11</td> <td>Duplicate tag RAM error in the primary data cache</td> </tr> </tbody> </table>	Encoding	Meaning	00	No tag or data RAM error detected for either the instruction or data caches	01	Primary tag RAM error in either the instruction or data cache	10	Primary instruction or data cache data RAM error	11	Duplicate tag RAM error in the primary data cache	R/W	00					
Encoding	Meaning																		
00	No tag or data RAM error detected for either the instruction or data caches																		
01	Primary tag RAM error in either the instruction or data cache																		
10	Primary instruction or data cache data RAM error																		
11	Duplicate tag RAM error in the primary data cache																		
ES	27	<p>Error source.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Error on internal request</td> </tr> <tr> <td>1</td> <td>Error on external request</td> </tr> </tbody> </table>	Encoding	Meaning	0	Error on internal request	1	Error on external request	R/W	0									
Encoding	Meaning																		
0	Error on internal request																		
1	Error on external request																		
EE	26	Error external:	R	0															

**Table 2.56 CacheErr Register Descriptions — Uncorrectable Primary Cache or Scratch Pad RAM Errors(continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
EB	25	<p>Error Both. Indicates that a cache error occurred in multiple instruction or data cache arrays.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No additional data cache error</td> </tr> <tr> <td>1</td> <td>Additional data cache error</td> </tr> </tbody> </table> <p>In the case of multiple errors, the Tag ram error has the highest priority, followed by the Data ram error, followed by the Way Select ram. Only the highest priority error information is recorded in the CacheErr register.</p>	Encoding	Meaning	0	No additional data cache error	1	Additional data cache error	R/W	0
Encoding	Meaning									
0	No additional data cache error									
1	Additional data cache error									
EF	24	<p>Error Fatal. Indicates that a fatal cache error has occurred.</p> <p>There are a few situations where software will not be able to get all information about a cache error from the CacheErr register. These situations are fatal because software cannot determine which memory locations have been affected by the error. To enable software to detect these cases, the EF bit (bit 24) has been added to the CacheErr register.</p> <p>The following cases are indicated as fatal cache errors by the EF bit:</p> <ol style="list-style-type: none"> <li>1 Dirty parity error in dirty victim (dirty bit cleared)</li> <li>2 Tag parity error in dirty victim</li> <li>3 Data parity error in dirty victim</li> <li>4 WB store miss and EW error at the requested index</li> <li>5 Dual/Triple errors from different transactions, e.g. scheduled and non-scheduled load.</li> <li>6 Multiple data cache errors detected before the first instruction of the cache error handler is issued.</li> </ol> <p>In addition to the above, simultaneous instruction and data cache errors as indicated by CacheErr<sub>EB</sub> will cause information about the data cache error to be unavailable. However, that situation is not indicated by CacheErr<sub>EF</sub>.</p>	R/W	0						
SP	23	<p>Scratchpad. Indicates Scratchpad RAM parity error.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No Scratchpad RAM error detected</td> </tr> <tr> <td>1</td> <td>Scratchpad RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No Scratchpad RAM error detected	1	Scratchpad RAM error detected	R/W	0
Encoding	Meaning									
0	No Scratchpad RAM error detected									
1	Scratchpad RAM error detected									
EW	22	<p>Error Way. Indicates a parity error on the dirty bits that are stored in the way selection RAM array:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No way selection RAM error detected</td> </tr> <tr> <td>1</td> <td>Way selection RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No way selection RAM error detected	1	Way selection RAM error detected	R	Undefined
Encoding	Meaning									
0	No way selection RAM error detected									
1	Way selection RAM error detected									
Way	21:20	<p>Way. Specifies the cache way in which the error was detected. It is not valid if a Tag RAM error is detected (ET=1) or Scratchpad RAM error is detected (SP=1).</p>	R	Undefined						

**Table 2.56 CacheErr Register Descriptions — Uncorrectable Primary Cache or Scratch Pad RAM Errors(continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
Index	19:0	Index. Specifies the cache or Scratchpad RAM index of the double word in which the error was detected. The way of the faulty cache is written by hardware in the <i>Way</i> field. Software must combine the <i>Way</i> and <i>Index</i> read in this register with cache configuration information in the <i>Config1</i> register in order to obtain an index which can be used in an indexed CACHE instruction to access the faulty cache data or tag. Note that <i>Index</i> is aligned as a byte index, so it does not need to be shifted by software before it is used in an indexed CACHE instruction. <i>Index</i> bits [4:3] are undefined upon tag RAM errors, and <i>Index</i> bits above the MSB actually used for cache indexing will also be undefined.  Bits [19:16] are only used used for errors in the Scratchpad RAM.	R	Undefined

**CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)**

This format is used when the following error types occur:

- Secondary (L2) cache error

**Figure 2.46 CacheErr Register — Secondary Cache**



**Table 2.57 CacheErr Register Field Descriptions — Secondary Cache**

Fields		Description	Read / Write	Reset State															
Name	Bits																		
ER:EC	31:30	The ER:EC field provides a first-level indication of the nature and location of a cache or scratch pad RAM error. This indication is encoded as follows.	R	00															
		<table border="1"> <thead> <tr> <th>ER</th> <th>EC</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>Uncorrected L1 instruction cache or instruction scratch pad error. Refer to <a href="#">Section “CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)”</a> for more information.  In this case, bits 29:0 of this register are used to provide additional information about the error.</td> </tr> <tr> <td>0</td> <td>1</td> <td>Secondary Cache Error. In this case, bits 29:0 of this register are used to provide additional information about the secondary cache error.</td> </tr> <tr> <td>1</td> <td>0</td> <td>Uncorrected L1 Data Cache or Data Scratch Pad Error.  Refer to <a href="#">Section “CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)”</a> for more information.</td> </tr> <tr> <td>1</td> <td>1</td> <td>Corrected L1 Cache or Scratch Pad Error.  Refer to <a href="#">Section “CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)”</a> for more information.</td> </tr> </tbody> </table>			ER	EC	Meaning	0	0	Uncorrected L1 instruction cache or instruction scratch pad error. Refer to <a href="#">Section “CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)”</a> for more information.  In this case, bits 29:0 of this register are used to provide additional information about the error.	0	1	Secondary Cache Error. In this case, bits 29:0 of this register are used to provide additional information about the secondary cache error.	1	0	Uncorrected L1 Data Cache or Data Scratch Pad Error.  Refer to <a href="#">Section “CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)”</a> for more information.	1	1	Corrected L1 Cache or Scratch Pad Error.  Refer to <a href="#">Section “CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)”</a> for more information.
		ER			EC	Meaning													
		0			0	Uncorrected L1 instruction cache or instruction scratch pad error. Refer to <a href="#">Section “CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)”</a> for more information.  In this case, bits 29:0 of this register are used to provide additional information about the error.													
		0			1	Secondary Cache Error. In this case, bits 29:0 of this register are used to provide additional information about the secondary cache error.													
1	0	Uncorrected L1 Data Cache or Data Scratch Pad Error.  Refer to <a href="#">Section “CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)”</a> for more information.																	
1	1	Corrected L1 Cache or Scratch Pad Error.  Refer to <a href="#">Section “CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)”</a> for more information.																	
ED	29	Error Data. Indicates a data RAM error.	R	0															
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No data RAM error detected</td> </tr> <tr> <td>1</td> <td>Data RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No data RAM error detected	1	Data RAM error detected											
Encoding	Meaning																		
0	No data RAM error detected																		
1	Data RAM error detected																		
ET	28	Error Tag. Indicates a tag RAM error.	R	0															
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No tag RAM error detected</td> </tr> <tr> <td>1</td> <td>Tag RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No tag RAM error detected	1	Tag RAM error detected											
Encoding	Meaning																		
0	No tag RAM error detected																		
1	Tag RAM error detected																		
Reserved	27:26	Reserved.	R	00															

**Table 2.57 CacheErr Register Field Descriptions — Secondary Cache (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
EM	25	Error Multi. Indicates that a cache error occurred in multiple L2 arrays.	R	0						
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No multi error</td> </tr> <tr> <td>1</td> <td>Multi error</td> </tr> </tbody> </table>			Encoding	Meaning	0	No multi error	1	Multi error
		Encoding			Meaning					
0	No multi error									
1	Multi error									
In the case of multiple errors, the Tag ram error has the highest priority, followed by the Data ram error, followed by the Way Select ram. Only the highest priority error information is recorded in the CacheErr register.										
EF	24	<p>Error Fatal. Indicates that a fatal cache error has occurred.</p> <p>There are a few situations where software will not be able to get all information about a cache error from the CacheErr register. These situations are fatal because software cannot determine which memory locations have been affected by the error. To enable software to detect these cases, the EF bit (bit 24) has been added to the CacheErr register.</p> <p>This bit is set when a second L2 error occurs before taking the exception for the first L2 error.</p>	R	0						
Reserved	23	Reserved.	R	0						
EW	22	Error Way. Indicates a way-selection RAM error.	R	Undefined						
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No way-selection RAM error detected</td> </tr> <tr> <td>1</td> <td>Way-selection RAM error detected</td> </tr> </tbody> </table>			Encoding	Meaning	0	No way-selection RAM error detected	1	Way-selection RAM error detected
		Encoding			Meaning					
0	No way-selection RAM error detected									
1	Way-selection RAM error detected									
Way	21:19	Way. Specifies the cache way in which the error was detected. It is not valid if a Tag RAM error is detected (ET = 1) or Scratchpad RAM error is detected (SP = 1).	R	Undefined						
Index	18:0	Index. Specifies the cache index of the double word in which the error was detected. The way of the faulty cache is written by hardware in the Way field. Software must combine the Way and Index read in this register with cache configuration information in the Config2 register in order to obtain an index which can be used in an indexed CACHE instruction to access the faulty cache data or tag. Note that Index is aligned as a byte index, so it does not need to be shifted by software before it is used in an indexed CACHE instruction. Index bits [4:3] are undefined upon tag RAM errors and Index bits above the MSB actually used for cache indexing will also be undefined.	R	Undefined						

**CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)**

This format is used when the following error types occur:

- Corrected L1 instruction cache error



- Corrected L1 data cache error
- Corrected Instruction scratch pad RAM (ISPRAM) error
- Corrected Data scratch pad RAM (DSPRAM) error

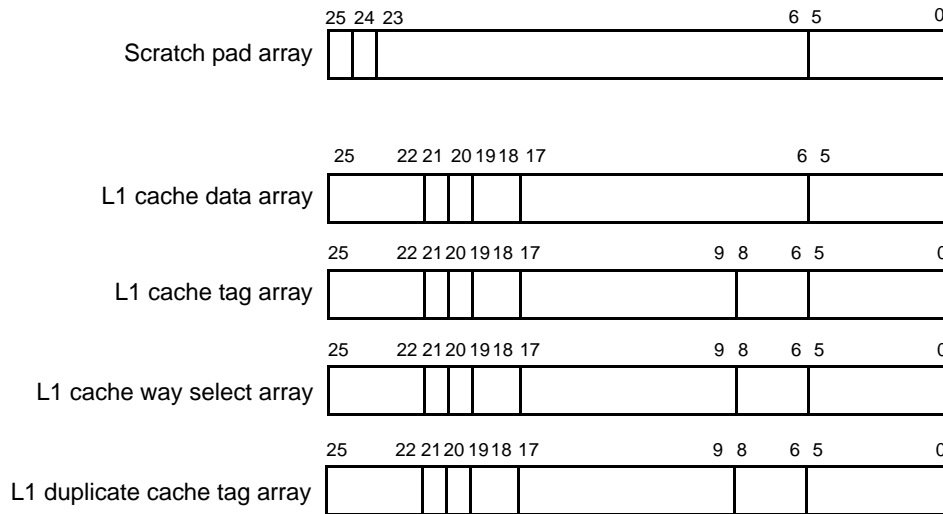
**Figure 2.47 CacheErr Register — Corrected L1 Caches or Scratch Pad RAM Error**

31	30	29	26	25	0	
ER	EC	0				Error Location Specifier

**Table 2.58 CacheErr Register Field Descriptions — Corrected L1 Caches or Scratch Pad RAM Error**

Fields		Description	Read / Write	Reset State															
Name	Bits																		
ER:EC	31:30	<p>Together, ER and EC provide a first-level indication of the nature and location of a cache or scratch pad error. This indication is encoded as follows.</p> <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>ER</th> <th>EC</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>Uncorrected L1 instruction cache or instruction scratch pad RAM error. Refer to the Section entitled 'CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)" on page 119 for more information.</td> </tr> <tr> <td>0</td> <td>1</td> <td>Secondary Cache Error. Refer to the Section entitled 'CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)" on page 122 for more information.</td> </tr> <tr> <td>1</td> <td>0</td> <td>Uncorrected L1 Data cache or data scratch pad RAM error. Refer to the Section entitled 'CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)" on page 119 for more information.</td> </tr> <tr> <td>1</td> <td>1</td> <td>Corrected L1 Cache or Scratch Pad Error. In this case, bits 25:0 of this register are used to determine the format of the ELS field (bits 25:0). Refer to Refer to the Section entitled 'CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)" on page 124 for more information.</td> </tr> </tbody> </table>	ER	EC	Meaning	0	0	Uncorrected L1 instruction cache or instruction scratch pad RAM error. Refer to the Section entitled 'CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)" on page 119 for more information.	0	1	Secondary Cache Error. Refer to the Section entitled 'CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)" on page 122 for more information.	1	0	Uncorrected L1 Data cache or data scratch pad RAM error. Refer to the Section entitled 'CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)" on page 119 for more information.	1	1	Corrected L1 Cache or Scratch Pad Error. In this case, bits 25:0 of this register are used to determine the format of the ELS field (bits 25:0). Refer to Refer to the Section entitled 'CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)" on page 124 for more information.	R	00
ER	EC	Meaning																	
0	0	Uncorrected L1 instruction cache or instruction scratch pad RAM error. Refer to the Section entitled 'CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)" on page 119 for more information.																	
0	1	Secondary Cache Error. Refer to the Section entitled 'CacheErr Register Format — Secondary (L2) Cache Error (ER:EC = 01)" on page 122 for more information.																	
1	0	Uncorrected L1 Data cache or data scratch pad RAM error. Refer to the Section entitled 'CacheErr Register Format — Uncorrectable L1 Instr, L1 Data , ISPRAM, or DSPRAM Error (ER:EC = 00 or 10)" on page 119 for more information.																	
1	1	Corrected L1 Cache or Scratch Pad Error. In this case, bits 25:0 of this register are used to determine the format of the ELS field (bits 25:0). Refer to Refer to the Section entitled 'CacheErr Register Format — Corrected L1 Instr, L1 Data, ISPRAM, or DSPRAM Error (ER:EC = 11)" on page 124 for more information.																	
0	29:26	Reserved. When an L1 cache or SPRAM error is corrected, bits 29:26 of this register are 0 and bits 25:0 are used to provide information about the error as described in <a href="#">Figure 2.48</a> .	R/W	00															
ELS	25:0	Error Location Specifier. See definition in <a href="#">Figure 2.48</a> and <a href="#">Table 2.59</a> below.	R	Undefined															

**Figure 2.48 CacheErr Register — Error Location Specifier (ELS Field) Formats**



**Table 2.59 Error Location Specifier Field Descriptions**

Field Name	Description
I/D	Instruction/Data Store <u>Value</u> 0 Data Store 1 Instruction Store (not used)
Way	Indicates the way number, in a multiple-way cache, where the error was detected.
Index	Indicates the storage array entry number where the error was detected.
Bit Number	Indicates the position of the corrected bit within a code word:  {checkbits, information-bits}.  For error reporting purposes, the right-most check bit is assigned bit number 56. The right-most information bit is assigned bit number 0.  For example, in the case of a cache tag, check bits [6:0] are assigned to bit numbers 62:56, respectively; and information bits [23:0] are assigned to bit numbers 23:0, respectively.  For correctable double-bit errors, which is the case for the duplicate tag and way select arrays, bit number is set to 63.

## 2.2.7 Thread Context and Shadow Control Registers

This section contains the following thread context registers.


- [Section 2.2.7.1, "TCStatus Register \(CP0 Register 2, Select 1\)" on page 127](#)
- [Section 2.2.7.2, "TCBind Register \(CP0 Register 2, Select 2\)" on page 129](#)
- [Section 2.2.7.3, "TCRestart Register \(CP0 Register 2, Select 3\)" on page 130](#)
- [Section 2.2.7.4, "TCHalt Register \(CP0 Register 2, Select 4\)" on page 130](#)
- [Section 2.2.7.5, "TCContext Register \(CP0 Register 2, Select 5\)" on page 131](#)
- [Section 2.2.7.6, "TCSchedule Register \(CP0 Register 2, Select 6\)" on page 131](#)
- [Section 2.2.7.7, "TCScheFBack Register \(CP0 Register 2, Select 7\)" on page 133](#)
- [Section 2.2.7.8, "TCOpt Register \(CP0 Register 3, Select 7\)" on page 134](#)
- [Section 2.2.7.9, "SRSCnf0 \(CP0 Register 6, Select 1\)" on page 134](#)
- [Section 2.2.7.10, "SRSCnf1-4 \(CP0 Register 6, Select 2-5\)" on page 135](#)
- [Section 2.2.7.11, "SRSCtl Register \(CP0 Register 12, Select 2\)" on page 135](#)
- [Section 2.2.7.12, "SRSSMap Register \(CP0 Register 12, Select 3\)" on page 137](#)

### 2.2.7.1 TCStatus Register (CP0 Register 2, Select 1)

The TCStatus register is instantiated per TC as part of the system coprocessor.

[Figure 2.49](#) shows the format of the TCStatus register; [Table 2.60](#) describes the TCStatus register fields.

**Figure 2.49 TCStatus Register Format**



**Table 2.60 TCStatus Register Field Descriptions**

Fields		Description	Read / Write	Reset State	Fork State
Name	Bits				
TCU (TCU3: TCU0)	31:28	Controls access of a TC to coprocessors 3,2,1, and 0 respectively. Status bits <i>CU3:CU0</i> are identical to TCStatus bits <i>TCU3:TCU0</i> of the thread referencing that Status with an MFC0 operation. The modification of either must be visible in both. <ul style="list-style-type: none"> <li>• When no FPU is present, TCU1 is read-only and hardwired to 0</li> <li>• When a single-threaded FPU is present, hardware enforced the rule that only 1 TC can have TCU1 set at a time. Attempts to set TCU1 on a second TC will be ignored.</li> <li>• When a multi-threaded FPU is present, there are no restrictions - TCU1 can be set or cleared by software for any TC.</li> </ul>	R/W	Undefined	Unchanged by FORK

**Table 2.60 TCStatus Register Field Descriptions**

Fields		Description	Read / Write	Reset State	Fork State		
Name	Bits						
TMX	27	DSP ASE Enable. If DSP ASE hardware is present, this field is read/write. If DSP ASE hardware is not present, this field is read-only. Controls access of a TC to extended media processing state, such as MDMX and DSP ASE accumulators. <b>Status</b> bit <i>MX</i> is identical to TCStatus bit <i>TMX</i> of the thread referencing that <b>Status</b> with an MFC0 operation. The modification of either must be visible in both.	Config Option	0	Unchanged by FORK		
TFR	26	This bit is used to control the floating point register mode for 64-bit floating point units.  <b>Status</b> bit <i>FR</i> is identical to TCStatus bit <i>TFR</i> of the thread referencing that <b>Status</b> with an MFC0 operation. The modification of either must be visible in both.	R/W	0	Unchanged by FORK		
RNST	24:23	Run State of TC. Indicates the Running vs. Blocked state of the TC and the reason for blockage. Value is stable only if TC is Halted and examined by another TC using an MFTR operation.	R	0	0		
						<b>Value</b>	<b>Meaning</b>
						0	Running
						1	Blocked on WAIT
2	Blocked on YIELD						
3	Blocked on Gating Storage						
TDS	21	Thread stopped in branch Delay Slot. If a TC is Halted such that the next instruction to issue would be an instruction in a branch delay slot, the TCRestart register will contain the address of the branch instruction, and the <i>TDS</i> bit will be set. Otherwise <i>TDS</i> is cleared on a Halt, or on a software write to the TCRestart register.	R	0	0		
DT	20	Dirty TC. This bit is set by hardware whenever an instruction is retired using the associated TC, and on successful dispatch of the TC via a FORK instruction. The setting of <i>DT</i> by the retirement of instructions is inhibited if the instructions are issued with the <i>EXL</i> or <i>ERL</i> bits of <b>Status</b> set, or with the processor in Debug mode.	R/W	0	1		
TCEE	17	Defined as per the <b>Status</b> register <i>CEE</i> field. This is the per-TC Core Extend Enable value. The <b>Status</b> <i>CEE</i> is identical to the TCStatus <i>TCEE</i> of the thread referencing <b>Status</b> with an MFC0 operation. The modification of either must be visible in both.	R/W	0	Unchanged by FORK		
DA	15	Dynamic Allocation enable. If set, TC may be allocated/deallocated/scheduled by the FORK and YIELD instructions.	R/W	0	FORK allocate only possible if DA = 1		
A	13	Thread Activated. Set automatically when a FORK instruction allocates the TC, and cleared automatically when a YIELD \$0 instruction deallocates it.	R/W	1 for TC 0, 0 for all others.	1		

**Table 2.60 TCStatus Register Field Descriptions**


Fields		Description	Read / Write	Reset State	Fork State
Name	Bits				
TKSU	12:11	Defined as per the Status register <i>KSU</i> field. This is the per-TC Kernel/Supervisor/User state. The Status <i>KSU</i> field is identical to the TCStatus <i>TKSU</i> field of the thread referencing Status. The modification of either must be visible in both.	R/W	Undefined	Copied from forking thread
IXMT	10	Interrupt Exempt. If set, the associated TC will not be used to handle Interrupt exceptions. Debug Interrupt exceptions are not affected.	R/W	0	Unchanged by FORK
TASID	7:0	Defined as per the EntryHi register <i>ASID</i> field. This is the per-TC <i>ASID</i> value. The EntryHi <i>ASID</i> is identical to the TCStatus <i>TASID</i> of the thread referencing EntryHi with an MFC0 operation. The modification of either must be visible in both.  This field is only relevant for the TLB based MMU and will be read-only 0 with an MPU.	R/W	Undefined	Copied from forking thread
0	26:25, 22, 14, 9:8	Must be written as zero; return zero on read.	R	0	0

The *(T)CUx*, *(T)MX*, and *(T)KSU* fields of the TCStatus and Status registers always display the correct state. That is, if the field is written via TCStatus, the new value may be read via Status, and vice-versa. Similarly, the *(T)ASID* field of the TCStatus and EntryHi always display the same current value for the TC.

### 2.2.7.2 TCBind Register (CP0 Register 2, Select 2)

The TCBind register is instantiated per-TC as part of the system co-processor. It defines the VPE affiliation and identification number of this TC.

**Figure 2.50 TCBind Register Format**



**Table 2.61 TCBind Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
0	31:29, 20:18, 16:4	Must be written as zeros; returns zeros on reads.	0	0
CurTC	28:21	Returns the ID number of this TC.	R	Preset
TBE	17	Thread Bus Error: A load instruction from this TC caused an error.	R/W	0
CurVPE	3:0	The ID number of the VPE affiliation of this TC. Externally set on reset based on <i>SL_Vpe0MaxTC</i> . In a two VPE system, all TCs between 0 and <i>SL_Vpe0MaxTC</i> inclusive are bound to VPE0 on reset and remaining ones are bound to VPE1. Writable when <i>MVPCntrolVPC</i> is set	R	External

### 2.2.7.3 TCRestart Register (CP0 Register 2, Select 3)

When a TC is in a Halted state, a read of the TCRestart register returns the instruction address at which the TC will start execution when it is restarted. The TCRestart register can be written while the associated TC is in a Halted state to change the address at which the TC will restart.

Reading the TCRestart register of a non-Halted TC will return the **UNSTABLE** address of some instruction that the TC was executing in the past, but which may no longer be valid. Writing the TCRestart register of a non-Halted TC will result in an **UNDEFINED** TC state.

In the case of branch and jump instructions with architectural delay slots, the restart address will advance beyond the address of the branch or jump instruction only after the instruction in the delay slot has been retired. If halted between the execution of a branch and the associated delay slot instruction, the branch delay slot is indicated by the *TDS* bit of the TCStatus register (see Section 2.2.7.1 “TCStatus Register (CP0 Register 2, Select 1)”).

Software writes to the TCRestart register cause the *TDS* bit of the TCStatus register to be cleared. If a software write of the TCRestart register of a TC intervenes between the execution of an LL instruction and an SC instruction on the target TC, the SC operation must fail.

Figure 2.51 shows the format of the TCRestart register. Table 2.62 describes the TCRestart register fields.

**Figure 2.51 TCRestart Register Format**



**Table 2.62 TCRestart Register Field Descriptions**

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Restart Address	31:0	Address at which execution of the TC is restarted.	R/W	Undefined	Required

#### **Special Handling of TCRestart Register in Processors Implementing MIPS16e™ ASE**

In processors that implement the MIPS16e™ ASE, the TCRestart register requires special handling.

When the processor writes the TCRestart register, it combines the address at which the TC will resume execution with the value of the ISAMode register:

$$\text{TCRestart} \leftarrow \text{resumePC}_{31:1} \parallel \text{ISAMode}_0$$

“resumePC” is the address at which the TC will resume execution, as described above.

When the processor reads the TCRestart register, it distributes the bits to the PC and ISAMode registers:

$$\begin{aligned} \text{PC} &\leftarrow \text{TCRestart}_{31:1} \parallel 0 \\ \text{ISAMode} &\leftarrow \text{TCRestart}_0 \end{aligned}$$


Software reads of the TCRestart register simply return to a GPR the last value written with no interpretation. Software writes to the TCRestart register store a new value which is interpreted by the processor as described above.

### 2.2.7.4 TCHalt Register (CP0 Register 2, Select 4)

The TCHalt register is instantiated per TC as part of the system coprocessor.

Figure 2.52 shows the format of the TCHalt register; Table 2.63 describes the TCHalt register fields.

**Figure 2.52 TCHalt Register Format**



**Table 2.63 TCHalt Register Field Descriptions**

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
H	0	Thread Halted. When set, the associated thread has been halted and cannot be allocated, activated, or scheduled.	R/W	0 for TC 0, 1 for all others	Required
0	31:1	Must be written as zero; return zero on read.	0	0	Reserved

Writing a one to the *Halted* bit of an activated TC causes the associated thread to cease fetching instructions and to set its Restart Address in the TCRestart register (see section 2.2.7.3) to the address of the next instruction to be issued. If the instruction stream associated with the TC is blocked waiting on a response from Gating Storage, the load or store is aborted, and the TC resolves to a state where the TCRestart register and *TDS* field of the TCStatus register (see section 2.2.7.1) reflect a restart at the blocked load or store. Similarly, if the TC is blocked on a WAIT or YIELD instruction, that instruction is cancelled and the state will reflect a restart at the WAIT or YIELD. If the TC was blocked at the time it is Halted, the *RNST* field of TCStatus indicates the blocked state, and the reason for blocking, even if that reason was an operation aborted by the Halt. Writing a zero to the *Halted* bit of an activated TC allows the associated thread of execution to be scheduled, fetching and executing as indicated by TCRestart. A one in the *Halted* bit (TCHalt<sub>H</sub>) of a TC prevents that TC from being allocated and activated by a FORK instruction.


The effect of writing a one to the *Halted* bit of a TC may not be instantaneous. An instruction hazard barrier, e.g. JR.HB, is required to guarantee that the target thread has been fully halted.

### 2.2.7.5 TCContext Register (CP0 Register 2, Select 5)

TCContext is purely a software read/write register, usable by the operating system as a pointer to thread-specific storage, e.g. a thread context save area.

Figure 2.53 shows the format of the TCContext register.

**Figure 2.53 TCContext Register Format**



### 2.2.7.6 TCSchedule Register (CP0 Register 2, Select 6)

The *Scheduler Hint* is a per-TC value whose interpretation is scheduler implementation-dependent. For example, it could encode a description of the requested issue bandwidth for the associated thread, as in the VPESchedule register, or it could encode a priority level.


A TCSchedule register value of zero is the default, and should result in a well-behaved default scheduling of the associated thread.

The VPESchedule register and the TCSchedule register create a hierarchy of issue bandwidth allocation. The set of VPESchedule registers assigns bandwidth to VPEs as a proportion of the total available on a processor or interAptiv,

while the TCSchedule register can only assign bandwidth to threads as a function of that which is available to the VPE containing the thread.

The WRR and WRR2 policy managers described in Chapter 10 implement the TCSchedule register. [Figure 2.54](#) and [Table 2.81](#) shows the format of the TCSchedule register.

**Figure 2.54 TCSchedule Register (CP0 Register2, Select 6)**



**Table 2.64 TCSchedule Register Field Descriptions**

Fields		Description	Read / Write	Reset State	Implemented?											
Name	Bits				WRR	WRR2										
0	31:24	Must be written as 0. Returns zero on reads.	0	0												
T1_STP	23	Throttle1 Stop Priority. When throttle1 is activated and throttle0 is not, the effective stop priority for this TC is set to this value.	R/W	Undef	No	Yes										
T1_GRP	22:21	Throttle1 Group. When throttle1 is activated and throttle0 is not, the effective group for this TC is set to this value.	R/W	Undef	No	Yes										
T1_TH	20:18	Throttle1 Threshold. When an enabled queue input is equal to or less than this value, throttle0 is activated for this TC. NOTE: Setting this value to 7 will disable the threshold check and activate this throttle permanently if there are any enabled queues.	R/W	Undef	No	Yes										
T1_QE	17:14	Throttle1 Queue Enable. When a bit is set in this vector, it sensitizes throttle1 to the available resource as follows: <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>T1_QE</th> <th>PM Input Signal</th> </tr> </thead> <tbody> <tr> <td>17</td> <td>PM_sys_avail[2:0]</td> </tr> <tr> <td>16</td> <td>PM_fsb_avail[2:0]</td> </tr> <tr> <td>15</td> <td>PM_ldq_avail[2:0]</td> </tr> <tr> <td>14</td> <td>PM_wbb_avail[2:0]</td> </tr> </tbody> </table>	T1_QE	PM Input Signal	17	PM_sys_avail[2:0]	16	PM_fsb_avail[2:0]	15	PM_ldq_avail[2:0]	14	PM_wbb_avail[2:0]	R/W	0	No	Yes
T1_QE	PM Input Signal															
17	PM_sys_avail[2:0]															
16	PM_fsb_avail[2:0]															
15	PM_ldq_avail[2:0]															
14	PM_wbb_avail[2:0]															
T0_STP	13	Throttle0 Group. When throttle0 is activated, the effective stop priority for this TC is set to this value.	R/W	Undef	No	Yes										
T0_GRP	12:11	Throttle0 Group. When throttle0 is activated, the effective group for this TC is set to the this value.	R/W	Undef	No	Yes										
T0_TH	10:8	Throttle0 Threshold. When an enabled queue input is equal to or less than this value, throttle0 is activated for this TC. NOTE: Setting this value to 7 will disable the threshold check and activate this throttle permanently if there are any enabled queues.	R/W	Undef	No	Yes										



**Table 2.64 TCSchedule Register Field Descriptions (continued)**


Fields		Description	Read / Write	Reset State	Implemented?											
Name	Bits				WRR	WRR2										
T0_QE	7:4	Throttle0 Queue Enable. When a bit is set in this vector, it sensitizes throttle0 to the available resource inputs as follows:  <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th><i>T0_QE</i></th> <th>PM Input Signal</th> </tr> </thead> <tbody> <tr> <td>7</td> <td><i>PM_sys_avail[2:0]</i></td> </tr> <tr> <td>6</td> <td><i>PM_fsb_avail[2:0]</i></td> </tr> <tr> <td>5</td> <td><i>PM_ldq_avail[2:0]</i></td> </tr> <tr> <td>4</td> <td><i>PM_wbb_avail[2:0]</i></td> </tr> </tbody> </table>	<i>T0_QE</i>	PM Input Signal	7	<i>PM_sys_avail[2:0]</i>	6	<i>PM_fsb_avail[2:0]</i>	5	<i>PM_ldq_avail[2:0]</i>	4	<i>PM_wbb_avail[2:0]</i>	R/W	0	No	Yes
<i>T0_QE</i>	PM Input Signal															
7	<i>PM_sys_avail[2:0]</i>															
6	<i>PM_fsb_avail[2:0]</i>															
5	<i>PM_ldq_avail[2:0]</i>															
4	<i>PM_wbb_avail[2:0]</i>															
STP	3	Stop Priority. Software sets this if this TC should never issue any instructions.	R/W	0	Yes	Yes										
0	2	Must be written as 0. Returns zero on reads.	0	0												
GRP	1:0	Group of the TC. Software sets this value to the group the TC should belong to.	R/W	0	Yes	Yes										

### 2.2.7.7 TCScheFBack Register (CP0 Register 2, Select 7)

The *Scheduler Feedback* is a per-TC feedback value from scheduler hardware to software, whose interpretation is scheduler implementation-dependent. For example, it might encode the number of instructions retired in the instruction stream corresponding to the TC since the last time the value was cleared by software.

The WRR and WRR2 policy managers described in Chapter 10 implement the TCScheFBack register for each TC. [Figure 2.55](#) and [Table 2.65](#) show the format of the TCScheFBack register.

**Figure 2.55 TCScheFBack Register (CP0 Register2, Select 7)**



**Table 2.65 TCScheFBack Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
COUNT	31:0	This is a count of the number of instructions completed by this TC. The value will saturate at 32'hffff_ffff rather than rolling over to 0.	R/W	Undefined

### 2.2.7.8 TCOpt Register (CP0 Register 3, Select 7)

The TCOpt register is instantiated per-TC. If way exclusion is enabled via the MVPCtrl<sub>CPA</sub> bit, the fields in this register will control which ways should be excluded from the replacement scheme for this TC. See also [Section 2.2.7.8, "TCOpt Register \(CP0 Register 3, Select 7\)."](#)

The Prefetch instruction with a hint of “Streamed” will always allocate in way0 regardless of TCOpt. Similarly, PREF/Retained will never allocate in way0 even if TCOpt restricts all other ways.

**Figure 2.56 TCOpt Register Format**



**Table 2.66 TCOpt Register Field Descriptions**


Fields		Description	Read / Write	Reset State															
Name	Bits																		
0	31:4	Must be written as zero; return zero on read.	R	0															
DWX3 : DWX0	3:0	Data cache way exclusion mask. If programmable cache allocation is enabled via the CPA bit in the MVPCtrl register, this field excludes ways of the primary data cache from allocation by the cache controller for any given TC. <table border="1" style="width: 100%; margin-top: 10px;"> <thead> <tr> <th>Bit</th> <th>Name</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>3</td> <td>DWX3</td> <td>If set, D-cache way 3 will not be allocated for the TC</td> </tr> <tr> <td>2</td> <td>DWX2</td> <td>If set, D-cache way 2 will not be allocated for the TC</td> </tr> <tr> <td>1</td> <td>DWX1</td> <td>If set, D-cache way 1 will not be allocated for the TC</td> </tr> <tr> <td>0</td> <td>DWX0</td> <td>If set, D-cache way 0 will not be allocated for the TC</td> </tr> </tbody> </table> <p>NOTE: Software is required to make at least one way available for replacement at all times.</p>	Bit	Name	Meaning	3	DWX3	If set, D-cache way 3 will not be allocated for the TC	2	DWX2	If set, D-cache way 2 will not be allocated for the TC	1	DWX1	If set, D-cache way 1 will not be allocated for the TC	0	DWX0	If set, D-cache way 0 will not be allocated for the TC	R/W	0
Bit	Name	Meaning																	
3	DWX3	If set, D-cache way 3 will not be allocated for the TC																	
2	DWX2	If set, D-cache way 2 will not be allocated for the TC																	
1	DWX1	If set, D-cache way 1 will not be allocated for the TC																	
0	DWX0	If set, D-cache way 0 will not be allocated for the TC																	

### 2.2.7.9 SRSCnf0 (CP0 Register 6, Select 1)

The SRSCnf0 register is instantiated per-VPE. It indicates the binding of TCs or other GPR resources to Shadow Register Sets 1 through 3.

When SRSCnf0 is written, SRSCtl<sub>HSS</sub> is automatically updated by hardware to indicate the highest numbered valid SRS. Software should ensure that the new *HSS* value is not less than the current value of the SRSCtl<sub>CSS</sub> or SRSCtl<sub>PSS</sub>.

**Figure 2.57 SRSCnf0 Register Format**



**Table 2.67 SRSCnf0 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
M	31	Continuation indication. Since there is no SRSCnf1 in the inter-Aptiv core, it will read zero.	R/W	0
0	30	Must be written with zero; returns zero on read	0	0
SRS3-1	29:20, 19:10, 9:0	Indicates the GPR set to be used for corresponding shadow set number (1-3). Shadow set 0 refers to the register set normally associated with the current TC.  Note if a particular SRS is instantiated, all other lower order SRSs must also be instantiated.  If set to 0x3ff indicates this SRS is not supported.  If set to 0x3fe indicates this SRS is not assigned (invalid).	R/W	0x3fe or 0x3ff


**2.2.7.10 SRSCnf1-4 (CP0 Register 6, Select 2-5)**

Not implemented on the interAptiv core.

**2.2.7.11 SRSCtl Register (CP0 Register 12, Select 2)**

The SRSCtl register controls the operation of GPR shadow sets in the processor.

**Figure 2.58 SRSCtl Register Format**



**Table 2.68 SRSCtl Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:30	Must be written as zeros; returns zero on read.	0	0

**Table 2.68 SRSCtl Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
HSS	29:26	<p>Highest Shadow Set. This field contains the highest shadow set number that is implemented by this processor. A value of zero in this field indicates that only the normal GPRs are implemented.</p> <p>Possible values of this field for the interAptiv core are:</p> <p>0x0: One shadow register set present            0x1: Two shadow register sets present            0x2: Three shadow register sets present            0x3: Four shadow register sets present            0x4 - 0xF: Reserved</p> <p>The value in this field also represents the highest value that can be written to the ESS, EICSS, PSS, and CSS fields of this register, or to any of the fields of the SRSSMap register. The operation of the processor is <b>UNDEFINED</b> if a value larger than the one in this field is written to any of these other fields. This field is automatically updated when SRSSConf0 is written.</p>	R	Preset
0	25:22	Must be written as zeros; returns zero on read.	0	0
EICSS	21:18	<p>EIC interrupt mode shadow set. If Config3<sub>VEIC</sub> is 1 (EIC interrupt mode is enabled), this field is loaded from the external interrupt controller for each interrupt request and is used in place of the SRSSMap register to select the current shadow set for the interrupt.</p> <p>If Config3<sub>VEIC</sub> is 0, this field returns zero on read.</p>	R	Undefined
0	17:16	Must be written as zeros; returns zero on read.	0	0
ESS	15:12	<p>Exception Shadow Set. This field specifies the shadow set to use on entry to Kernel Mode caused by any exception other than a vectored interrupt.</p> <p>The operation of the processor is <b>UNDEFINED</b> if software writes a value into this field that is greater than the value in the HSS field.</p>	R/W	0
0	11:10	Must be written as zeros; returns zero on read.	0	0
PSS	9:6	<p>Previous Shadow Set. If GPR shadow registers are implemented, and with the exclusions noted in the next paragraph, this field is copied from the CSS field when an exception or interrupt occurs. An ERET instruction copies this value back into the CSS field if Status<sub>BEV</sub> = 0.</p> <p>This field is not updated on any exception which sets Status<sub>ERL</sub> to 1 (i.e., Reset, Soft Reset, NMI, cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with Status<sub>EXL</sub> = 1, or Status<sub>BEV</sub> = 1. This field is not updated on an exception that occurs while Status<sub>ERL</sub> = 1.</p> <p>The operation of the processor is <b>UNDEFINED</b> if software writes a value into this field that is greater than the value in the HSS field.</p>	R/W	0
0	5:4	Must be written as zeros; returns zero on read.	0	0

**Table 2.68 SRSCtl Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
CSS	3:0	<p>Current Shadow Set. If GPR shadow registers are implemented, this field is the number of the current GPR set. With the exclusions noted in the next paragraph, this field is updated with a new value on any interrupt or exception, and restored from the PSS field on an ERET. Table 2.69 describes the various sources from which the CSS field is updated on an exception or interrupt.</p> <p>This field is not updated on any exception which sets <code>StatusERL</code> to 1 (i.e., Reset, Soft Reset, NMI, cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with <code>StatusEXL = 1</code>, or <code>StatusBEV = 1</code>. Neither is it updated on an ERET with <code>StatusERL = 1</code> or <code>StatusBEV = 1</code>. This field is not updated on an exception that occurs while <code>StatusERL = 1</code>.</p> <p>The value of CSS can be changed directly by software only by writing the PSS field and executing an ERET instruction.</p>	R	0

**Table 2.69 Sources for new SRSCtl<sub>CSS</sub> on an Exception or Interrupt**

Exception Type	Condition	SRSCtl <sub>CSS</sub> Source	Comment
Exception	All	SRSCtl <sub>ESS</sub>	
Non-Vectored Interrupt	<code>Cause<sub>IV</sub> = 0</code>	SRSCtl <sub>ESS</sub>	Treat as exception
Vectored Interrupt	<code>Cause<sub>IV</sub> = 1</code> and <code>Config3<sub>VEIC</sub> = 0</code> and <code>Config3<sub>VInt</sub> = 1</code>	SRSSMap <sub>VECTNUM</sub>	Source is internal map register.
Vectored EIC Interrupt	<code>Cause<sub>IV</sub> = 1</code> and <code>Config3<sub>VEIC</sub> = 1</code>	SRSCtl <sub>EICSS</sub>	Source is external interrupt controller.

### 2.2.7.12 SRSSMap Register (CP0 Register 12, Select 3)


The SRSSMap register contains 8 4-bit fields that provide the mapping from an vector number to the shadow set number to use when servicing such an interrupt. The values from this register are not used for a non-interrupt exception, or a non-vectored interrupt (`CauseIV = 0` or `IntCtlVS = 0`). In such cases, the shadow set number comes from SRSCtl<sub>ESS</sub>.

If SRSCtl<sub>HSS</sub> is zero, the results of a software read or write of this register are **UNPREDICTABLE**.

The operation of the processor is **UNDEFINED** if a value is written to any field in this register that is greater than the value of SRSCtl<sub>HSS</sub>.

The SRSSMap register contains the shadow register set numbers for vector numbers 7:0. The same shadow set number can be established for multiple interrupt vectors, creating a many-to-one mapping from a vector to a single shadow register set number.

**Figure 2.59 SRSSMap Register Format**



**Table 2.70 SRSSMap Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
SSV7	31:28	Shadow register set number for Vector Number 7.	R/W	0
SSV6	27:24	Shadow register set number for Vector Number 6.	R/W	0
SSV5	23:20	Shadow register set number for Vector Number 5.	R/W	0
SSV4	19:16	Shadow register set number for Vector Number 4.	R/W	0
SSV3	15:12	Shadow register set number for Vector Number 3.	R/W	0
SSV2	11:8	Shadow register set number for Vector Number 2.	R/W	0
SSV1	7:4	Shadow register set number for Vector Number 1.	R/W	0
SSV0	3:0	Shadow register set number for Vector Number 0.	R/W	0

## 2.2.8 Virtual Processing Element (VPE) Registers

This section contains the following VPE registers.


- [Section 2.2.8.1, "VPE Control Register \(CP0 Register 1, Select 1\)" on page 138](#)
- [Section 2.2.8.2, "VPE Conf0 Register \(CP0 Register 1, Select 2\)" on page 139](#)
- [Section 2.2.8.3, "VPE Conf1 Register \(CP0 Register 1, Select 3\)" on page 140](#)
- [Section 2.2.8.4, "VPE Schedule Register \(CP0 Register 1, Select 5\)" on page 141](#)
- [Section 2.2.8.5, "VPE ScheFBack Register \(CP0 Register 1, Select 6\)" on page 142](#)
- [Section 2.2.8.6, "VPE Opt Register \(CP0 Register 1, Select 7\)" on page 143](#)
- [Section 2.2.8.7, "MVPCControl Register \(CP0 Register 0, Select 1\)" on page 144](#)
- [Section 2.2.8.8, "MVPCConf0 Register \(CP0 Register 0, Select 2\)" on page 144](#)
- [Section 2.2.8.9, "MVPCConf1 Register \(CP0 Register 0, Select 3\)" on page 145](#)

### 2.2.8.1 VPE Control Register (CP0 Register 1, Select 1)

The VPEControl register is instantiated per VPE as part of the system coprocessor.

[Figure 2.60](#) shows the format of the VPEControl register; [Table 2.71](#) describes the VPEControl register fields.

**Figure 2.60 VPEControl Register Format**



**Table 2.71 VPEControl Register Field Descriptions**

Fields		Description	Read / Write	Reset State		
Name	Bits					
YSI	21	YIELD Scheduler Intercept. If set, and the TCStatus <i>DT</i> bit is also set, valid YIELD instructions that could otherwise cause a rescheduling cause a Thread exception with a YIELD Scheduler Exception sub-code (see below).	R/W	0		
GSI	20	Gating Storage Scheduler Intercept. If set, and the TCStatus <i>DT</i> bit is also set, Gating Storage load and store operations that would otherwise block the issuing TC cause a Thread exception with a GS Scheduler Exception sub-code (see below).	R/W	0		
EXCPT	18:16	Exception sub-code of most recently dispatched Thread exception	R	Undefined	<b>Value</b>	<b>Meaning</b>
					0	Thread Underflow
					1	Thread Overflow
					2	Invalid YIELD Qualifier
					3	Gating Storage Exception
					4	YIELD Scheduler Exception
					5	GS Scheduler Exception
					6-7	Reserved
TE	15	Threads Enabled. Set by EMT instruction, cleared by DMT instruction. If set, multiple TCs may be simultaneously active. If cleared, only one thread may execute on the VPE.	R/W	0		
TargTC	7:0	TC number to be used on MTTR and MFTR instructions.	R/W	Undefined		
0	31:22, 19,14:8	Must be written as zero; return zero on read.	0	0		

So long as the *TE* bit is zero, no thread scheduling will be performed by the VPE. On a processor reset, only the reset thread, TC 0, will execute. If *TE* is cleared by software, only the thread which issued the DMT or MTC0 instruction which cleared the bit will issue further instructions. All other TCs of the VPE are suspended.

The effect of clearing *TE* in software may not be instantaneous. An instruction hazard barrier, e.g. JR.HB, is required to guarantee that all other threads have been quiesced.

**2.2.8.2 VPE Conf0 Register (CP0 Register 1, Select 2)**

The VPEConf0 register is instantiated per VPE. It indicates the activation state and privilege level of the VPE. All fields in the VPEConf0 register are read-only in normal execution, but the *MVP* and *VPA* fields are writable while the *MVP* bit is set for the VPE performing the modification.

Figure 2.61 shows the format of the VPEConf0 register; Table 2.72 describes the VPEConf0 register fields.

**Figure 2.61 VPEConf0 Register Format**

31	30	29	28	21	20	19	18	17	16	15	2	1	0	
M	0	XTC			0	TCS	SCS	DCS	ICS	0			MVP	VPA

**Table 2.72 VPEConf0 Register Field Descriptions**

Fields		Description	Read/Write		Reset State
Name	Bits		MVP=0	MVP=1	
M	31	This bit is reserved to indicate that a VPEConf1 register is present. If the VPEConf1 register is not implemented, this bit should read as a 0. If the VPEConf1 register is implemented, this bit should read as a 1.	R		Preset
XTC	28:21	Exclusive TC. Set by hardware when execution is restricted within a VPE to a single TC, due to <i>EXL/ERL</i> being set in the <i>Status</i> register, or <i>TE</i> being cleared in the <i>VPEControl</i> register, this field contains the TC number of the TC eligible to run. Read by hardware when the <i>VPA</i> bit is written set by software. For cross-VPE initialization, <i>XTC</i> is writable by MTTR if the issuing VPE has <i>MVP</i> set <b>and</b> the target VPE has <i>VPA</i> clear.	R	R/W (if <i>VPA</i> not set for target)	0 for VPE 0, Undefined for all others
TCS	19	Tertiary Cache Shared. Indicates that the tertiary cache described in the <i>Config2</i> register is shared with at least one other VPE.	R		Preset
SCS	18	Secondary Cache Shared. Indicates that the secondary cache described in the <i>Config2</i> register is shared with at least one other VPE.	R		Preset
DCS	17	Data Cache Shared. Indicates that the primary data cache described in the <i>Config1</i> register is shared with at least one other VPE.	R		Preset
ICS	16	Instruction Cache Shared. Indicates that the primary instruction cache described in the <i>Config1</i> register is shared with at least one other VPE.	R		Preset
MVP	1	Master Virtual Processor. If set, the VPE can access the registers of other VPEs of the same processor, using MTTR/MFTR, and can modify the contents of the <i>MVPCControl</i> and <i>VPEConf0</i> registers, thus acquiring the capability to manipulate and configure other VPEs sharing the same processor.	R	R/W	1 for VPE 0, 0 for all others
VPA	0	Virtual Processor Activated. If set, the VPE will schedule threads and execute instructions so long as the <i>EVP</i> bit of the <i>MVPCControl</i> register enables multi-VPE execution.	R	R/W	1 for VPE 0, 0 for all others
0	30:29, 20, 15:2	Reserved. Reads as zero, must be written as zero.	R		0

The *XTC* field is set by hardware on an exception setting *EXL* or *ERL* of the *Status* register, or on an MTC0 or DMT instruction clearing the *TE* bit of *VPEControl*. It may be set by software if and only if both *MVP* of the writing VPE is set and *VPA* of the written VPE is clear, which implies a cross-VPE MTTR operation. It is read by hardware when *VPA* is set, and if the initial state of the VPE is such that only one activated TC may issue, i.e. if *EXL* or *ERL* are set, or *TE* is clear, the TC designated by the *XTC* field will be the TC selected for exclusive execution on the VPE. This allows initialization of one VPE by another, such that the initialized VPE can begin execution in an exception or single-threaded state, and the full context save/restore of one VPE by another, even if the target VPE is in an exception or single-threaded state.

### 2.2.8.3 VPE Conf1 Register (CP0 Register 1, Select 3)

The *VPEConf1* register is instantiated per VPE. It indicates the coprocessor and UDI resources available to the VPE. All fields in the *VPEConf1* register are read-only in normal operation, but is writable while the *MVPCControl* *VPC* bit is set. See section [2.2.8.7](#).





**Table 2.74 VPESchedule Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:6	Must be written as 0. Returns zero on reads.	0	0
GPO	5	Group priority override. When set, the priorities of the groups will be fixed as follows: group3 will be priority3 group2 will be priority2 group1 will be priority1 group0 will be priority0  When cleared, the priorities of the groups will be rotated as described in the Group Rotation Schedule in Chapter 8.  NOTE: GPO is a per-processor field. There is only one GPO register, which is accessible from both GPO fields in a dual-VPE system.	R/W	1
0	4:0	Must be written as 0. Returns zero on reads.	0	0

**2.2.8.5 VPE ScheFBack Register (CP0 Register 1, Select 6)**

The Scheduler Feedback is a per-VPE feedback value from scheduler hardware to software. The interpretation is scheduler implementation-dependent. For example, it might encode the total number of instructions retired in the instruction streams on the associated VPE since the last time the value was cleared by software.

The WRR2 policy manager described in Chapter 10 implements theVPEScheFBack register. [Figure 2.64](#) and [Table 2.75](#) show the format of the VPEScheFBack register.

**Figure 2.64 VPEScheFBack Register (CP0 Register1, Select 6)**



**Table 2.75 VPEScheFBack Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:16	Reserved. Must be written as 0.	R	0
SYSS	15:12	System Buffer Size. This is a reflection of the value on PM_sys_size[3:0].	R	Preset
WBBS	11:8	WBB Size	R	Preset
FSBS	7:4	FSB Size.	R	Preset
LDQS	3:0	LDQ Size.	R	Preset

### 2.2.8.6 VPE Opt Register (CP0 Register 1, Select 7)

The optional VPEOpt register is instantiated per-VPE. If way exclusion is enabled via the MVPControl<sub>CPA</sub> bit, the fields in this register will control which ways should be excluded from the replacement scheme for this VPE.

The Prefetch instruction with a hint of “Streamed” will always allocate in way0 regardless of VPEOpt. Similarly, PREF/Retained will never allocate in way0 even if VPEOpt restricts all other ways.

**Figure 2.65 VPEOpt Register Format**



**Table 2.76 VPEOpt Register Field Descriptions**


Fields		Description	Read / Write	Reset State															
Name	Bits																		
0	31:12, 7:4	Must be written as zero; return zero on read.	R	0															
IWX3: IWX0	11:8	<p>Instruction cache way exclusion mask. If programmable cache allocation is enabled via the CPA bit in the MVPControl register, a VPE can exclude an arbitrary subset of the ways of the primary instruction cache from allocation by the cache controller on behalf of the VPE</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-bottom: 5px;"> <thead> <tr> <th style="text-align: center;">Bit</th> <th style="text-align: center;">Name</th> <th style="text-align: center;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">11</td> <td style="text-align: center;">IWX3</td> <td>If set, I-cache way 3 will not be allocated for the VPE</td> </tr> <tr> <td style="text-align: center;">10</td> <td style="text-align: center;">IWX2</td> <td>If set, I-cache way 2 will not be allocated for the VPE</td> </tr> <tr> <td style="text-align: center;">9</td> <td style="text-align: center;">IWX1</td> <td>If set, I-cache way 1 will not be allocated for the VPE</td> </tr> <tr> <td style="text-align: center;">8</td> <td style="text-align: center;">IWX0</td> <td>If set, I-cache way 0 will not be allocated for the VPE</td> </tr> </tbody> </table> <p>NOTE: Software is required to make at least one way available for replacement at all times.</p>	Bit	Name	Meaning	11	IWX3	If set, I-cache way 3 will not be allocated for the VPE	10	IWX2	If set, I-cache way 2 will not be allocated for the VPE	9	IWX1	If set, I-cache way 1 will not be allocated for the VPE	8	IWX0	If set, I-cache way 0 will not be allocated for the VPE	R/W	0
Bit	Name	Meaning																	
11	IWX3	If set, I-cache way 3 will not be allocated for the VPE																	
10	IWX2	If set, I-cache way 2 will not be allocated for the VPE																	
9	IWX1	If set, I-cache way 1 will not be allocated for the VPE																	
8	IWX0	If set, I-cache way 0 will not be allocated for the VPE																	
DWX3: DWX0	3:0	<p>Data cache way exclusion mask. If programmable cache allocation is enabled via the CPA bit in the MVPControl register, a VPE can exclude an arbitrary subset of the ways of the primary data cache from allocation by the cache controller on behalf of the VPE</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-bottom: 5px;"> <thead> <tr> <th style="text-align: center;">Bit</th> <th style="text-align: center;">Name</th> <th style="text-align: center;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">3</td> <td style="text-align: center;">DWX3</td> <td>If set, D-cache way 3 will not be allocated for the VPE</td> </tr> <tr> <td style="text-align: center;">2</td> <td style="text-align: center;">DWX2</td> <td>If set, D-cache way 2 will not be allocated for the VPE</td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">DWX1</td> <td>If set, D-cache way 1 will not be allocated for the VPE</td> </tr> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">DWX0</td> <td>If set, D-cache way 0 will not be allocated for the VPE</td> </tr> </tbody> </table> <p>NOTE: Software is required to make at least one way available for replacement at all times.</p>	Bit	Name	Meaning	3	DWX3	If set, D-cache way 3 will not be allocated for the VPE	2	DWX2	If set, D-cache way 2 will not be allocated for the VPE	1	DWX1	If set, D-cache way 1 will not be allocated for the VPE	0	DWX0	If set, D-cache way 0 will not be allocated for the VPE	R/W	0
Bit	Name	Meaning																	
3	DWX3	If set, D-cache way 3 will not be allocated for the VPE																	
2	DWX2	If set, D-cache way 2 will not be allocated for the VPE																	
1	DWX1	If set, D-cache way 1 will not be allocated for the VPE																	
0	DWX0	If set, D-cache way 0 will not be allocated for the VPE																	

### 2.2.8.7 MVPControl Register (CP0 Register 0, Select 1)

The MVPControl register is instantiated per-processor, and provides an interface for global control and configuration of a multi-VPE MIPS MT interAptiv core.

Figure 2.66 shows the format of the MVPControl register; Table 2.77 describes the MVPControl register fields.

**Figure 2.66 MVPControl Register Format**




**Table 2.77 MVPControl Register Field Descriptions**

Fields		Description	Read/Write		Reset State
Name	Bits		MVP=0	MVP=1	
0	31:4	Must be written as zero; return zero on read.	0		0
CPA	3	Cache Partitioning Active. If set, the <i>IWX</i> and <i>DWX</i> fields of the <i>VPEOpt</i> register control the allocation of cache lines. If clear, <i>IWX</i> and <i>DWX</i> are ignored.	R	R/W	0
STLB	2	The shared TLB function is not supported in the interAptiv architecture.	R	R	0
VPC	1	Indicates that Processor is in a VPE Configuration State. When <i>VPC</i> is set, some normally “Preset” configuration register fields become writable, to allow for dynamic configuration of processor resources .  Writable by software only if the <i>VPEConf0<sub>MVP</sub></i> bit is set for the VPE issuing the modifying instruction.  Processor behavior is UNDEFINED if <i>VPC</i> and <i>EVP</i> are both in a set state at the same time.	R	R/W	0
EVP	0	Enable Virtual Processors. Modifiable only if the <i>VPEConf0<sub>MVP</sub></i> bit is set for the VPE issuing the modifying instruction. Set by <i>EVPE</i> instruction and cleared by <i>DVPE</i> instruction. If set, all activated (see section 2.2.8.2) VPEs on a processor fetch and execute independently. If cleared, only a single instruction stream on a single VPE can run.	R	R/W	0

### 2.2.8.8 MVPConf0 Register (CP0 Register 0, Select 2)

The MVPConf0-1 registers provide read-only multithreading-specific configuration information.

**Figure 2.67 MVPConf0 Register Format**




**Table 2.78 MVPConf0 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
M	31	This bit reads 1 if the MVPConf1 register is present, otherwise it reads 0.	R	1
TLBS	29	TLB Sharable: This bit is always zero as the TLB is not shared in the interAptiv architecture.	R	0
GS	28	Gating Storage Present. Indicates that the processor is configured to support gating storage operations. Externally set on reset based on the state of the <i>IT_num_entries</i> InterThread input. If <i>IT_num_entries</i> is greater than zero, this bit is set to 1.	R	Preset
PCP	27	Programmable Cache Partitioning: If set, indicates that the allocation behavior of the “ways” of the primary instruction and data caches can be controlled via the VPEOpt register’s <i>IWX</i> and <i>DWX</i> fields.	R	1 if multiple VPEs
PTLBE	25:16	Total processor complement of allocatable TLB entry pairs. TLB configuration is fixed, so PTLBE is zero.	R	0
TCA	15	TCs Allocatable: If set, TCs may be assigned to VPEs by writing the CurVPE field of the TCBind register of each TC while the <i>VPC</i> bit of MVPControl is set.	R	1
PVPE	13:10	Total processor complement of VPE contexts - 1. This field reflects the number of VPEs present after subtracting the value of the static input <i>SI_DisableVPE</i> .	R	Preset: 0 or 1
PTC	7:0	Total processor complement of TCs - 1. This field reflects the number of TCs present after subtracting the value of the static input <i>SI_DisableTCs</i> .	R	Preset: 0 to 8
0	30, 26, 14, 9:8	Must be written as zeros; returns zeros on reads.	0	0

**2.2.8.9 MVPConf1 Register (CP0 Register 0, Select 3)**

The MVPConf0-1 registers provide read-only multithreading-specific configuration information.

**Figure 2.68 MVPConf1 Register Format**



**Table 2.79 MVPConf1 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
C1M	31	If set, floating point unit (co-processor 1) implements the MDMX™ extension to the instruction set.	R	Preset
C1F	30	If set, floating point unit (co-processor 1) implements 64-bit instructions	R	Preset

**Table 2.79 MVPConf1 Register Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
PCX	27:20	Number of register set contexts available for CorExtend. This field is Preset to 0 for any of the following cases: <ul style="list-style-type: none"> <li>Whenever the <i>UDI_present</i> input is deasserted</li> <li>Whenever the <i>UDI_context_present</i> input is deasserted, indicating that the CorExtend module has no state associated with it.</li> </ul> If none of the above are true, then if <i>UDI_mt_context_per_tc</i> is asserted, this field will be set to the number of TCs available in the core, otherwise it will be set to 1.	R	Preset
PCP2	17:10	Number of register set contexts available for co-processor 2. This field represents the value on the <i>CP2_maxtc[3:0]</i> input.	R	Preset
PCP1	7:0	Number of integrated and allocatable FPU contexts. <ul style="list-style-type: none"> <li>If no FPU is present, this will be 0.</li> <li>If a single-threaded FPU is present, this will be 1</li> <li>If the multi-threaded FPU is present, then this will be 0 because the FPU contexts are not allocatable.</li> </ul>	R	Preset
0	29:28, 19:18, 9:8	Must be written as zeros; returns zeros on reads.	0	0

## 2.2.9 Performance Monitoring Registers

This section contains the following performance registers.

- [Section 2.2.9.1, "Performance Counter 0 and 1 Control Registers \(CP0 Register 25, select 0, 2\)" on page 146](#)
- [Section 2.2.9.2, "Performance Counter 0 and 1 Count Registers \(CP0 Register 25, select 1, 3\)" on page 156](#)

### 2.2.9.1 Performance Counter 0 and 1 Control Registers (CP0 Register 25, select 0, 2)

If the processor is configured without performance counter logic, then reading these registers return -1 and writing them has no effect. If the processor is configured with performance counters, there are two performance counters and two associated control registers per TC, which are mapped to CP0 register 25. The select field of the MTC0/MFC0 instructions are used to select the specific register accessed by the instruction, as shown in [Table 2.80](#).

**Table 2.80 Performance Counter Register Selects**

Select[2:0]	Register
0	Register 0 Control
1	Register 0 Count
2	Register 1 Control
3	Register 1 Count

Each counter is a 32-bit read/write register and is incremented by one each time the countable event, specified in its associated control register, occurs. Each counter can independently count one type of event at a time.

Bit 31 of each of the counters are AND'ed with an interrupt enable bit, *IE*, of their respective control register to determine if a performance counter interrupt should be signalled. The performance counter interrupt is implemented per VPE and signalled via the *SI\_PCI* and *SI\_PCI\_1* outputs. The interrupt will be sent to the last VPE to write to the control register. This signal is combined with one of the *SI\_Int* pins to signal an interrupt to the CPU. Counting is not affected by the interrupt indication. This interrupt is deasserted when the above conditions are no longer met - bit 31 of the counter is no longer zero because a value written by software or the count wrapped to zero, or if the IE bit of the control register was cleared.

NOTE: the performance counter registers are connected to a clock that is stopped when the processor is in sleep mode (if the top level clock gater is present). Most events would not be active during that time, but others would be, notably the cycle count. This behavior should be considered when analyzing measurements taken on a system. Further, note that FPGA implementations of the core would generally not have the clock gater present and thus would have different behavior than a typical ASIC implementation.

**Figure 2.69 Performance Counter 0/1 Control Register**

31	30	29		22	21	20	19		16	15	14		12	11		5	4	3	2	1	0
M	0		TCID		MT_EN		VPEID		PCTD		0			Event		IE	U	S	K	EXL	

**Table 2.81 Performance Counter Control Register Field Descriptions**

Fields		Description	Read/Write	Reset State										
Name	Bits													
M	31	If this bit is one, another pair of Performance Control and Counter registers is implemented at a MTC0 or MFC0 select field value of 'n+2' and 'n+3'.	R	Preset										
TCID	29:22	Specifies which TC events should be counted for if per-TC counting is enabled.	R/W	Undefined										
MT_EN	21:20	Specifies which events should be counted: <table border="1" data-bbox="457 1129 1091 1352"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Count events from all TCs and VPEs</td> </tr> <tr> <td>01</td> <td>Count events from all TCs of the VPE specified in <i>VPEID</i></td> </tr> <tr> <td>10</td> <td>Count events from the TC specified in <i>TCID</i></td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	00	Count events from all TCs and VPEs	01	Count events from all TCs of the VPE specified in <i>VPEID</i>	10	Count events from the TC specified in <i>TCID</i>	11	Reserved	R/W	Undefined
Encoding	Meaning													
00	Count events from all TCs and VPEs													
01	Count events from all TCs of the VPE specified in <i>VPEID</i>													
10	Count events from the TC specified in <i>TCID</i>													
11	Reserved													
VPEID	19:16	Specifies which VPE events should be counted if per-VPE counting is enabled.	R/W	Undefined										
PCTD	15	Performance Counter Trace Disable. Setting this bit will prevent the tracing of data from this performance counter when performance count trace mode in PDtrace is enabled.	R/W	0										
Event	11:5	Counter event enabled for this counter. Possible events are listed in <a href="#">Table 2.82</a> .	R/W	Undefined										
IE	4	Counter Interrupt Enable. This bit masks bit 31 of the associated count register from the interrupt exception request output.	R/W	0										
U	3	Count in User Mode. When this bit is set, the specified event is counted in User Mode.	R/W	Undefined										
S	2	Count in Supervisor Mode. When this bit is set, the specified event is counted in Supervisor Mode.	R/W	Undefined										

**Table 2.81 Performance Counter Control Register Field Descriptions(continued)**

Fields		Description	Read/Write	Reset State
Name	Bits			
K	1	Count in Kernel Mode. When this bit is set, count the event in Kernel Mode when <i>EXL</i> and <i>ERL</i> both are 0.	R/W	Undefined
EXL	0	Count when <i>EXL</i> . When this bit is set, count the event when <i>EXL</i> = 1 and <i>ERL</i> = 0.	R/W	Undefined
0	30, 14:12	Must be written as zeroes; returns zeroes when read.	0	0

Table 2.82 describes the events countable with the two performance counters. The mode column indicates whether the event counting is influenced by the mode bits (U,S,K,EXL) The type field indicates whether the event can be per-TC (T), per-VPE (V), or per-Processor (P). TC countable events can also be counted in VPE or Processor modes, and VPE countable events can also be counted in Processor mode. While the counters are implemented per-TC, they are not restricted to counting events for that TC - events from the other VPE or other TCs can be counted. The operation of a counter is **UNPREDICTABLE** for events which are specified as Reserved.

**Table 2.82 Performance Counter Event Types and Codes**

Event Num	Counter 0	Mode	Type	Counter 1	Mode	Type
0	Cycles	No	P	Cycles	No	P
1	Instructions completed	Yes	T	Instructions completed	Yes	T
2	branch instructions	Yes	T	Branch mispredictions	Yes	T
3	JR r31 (return) instructions	Yes	T	JR r31 mispredictions	Yes	T
4	JR (not r31) instructions	Yes	T	JR r31 not predicted	Yes	T
5	ITLB accesses	Yes	T	ITLB misses	Yes	T
6	DTLB accesses	Yes	T	DTLB misses	Yes	T
7	JTLB instruction accesses	Yes	T	JTLB instruction misses	Yes	T
8	JTLB data accesses	Yes	T	JTLB data misses	Yes	T
9	Instruction Cache accesses	Yes	T	Instruction cache misses	Yes	T
10	Data cache accesses	Yes	T	Data cache writebacks	Yes	T
11	Data cache misses	Yes	T	Data cache misses	Yes	T
12	Reserved			Reserved		
13	Store Misses	Yes	T	Load Misses	Yes	T
14	integer instructions completed	Yes	T	FPU instructions completed	Yes	T
15	loads completed	Yes	T	stores completed	Yes	T
16	J/JAL completed	Yes	T	MIPS16 instructions completed	Yes	T
17	no-ops completed	Yes	T	integer multiply/divide completed	Yes	T
18	Stall cycles	No	P	replay traps (other than uTLB)	Yes	T



**Table 2.82 Performance Counter Event Types and Codes (continued)**

Event Num	Counter 0	Mode	Type	Counter 1	Mode	Type
19	SC instructions completed	Yes	T	SC instructions failed	Yes	T
20	Prefetch instructions completed	Yes	T	Prefetch instructions completed with cache hit	Yes	T
21	L2 cache writebacks	No	P	L2 cache accesses	No	P
22	L2 cache misses	No	P	L2 cache single bit errors corrected	No	P
23	Exceptions taken	Yes	T	Single Threaded Mode	Yes	T
24	cache fixup	Yes	T	Refetches	Yes	T
25	IFU stall cycles	No	P	ALU stall cycles	No	P
26	DSP Instructions Completed	Yes	T	ALU-DSP Saturations Done	Yes	T
27	Reserved			MDU-DSP Saturations Done	Yes	T
28	Impl. specific PM event	Yes	T	Impl. specific Cp2 event	Yes	T
29	Impl. specific ISPRAM event	Yes	T	Impl. specific DSPRAM event	Yes	T
30	Impl. specific CorExtend event	Yes	T	Reserved		
31	Impl. specific XYM event	Yes	T	Impl. specific ITC event	Yes	T
32	ITC Loads	Yes	T	ITC Stores	Yes	T
33	Uncached Loads	Yes	T	Uncached Stores	Yes	T
34	FORK Instructions completed	Yes	T	YIELD instruction completed	Yes	T
35	CP2 Arithmetic Instns Completed	Yes	T	CP2 To/From Instns completed	Yes	T
36	Intervention stall main pipe	No	P	Intervention response stalled on miss	No	P
37	I\$ Miss stall cycles	Yes	T	D\$ miss stall cycles	Yes	T
38	SYNC stall cycles	Yes		FSB stall cycles	Yes	
39	D\$ miss cycles	No	P	L2 miss cycles	No	P
40	Uncached stall cycles	Yes	T	ITC stall cycles	Yes	T
41	MDU stall cycles	Yes	T	FPU stall cycles	Yes	T
42	CP2 stall cycles	Yes	T	CorExtend stall cycles	Yes	T
43	ISPRAM stall cycles	Yes	T	DSPRAM stall cycles	Yes	T
44	CACHE Instn stall cycles	No	P	Long stall cycles	Yes	T
45	Load to Use stall cycles	Yes	T	ALU to AGEN stalls cycles	Yes	T
46	Other interlock stall cycles	Yes	T	Branch mispredict stall cycles	No	P
47	Relax bubbles	Yes	V	Number of corrected ECC errors in the L1 Data Cache or DSPRAM	Yes	V

**Table 2.82 Performance Counter Event Types and Codes (continued)**

Event Num	Counter 0	Mode	Type	Counter 1	Mode	Type
48	IFU FB full refetches	Yes	T	FB entry allocated	No	P
49	EJTAG Instruction Triggerpoints	Yes	T	EJTAG Data Triggerpoints	Yes	T
50	FSB < 1/4 full	No	P	FSB 1/4-1/2 full	No	P
51	FSB > 1/2 full	No	P	FSB full pipeline stall cycles	No	P
52	LDQ < 1/4 full	No	P	LDQ 1/4-1/2 full	No	P
53	LDQ > 1/2 full	No	P	LDQ full pipeline stall cycles	No	P
54	WBB < 1/4 full	No	P	WBB 1/4-1/2 full	No	P
55	WBB > 1/2 full	No	P	WBB full pipeline stall cycles	No	P
56	Intervention Hits	No	P	All Interventions	No	P
57	All Invalidates	No	P	Invalidate Hits	No	P
58	Evictions	No	P	Writebacks	No	P
59	ST_Inval	No	P	ST_Exclusive	No	P
60	ST_Store_to_S	Yes	T	ST_Downgrade	No	P
61	Request Latency to Self Intervention	Yes	P	Request Count for SI Latency	Yes	P
62	Request Latency to Read Response		P	Request Count for Resp. Latency		P
63	Reserved					
64	SI_PCEvent[0] - System specific event 0	No	P	SI_PCEvent[1] - System specific event 1	No	P
65	SI_PCEvent[2] - System specific event 2	No	P	SI_PCEvent[3] - System specific event 3	No	P
66	SI_PCEvent[4] - System specific event 4	No	P	SI_PCEvent[5] - System specific event 5	No	P
67	SI_PCEvent[6] - System specific event 6	No	P	SI_PCEvent[7] - System specific event 7	No	P
68-127	Reserved					

**Table 2.83 Event Descriptions**

Event Name	Counter	Event Number	Description
Cycles	0/1	0	Total number of cycles. The performance counters are clocked by the top-level gated clock. If the CPU is built with that clock gater present, none of the counters will increment while the clock is stopped.

**Table 2.83 Event Descriptions (continued)**

Event Name	Counter	Event Number	Description
<b>Instruction Completion:</b> The following events indicate completion of various types of instructions			
Instructions	0/1	1	Total number of instructions completed.
Branch instns	0	2	Counts all branch instructions that completed.
JR R31 (return) instns	0	3	Counts all JR R31 instructions that completed.
JR (not R31)	0	4	Counts all JR \$xx (not \$31) and JALR instructions (indirect jumps).
Integer instns	0	14	Non-floating point, non-Coprocessor 2 instructions.
FPU instns	1	14	Floating point instructions.
Loads	0	15	Includes both integer and coprocessor loads.
Stores	1	15	Includes both integer and coprocessor stores.
J/JAL	0	16	Direct Jump (And Link) instruction.
MIPS16e	1	16	All MIPS16e instruction.
no-ops	0	17	This includes all instructions that normally write to a GPR, but where the destination register was set to r0.
Integer Multiply/Divide	1	17	Counts all Integer Multiply/Divide instructions (MULxx, DIVx, MADDx, MSUBx).
SC	0	19	Counts conditional stores regardless of whether they succeeded.
PREF	0	20	Note that this only counts PREFs that are actually attempted. PREFs to uncached addresses or ones with translation errors are not counted
DSPinstns	0	26	Counts DSP ASE instructions.
ITC Loads	0	32	Counts loads issued to ITC. This includes loads that are rolled back due to the parent TC getting halted or taking an exception.
ITC Stores	1	32	Counts stores issued to ITC. This includes stores that are rolled back due to the parent TC getting halted or taking an exception.
Uncached Loads	0	33	Include both Uncached and Uncached Accelerated CCAs.
Uncached Stores	1	33	
FORK instns	0	34	MT ASE Fork instruction.
YIELD instns	1	34	MT ASE YIELD instruction.
Cp2 Arithmetic instns	0	35	Counts Coprocessor 2 register-to-register instructions.
Cp2 To/From instns	1	35	Includes move to/from, control to/from, and cop2 loads and stores.
<b>Instruction execution events</b>			
Branch mispredicts	1	2	Counts all branch instructions which completed, but were mispredicted.
JR r31 mispredicts	1	3	Counts all JR \$31 instructions which completed, used the RPS for a prediction, but were mispredicted.
JR r31 not-predicted	1	4	RPS will be dynamically associated with only one TC,; returns on other TCs will not be predicted.

**Table 2.83 Event Descriptions (continued)**

Event Name	Counter	Event Number	Description
ITLB accesses	0	5	Counts ITLB accesses that are due to fetches showing up in IF stage of the pipe and do not use fixed mapping or are not in unmapped space. If an address is fetched twice down the pipe (as in the case of a cache miss), that instruction will count 2 ITLB accesses. Also, since each fetch gets us 2 instructions, there is one access marked per double word.
ITLB misses	1	5	Counts all misses in ITLB except ones that are on the back of another miss. We cannot process back to back misses and thus those are ignored for this purpose. Also ignored if there is some form of address error.
DTLB accesses	0	6	Counts DTLB access including those in unmapped address spaces.
DTLB misses	1	6	Counts DTLB misses. Back to back misses that result in only one DTLB entry getting refilled are counted as a single miss.
JTLB instruction accesses	0	7	Instruction JTLB accesses are counted exactly the same as ITLB misses.
JTLB instruction misses	1	7	Counts instruction JTLB accesses that result in no match or a match on an invalid translation.
JTLB data accesses	0	8	Data JTLB accesses.
JTLB data misses	1	8	Counts data JTLB accesses that result in no match or a match on an invalid translation.
I\$ accesses	0	9	Counts every time the instruction cache is accessed. All replays, wasted fetches etc. are counted. For example, following a branch, even the prediction is taken, the fall through access is counted.
I\$ misses	1	9	Counts all instruction cache misses that result in a bus request.
D\$ accesses	0	10	Counts cached loads and stores.
D\$ writebacks	1	10	Counts cache lines written back to memory due to replacement or cacheops.
D\$ misses	0/1	11	Counts loads and stores that miss in the cache
Load Misses	0	13	Counts number of cacheable loads that miss in the cache.
Store Misses	1	13	Counts number of cacheable stores that miss in the cache. Includes stores that hit on a Shared line.
SC instructons failed	1	19	SC instruction that did not update memory Note: While this event and the SC instruction count event can be configured to count in specific operating modes, the timing of the events is much different and the observed operating mode could change between them, causing some inaccuracy in the measured ratio.
PREF completed with cache hit	1	20	Counts PREF instructions that hit in the cache
L2 Cache Writebacks	0	21	Counts cache lines written back to memory due to replacement or cacheops
L2 Cache Accesses	1	21	Number of accesses to L2 Cache
L2 Cache Misses	0	22	Number of accesses that missed in the L2 cache
L2 Cache Single Bit Errors Corrected	1	22	Single bit errors in L2 Cache that were detected and corrected
Exceptions Taken	0	23	Any type of exception taken

**Table 2.83 Event Descriptions (continued)**

Event Name	Counter	Event Number	Description
ALU-DSP Saturations Done	1	26	Number of times a DSP instruction caused an ALU accumulator to saturate
MDU-DSP Saturations Done	1	27	Number of times a DSP instruction caused an MDU accumulator to saturate
EJTAG instruction triggers	0	49	Number of times an EJTAG Instruction Trigger Point condition matched
EJTAG data triggers	1	49	Number of times an EJTAG Data Trigger Point condition matched
Pipeline Events			
Replays	1	18	Counts all replayed instructions. When a long stall condition is detected, instructions are flushed back to the instruction buffer to allow other TCs to advance. The flushed instructions must then be replayed. Count includes instructions that have been replayed multiple times.
Single Threaded mode	1	23	Counts all cycles where one and only one TC is eligible for scheduling instructions. Other TCs can be ineligible based on architectural (cop0) state as well as dynamically detected conditions (long stall, blocked ITC access, WAIT executed, etc)  When counted per-TC or per-VPE, it will count cycles when the specified TC(s) are the one and only TC eligible for scheduling.
Refetches	1	24	Counts the number of replayed instructions that are sent back to IFU to be refetched. If a replay condition is detected, but the instruction is no longer in the instruction buffer, the IFU will need to refetch it.
Cache fixup	0	24	Counts cycles where the LSU is in fixup and cannot accept a new instruction from the ALU. Fixups are replays within the LSU that occur when an instruction needs to re-access the cache or the DTLB. If this event is enabled per TC, the counter will increment if the replayed instruction belongs to the selected TC regardless of which TC caused the replay.
General Stalls			
IFU stall cycles	0	25	Counts the number of cycles where the fetch unit is not providing a valid instruction to the ALU.
ALU stall cycles	1	25	Counts the number of cycles where the ALU pipeline cannot advance.
Stall cycles	0	18	Counts the total number of cycles where no instructions are issued by IFU to ALU (the RF stage does not advance). This includes both of the previous two events. This is different than the sum of them though because cycles when both stalls are active will only be counted once.
Long stall cycles	1	44	This measures stall cycles due to 'long stall' conditions. These are stalls that would be flushed out of the execution pipeline if other TCs were runnable.
<b>Specific stalls</b> - these events will count the number of cycles lost due to this. This will include bubbles introduced by replays within the pipe. If multiple stall sources are active simultaneously, the counters for each of the active events will be incremented.			
Intervention processing stall cycles	0	36	Cycles where the main pipeline is stalled because of intervention processing for cache coherence
SYNC stall cycles	0	38	Cycles where the main pipeline is stalled waiting for a SYNC to complete.

**Table 2.83 Event Descriptions (continued)**

Event Name	Counter	Event Number	Description
FSB index conflict stall cycles	1	38	Cycles where the main pipeline is stalled because of an index conflict in the Fill Store Buffer.
I\$ miss stall cycles	0	37	Cycles when IFU stalls because an I\$ miss caused not to have any runnable instructions. Ignores the stalls due to ITLB misses as well as the 4 cycles following a redirect.
D\$ miss stall cycles	1	37	Counts all cycles where integer pipeline waits on Load return data due to a D-cache miss. The LSU can signal a “long stall” on D-cache misses, in which case the waiting TC might be rescheduled so other TCs can execute instructions till the data returns.
D\$ miss cycles	0	39	D\$ miss is outstanding, but not necessarily stalling the pipeline. The difference between this and D\$ miss stall cycles can show the gain from non-blocking cache misses.
L2 miss cycles	1	39	L2 miss is outstanding, but not necessarily stalling the pipeline.
Uncached stall cycles	0	40	Cycles where the processor is stalled on an uncached fetch, load, or store.
ITC Load/Store stall cycles	1	40	Counts all cycles where a TC is waiting on a ITC load or store to complete and there are no other TCs that can execute.
MDU stall cycles	0	41	Counts all cycles where integer pipeline waits on MDU return data. MDU block can signal a “long stall”, in which case the waiting TC might be rescheduled so other TCs can execute instructions till the data returns.
FPU stall cycles	1	41	Counts all cycles where integer pipeline waits on FPU return data. FPU block can signal a “long stall”, in which case the waiting TC might be rescheduled so other TCs can execute instructions till the data returns.
Cp2 stall cycles	0	42	Counts all cycles where integer pipeline waits on CP2 return data. CP2 block can signal a “long stall”, in which case the waiting TC might be rescheduled so other TCs can execute instructions till the data returns.
CorExtend stall cycles	1	42	Counts all cycles where integer pipeline waits on CorExtend return data. CorExtend block can signal a “long stall”, in which case TC might be rescheduled so other TCs can execute instructions till the data returns.
ISPRAM stall cycles	0	43	Count all pipeline bubbles that are a result of multicycle ISPRAM access. Pipeline bubbles are defined as all cycles that IFU doesn't present an instruction to ALU. The four cycles after a redirect are not counted.
DSPRAM stall cycles	1	43	Counts stall cycles created by an instruction waiting for access to DSPRAM.
CACHE instn stall cycles	0	44	Counts all cycles where pipeline is stalled due to CACHE instructions. Includes cycles where CACHE instructions themselves are stalled in the ALU, and cycles where CACHE instructions cause subsequent instructions to be stalled.
Load to Use stall cycles	0	45	Counts all cycles where integer pipeline waits on Load return data. LSU block can signal a “long stall”, in which case the waiting TC might be rescheduled so other TCs can execute instructions till the data returns.

**Table 2.83 Event Descriptions (continued)**

Event Name	Counter	Event Number	Description
ALU to AGEN stall cycles	1	45	Counts stall cycles due to skewed ALU where the bypass to the address generation takes an extra cycle.
Other interlocks stall cycles	0	46	Counts all cycles where integer pipeline waits on return data from MFC0, RDHWR, and MFTR instructions.
Branch mispredict stalls cycles	1	46	This counts the number of cycles from a mispredicted branch until the next non-delay slot instruction executes. Count is not very meaningful when executing from multiple TCs.
Relax stall cycles	0	47	Number of cycles that a low power op is 'executed' as requested by Policy Manager.
Number of corrected ECC errors	1	47	Number of corrected ECC errors.
FSB full pipeline stall cycles	1	51	Cycles where the pipeline is stalled because the Fill-Store Buffer in LSU is full.
LDQ full pipeline stall cycles	1	53	Cycles where the pipeline is stalled because the Load Data Queue in the LSU is full.
Write Back Buffer full stall cycles	1	55	Cycles where the pipeline is stalled because the WriteBack Buffer in the BIU is full.
<b>Coherence Events</b> - these are events related to cache coherence			
Intervention response pending	1	36	Cycles where an intervention response is delayed because it is waiting for return data from an earlier miss
All Interventions	1	56	These events count external intervention requests and the number of them that hit in the cache. This does not include self-interventions or interventions directed only to the Instruction Cache. It includes all intervention types.
Intervention Hits	0	56	
All Invalidates	0	57	These events count external interventions of types that will leave a cache line in the invalid state. These do not include self-interventions or intervention directed to the Instruction Cache.
Invalidate Hits	1	57	
Evictions	0	58	The core writes back a dirty line to memory as a result of cache replacement or a non-coherent cache operation.
Writebacks	1	58	The core writes back a dirty line to memory as a result of cache replacement or a non-coherent cache operation, self or external memory operation.
ST_Inval	0	59	Counts the number of transitions into the I state from any other state.
ST_Exclusive	1	59	Counts the number of transitions into the E state from I or S states.
ST_Store_to_S	0	60	Counts the number of transitions from S to M due to a store hitting on a shared line
ST_Downgrade	1	60	Counts transitions to S state from M or E.
<b>Latency Events</b> - These events provide a statistical sampling of latencies within the system. One particular FSB entry is monitored. The latency event increments each cycle from the time a request is generated until the self-intervention or response is seen. The count events are incremented once for each request that we are counting the latency for.			
Request Latency to Self Intervention	0	61	Measures latency from miss detection to self intervention. Only counts for coherent requests.
Request Count for SI Latency	1	61	Counts number of coherent requests used for above latency counter
Request Latency to Read Response	0	62	Measures latency from miss detection until critical dword of response is returned, Only counts for cacheable reads.


**Table 2.83 Event Descriptions (continued)**

Event Name	Counter	Event Number	Description												
Request Count for RR Latency	1	62	Counts number of cacheable read requests used for previous latency counter.												
<b>Implementation specific events</b> - Modules that can be replaced by the customer will have an event signal associated with them.															
Policy Manager	0	28	Implementation-specific.												
Cp2	1	28													
ISPRAM	0	29													
DSPRAM	1	29													
XYM	0	31													
ITC	1	31													
CorExtend	0	30													
SI_PCEvent[7:0]	0/1	64-67													
<b>Buffer usage events</b> - These count the number of cycles that buffers within the CPU spend at various levels of fullness. These events cannot be qualified by TC or VPE number															
Fill Store Buffer < 1/4 full	0	50	Buffer Occupancy: The following table shows what values fall into each of the bins for the different buffer sizes that can be chosen. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>State</th> <th>4 Entry Buffer</th> <th>8/9 Entry Buffer</th> </tr> </thead> <tbody> <tr> <td>&lt; 1/4</td> <td>0</td> <td>0-1</td> </tr> <tr> <td>1/4-1/2</td> <td>1-2</td> <td>2-4</td> </tr> <tr> <td>&gt; 1/2</td> <td>3+</td> <td>5+</td> </tr> </tbody> </table>	State	4 Entry Buffer	8/9 Entry Buffer	< 1/4	0	0-1	1/4-1/2	1-2	2-4	> 1/2	3+	5+
State	4 Entry Buffer	8/9 Entry Buffer													
< 1/4	0	0-1													
1/4-1/2	1-2	2-4													
> 1/2	3+	5+													
Fill Store Buffer 1/4 to 1/2 full	1	50													
Fill Store Buffer > 1/2 full	0	51													
Load Data Queue < 1/4 full	0	52													
Load Data Queue 1/4 to 1/2 full	1	52													
Load Data Queue > 1/2 full	0	53													
Write Back Buffer < 1/4 full	0	54													
Write Back Buffer 1/4 to 1/2 full	1	54													
Write Back Buffer > 1/2 full	0	55													
IFU Fill buffer allocated	1	48	Number of cycles where at least one of the IFU fill buffers is allocated (miss pending)												
Refetches due to all IFU Fill Buffers allocated	0	48	Counts the number of times an instruction cache miss was detected, but both fill buffers were already allocated.												

**2.2.9.2 Performance Counter 0 and 1 Count Registers (CP0 Register 25, select 1, 3)**

The performance counter resets to a low-power state, in which none of the counters will start counting events until software has enabled event counting, using an MTC0 instruction to the Performance Counter Control Registers.

**Figure 2.70 Performance Counter Count Register**





**Table 2.84 Performance Counter Count Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
Counter	31:0	Counter	R/W	Undefined

## 2.2.10 Debug Registers

This section contains the following debug registers.

- [Section 2.2.10.1, "Debug \(CP0 Register 23, Select 0\)" on page 157](#)
- [Section 2.2.10.2, "Debug Exception Program Counter — DEPC \(CP0 Register 24, Select 0\)" on page 160](#)
- [Section 2.2.10.3, "Debug Save — DESAVE \(CP0 Register 31, Select 0\)" on page 161](#)
- [Section 2.2.10.4, "Watch Low 0 - 3 — WatchLo0-3 \(CP0 Register 18, Select 0-3\)" on page 162](#)
- [Section 2.2.10.5, "Watch High 0 - 3 — WatchHi0-3 \(CP0 Register 19, Select 0-3\)" on page 162](#)

### 2.2.10.1 Debug (CP0 Register 23, Select 0)

The Debug register provides control and status information while in debug mode. During normal operation (non-debug mode), this register may not be written at all, and only the DM bit and the EJTAGver field returns valid data.

The read-only bits are updated by hardware every time the debug exception is taken, or when a normal exception is taken when already in debug mode (a "nested exception"). Not all fields are valid in both circumstances: Halt and Doze are not defined after a nested exception, and the nested-exception-type field DExcCode is undefined from a debug exception.

Some of the bits and fields are only updated on debug exceptions and/or exceptions in debug mode, as shown below:

- DSS, DBp, DDBL, DDBS, DIB, DINT are updated on both debug exceptions and on exceptions in debug modes
- DExcCode is updated on exceptions in debug mode, and is undefined after a debug exception
- Halt and Doze are updated on a debug exception, and are undefined after an exception in debug mode
- DBD is updated on both debug and on exceptions in debug modes

All bits and fields are undefined when read from normal mode, except those explicitly described to be defined, e.g. EJTAGver and DM.

**Figure 2.71 Debug Register Format**

31	30	29	28	27	26	25	24	23	22	21	20				
DBD	DM	NoDCR	LSNM	Doze	Halt	CountDM	IBusEP	MCheckP	CacheEP	DBusEP	IEXI				
19	18	17	15	14	10	9	8	7	6	5	4	3	2	1	0
DDBSImp	DDBLImp	EJTAGver	DExcCode	NoSSt	SSt	Offline	0	DINT	DIB	DDBS	DDBL	DBp	DSS		

**Table 2.85 Field Descriptions for Debug Register**

Name	Bit(s)	Description	Read/Write	Reset State
DBD	31	Indicates if the last debug exception or exception in debug mode occurred in a branch delay slot.  0: Not in delay slot 1: In delay slot  When set to 1, the Debug Exception Program Counter (DEPC) points to the branch instruction, which is usually the correct place to restart.	R	Undefined
DM	30	Indicates if the processor is operating in debug mode.  0: Processor is operating in non-debug mode 1: Processor is operating in debug mode  In debug mode, this bit is set on any debug exception and is cleared by <b>deret</b> .	R	0
NoDCR	29	Indicates if the dseg memory segment and a memory-mapped DCR register is present.  0: dseg address space is present 1: dseg address space is not present	R	0
LSNM	28	Controls access of load/store between dseg and main memory.  0: Load/stores in dseg address range goes to dseg 1: Load/stores in dseg address range goes to main memory  Setting this bit causes debug-mode accesses to dseg addresses to be sent to system memory. This makes most of the EJTAG unit's control systems unavailable, so will probably only be done around a particular load/store.	R/W	0
Doze	27	Indicates that the processor was in any kind of low power mode when a debug exception occurred.  0: Processor not in low power mode when debug exception occurred 1: Processor in low power mode when debug exception occurred  Before the debug exception, CPU was in one of the reduced power mode.	R	Undefined
Halt	26	Indicates that the internal system bus clock was stopped when the debug exception occurred.  0: Internal system bus clock running 1: Internal system bus clock stopped  Before the debug exception, the CPU was stopped — probably asleep following a <b>wait</b> instruction.	R	Undefined
CountDM	25	Controls or indicates the Count register behavior in debug mode.  0: Count register stopped in debug mode 1: Count register is running in debug mode	R/W	1

**Table 2.85 Field Descriptions for Debug Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
IBusEP	24	<p>These "pending exception" flags remember exception events caused by instructions run in debug mode, but which have not yet occurred because they are imprecise and <code>DebugIEXI</code> is set. Note that you can write a 1 to any of these at any time, so they survive writes to the whole <code>Debug</code> register; but a write of zero to a field is ignored.</p> <p>They remain set until <code>DebugIEXI</code> is cleared explicitly, or implicitly by a <b>deret</b>. If the <b>deret</b> clears the bit, the exception is taken and the pending bit cleared.</p> <p><code>IBusEP</code> remembers a bus error on an instruction fetch. This exception is precise, so it cannot occur and the field is always zero. If <code>IBusEP</code> is set when <code>IEXI</code> is cleared, a Bus Error exception on instruction fetch is taken by the processor, and <code>IBusEP</code> is cleared.</p> <p><code>MCheckP</code> machine check condition. There are no machine check exceptions in the interAptiv core, so this bit is read only and tied to 0.</p> <p><code>CacheEP</code> indicates a precise cache parity error is pending.</p> <p>Data access Bus Error exception Pending: <code>DBusEP</code> remembers a bus error on a data access. Set when an data bus error event occurs or if a 1 is written to the bit by software. Cleared when a Data Bus Error exception is taken by the processor, and by reset. If <code>DBusEP</code> is set when <code>IEXI</code> is cleared, a Data Bus Error exception is taken by the processor, and <code>DBusEP</code> is cleared</p>	R	0
MCheckP	23		R	0
CacheEP	22		R/W	0
DBusEP	21		R/W	0
IEXI	20	Imprecise Error eXception Inhibit. Set when the processor takes a debug exception. Cleared by execution of the <code>DERET</code> instruction; otherwise modifiable by debug mode software. When <code>IEXI</code> is set, the imprecise error exception from a bus error on an instruction fetch or data access, cache error, or machine check is inhibited and deferred until the bit is cleared.	R/W	0
DDBSImpr	19	Imprecise store breakpoint. <code>DEPC</code> probably points to an instruction some time later in the sequence than the store which triggered the breakpoint.	R	0
DDBLImpr	18	Imprecise load breakpoint. <code>DEPC</code> probably points to an instruction some time later in the sequence than the store which triggered the breakpoint. The debugger or user (or both) have to cope as best they can.	R	0
EJTAgver	17:15	These read-only bits encode the revision of the EJTAG specification to which this implementation conforms. The legal values are: 101: Version 5.0 All other values are reserved.	R	5
DExcCode	14:10	Indicates the cause of the latest exception in debug mode. Following initial entry to debug mode, this field is undefined. The subsequent value will be one of those defined in <code>CauseExcCode</code> . See <a href="#">Table 2.35</a> for a list of values. Value is undefined after a debug exception.	R	Undefined
NoSSt	9	Indicates whether the single-step feature controllable by the <code>SSt</code> bit is available in this implementation. This read-only bit is always zero on the interAptiv core because single-step is implemented.	R	0
SSt	8	Controls if debug single step exception is enabled. 0 = No debug single-step exception enabled 1 = Debug single-step exception enabled	R/W	0

**Table 2.85 Field Descriptions for Debug Register (continued)**

Name	Bit(s)	Description	Read/Write	Reset State
Offline	7	Implemented per-TC. When this bit is 1, TC is allowed to execute only in Debug mode.	R/W	0
R	6	Reserved. Must be written as zeros; returns zeros on reads.	R	0
DINT	5	Indicates that a debug interrupt exception (from EJTAG pin) occurred. Cleared on exception in debug mode. 0: No debug interrupt exception 1: Debug interrupt exception	R	Undefined
DIB	4	Instruction breakpoint. This bit is set by hardware when an instruction breakpoint occurs. 0: No debug exception breakpoint 1: Debug exception breakpoint occurred	R	Undefined
DDBS	3	Indicates that a debug data break exception occurred on a store. Cleared on exception in debug mode. 0: No debug data exception on a store 1: Debug instruction exception on a store	R	Undefined
DDBL	2	Indicates that a debug data break exception occurred on a load. Cleared on exception in debug mode. 0: No debug data exception on a load 1: Debug instruction exception on a load	R	Undefined
DBp	1	Indicates that a debug software breakpoint exception occurred. Cleared on exception in debug mode. 0: No debug software breakpoint exception 1: Debug software breakpoint exception	R	Undefined
DSS	0	Indicates that a debug single-step exception occurred. Cleared on exception in debug mode. 0: No debug single-step exception 1: Debug single-step exception	R	Undefined

### 2.2.10.2 Debug Exception Program Counter — DEPC (CP0 Register 24, Select 0)

The Debug Exception Program Counter (DEPC) points to the instruction to restart when a **deret** is executed to exit debug mode. When `DebugDBD` is set, it means that the "real" return address is in a branch delay slot, and DEPC points to the preceding branch.

For synchronous (precise) debug and debug mode exceptions, the DEPC contains either:

- The virtual address of the instruction that was the direct cause of the debug exception, or
- The virtual address of the immediately preceding branch or jump instruction, when the debug exception causing instruction is in a branch delay slot, and the Debug Branch Delay (DBD) bit in the Debug register is set.

For asynchronous debug exceptions (debug interrupt), the DEPC contains the virtual address of the instruction where execution should resume after the debug handler code is executed.

In processors that implement the MIPS16 ASE, a read of the DEPC register (via MFC0) returns the following value in the destination GPR:

$$\text{GPR}[rt] \leftarrow \text{DebugExceptionPC}_{31:1} \parallel \text{ISAMode}_0$$

That is, the upper 31 bits of the debug exception PC are combined with the lower bit of the ISAMode field and written to the GPR.

Similarly, a write to the DEPC register (via MTC0) takes the value from the GPR and distributes that value to the debug exception PC and the ISA Mode field, as follows

$$\begin{aligned} \text{DebugExceptionPC} &\leftarrow \text{GPR}[rt]_{31:1} \parallel 0 \\ \text{ISAMode} &\leftarrow 2\#0 \parallel \text{GPR}[rt]_0 \end{aligned}$$

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the debug exception PC, and the lower bit of the debug exception PC is cleared. The upper bit of the ISAMode field is cleared and the lower bit is loaded from the lower bit of the GPR.

**Figure 2.72 DEPC Register Format**



**Table 2.86 DEPC Register Formats**

Field		Description	Read / Write	Reset
Name	Bit(s)			
DEPC	31:0	The DEPC register is updated with the virtual address of the instruction that caused the debug exception. If the instruction is in the branch delay slot, then the virtual address of the immediately preceding branch or jump instruction is placed in this register.  Execution of the <b>deret</b> instruction causes a jump to the address in the DEPC.	R/W	Undefined

### 2.2.10.3 Debug Save — DESAVE (CP0 Register 31, Select 0)

Software-only register, with no hardware effect. Provided because the debug exception handler can't use the k0-1 GP registers, used by ordinary exception handlers to bootstrap themselves: but a debug handler can save a GPR into DESAVE, and then use that GPR register in code which saves everything else.

**Figure 2.73 DeSave Register Format**




**Table 2.87 DeSave Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DESAVE	31:0	Debug exception save contents.	SO	Undefined

### 2.2.10.4 Watch Low 0 - 3 — WatchLo0-3 (CP0 Register 18, Select 0-3)

Used in conjunction with WatchHi0-3 respectively, each of these registers carries the virtual address and what-to-match fields for a CP0 watchpoint. WatchLo0-1 are used for instruction side accesses and WatchLo2-3 are used for data side accesses. The bit assignments for each of the WatchLo registers is identical. Hence, only one register is shown below.

**Figure 2.74 WatchLo0-3 Register Format**



**Table 2.88 Field Descriptions for WatchLo0-3 Register**

Name	Bit(s)	Description	Read/Write	Reset State
VAddr	31:3	The address to match on, with a resolution of a doubleword.	R/W	Undefined
I	2	Accesses to match: I = Instruction fetches R = Reads (loads) W = Writes (stores)  WatchLo0-1 <sub>R</sub> and WatchLo0-1 <sub>W</sub> are fixed to zero, while WatchLo2-3 <sub>I</sub> will be zero.	R/W	0
R	1		R/W	0
W	0		R/W	0

### 2.2.10.5 Watch High 0 - 3 — WatchHi0-3 (CP0 Register 19, Select 0-3)

These registers provide the interface to a debug facility that causes an exception if an instruction or data access matches the address specified in the registers. Watch exceptions are not taken if the CPU is already in exception mode (that is if Status<sub>EXL</sub> or Status<sub>ERL</sub> is already set).


Watch events which trigger in exception mode are remembered, and result in a "deferred" exception, taken as soon as the CPU leaves exception mode.

WatchHi0-1 are used for instruction side accesses and WatchHi2-3 are used for data side accesses.

This CP0 watchpoint system is independent of the EJTAG debug system (which provides more sophisticated hardware breakpoints).

The WatchLo0-3 registers hold the address to match, while WatchHi0-3 hold a bundle of control fields.

**Figure 2.75 WatchHi0-3 Register Format**



**Table 2.89 Field Descriptions for WatchHi0-3 Register**

Name	Bit(s)	Description	Read/Write	Reset State
M	31	The WatchHi0-3 <sub>M</sub> bit is set whenever there is one more watchpoint register pair to find. Software can use these four bits (starting with WatchHi0) to determine how many watchpoints there are. This field is set for WatchHi0-2 and cleared on WatchHi3.	R	1 (WatchHi0-2) 0 (WatchHi3)
G	30	If the WatchHi0-3 <sub>G</sub> bit is set, any address that matches that specified in the corresponding WatchLo register causes a watch exception. If this bit is zero, the ASID field of the WatchHi register must match the ASID field of the EntryHi register to cause a watch exception.	R/W	Undefined
0	29:24	Reserved. Write as zero. Ignored on reads.	R	0
ASID	23:16	WatchHi0-3 <sub>ASID</sub> matches addresses from a particular address space (the "ASID" is like that in TLB entries) — except that you can set WatchHi0-3 <sub>G</sub> ("global") to match the address in any address space.  The match a particular address, the WatchHi0-3 <sub>G</sub> bit is cleared and the WatchHi0-3 <sub>ASID</sub> value is used to ensure that the match is to the correct address space. If the If the WatchHi0-3 <sub>G</sub> bit is set, the address is always matched, regardless of the WatchHi0-3 <sub>ASID</sub> value.	R/W	Undefined
0	15:12	Reserved. Write as zero. Ignored on reads.	R	0
Mask	11:3	Watch mask. This field marks the corresponding WatchLo0-3 <sub>VAddr</sub> address bits to be ignored when deciding whether this is a match.	R/W	Undefined
I	2	Watch exception type. These bits indicate what type of access (if any) matched after a watch exception.  I = Instruction fetches R = Reads (loads) W = Writes (stores)  Write a 1 to any of these bits in order to <i>clear</i> it (and therefore prevent the exception from immediately happening again). This behavior is unusual among CP0 registers, but it is quite convenient: to clear a watchpoint of all the exception causes you've seen, just read the value of WatchHi0-3 and write it back again. WatchHi0-1 <sub>R</sub> and WatchHi0-1 <sub>W</sub> should always read 0 and WatchHi2-3 <sub>I</sub> should always read 0	W1C	Undefined
R	1		W1C	Undefined
W	0		W1C	Undefined

## 2.2.11 PDTrace Registers

This section contains the following MIPS PDTrace registers.

- [Section 2.2.11.1, "Trace Control Register — TraceControl \(CP0 Register 23, Select 1\)" on page 164](#)
- [Section 2.2.11.2, "Trace Control 2 Register — TraceControl2 \(CP0 Register 23, Select 2\)" on page 166](#)
- [Section 2.2.11.3, "Trace Control 3 Register — TraceControl3 \(CP0 Register 24, Select 2\)" on page 167](#)
- [Section 2.2.11.4, "User Trace Data 1 Register — UserTraceData1 \(CP0 Register 23, Select 3\)" on page 168](#)
- [Section 2.2.11.5, "User Trace Data 2 Register — UserDataTrace2 \(CP0 Register 24, Select 3\)" on page 169](#)
- [Section 2.2.11.6, "Trace Instruction Breakpoint Condition Register — TraceIBPC \(CP0 Register 23, Select 4\)" on page 169](#)

- [Section 2.2.11.7, "Trace Data Breakpoint Condition Register — TraceDBPC \(CP0 Register 23, Select 5\)" on page 170](#)

### 2.2.11.1 Trace Control Register — TraceControl (CP0 Register 23, Select 1)

The TraceControl register configuration is shown below.

**Figure 2.76 TraceControl Register Format**



**Table 2.90 TraceControl Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
TS	31	The trace select bit is used to select between the hardware and the software trace control bits. A value of zero selects the external hardware trace block signals, and a value of one selects the trace control bits in the TraceControl register.	R/W	0
UT	30	This bit has been deprecated and is no longer used since there are now two explicit trace registers, UserTraceData1 and UserTraceData2. This bit is tied to 0 internally.	R	0
0	29:28	Reserved for future use; Must be written as zero; returns zero on read.	R	0
TB	27	Trace All Branch. When set to 1, this tells the processor to trace the PC value for all branches taken, not just the ones whose branch target address is statically unpredictable.	R/W	Undefined
IO	26	Inhibit Overflow. This signal is used to indicate to the interAptiv trace logic that slow but complete tracing is desired. Hence, the interAptiv tracing logic must not allow a FIFO overflow and discard trace data. This is achieved by stalling the pipeline when the FIFO is nearly full, so that no trace records are ever lost.	R/W	Undefined
D	25	Debug mode. When set to one, this enables tracing in debug mode. For a trace to be enabled in Debug mode, the On bit must also be set, and either the G bit must be set, or the current process ASID must match the ASID field in this register.  When set to zero, trace is disabled in debug mode.	R/W	Undefined
E	24	Exception mode. When set to one, tracing is enabled in Exception mode. For a trace to be enabled in Exception mode, the On bit must be set, and either the G bit must be set, or the current process ASID must match the ASID field in this register.  When set to zero, trace is disabled in Exception Mode.	R/W	Undefined
K	23	Kernel mode. When set to one, enables tracing in Kernel mode. For a trace to be enabled in Kernel mode, the On bit must be set, and either the G bit must be set, or the current process ASID must match the ASID field in this register.  When set to zero, trace is disabled in Kernel Mode.	R/W	Undefined



**Table 2.90 TraceControl Register Field Descriptions (continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
S	22	Supervisor mode. When set to one, tracing is enabled in Supervisor Mode. For a trace to be enabled in Supervisor mode, the On bit must be set, and either the G bit must be set, or the current process ASID must match the ASID field in this register.  When set to zero, trace is disabled in Supervisor Mode, regardless of other bits.  If the processor does not implement Supervisor Mode, this bit is ignored on write and returns zero on read.	R/W	Undefined
U	21	User mode. When set to one, tracing is enabled in User mode. For a trace to be enabled in User mode, the On bit must be set, and either the G bit must be set, or the current process ASID must match the ASID field in this register.  When set to zero, trace is disabled in User Mode, regardless of the setting of other bits.	R/W	Undefined
ASID_M	20:13	ASID mask. This is a mask value applied to the ASID comparison (done when the G bit is zero). A “1” in any bit in this field inhibits the corresponding ASID bit from participating in the match. As such, a value of zero in this field compares all bits of ASID.  Note that the ability to mask the ASID value is not available in the hardware signal bit; it is only available via the software control register.	R/W	Undefined
ASID	12:5	Address space identifier. This field stores the ASID field to match when the G bit is zero. When the G bit is one, this field is ignored.	R/W	Undefined
G	4	Global enable. When set, tracing is to be enabled for all processes, provided that other enabling functions (like U, S, etc.) are also true.	R/W	Undefined
TFCR	3	When set, this bit indicates to the PDtrace interface that the optional Fcr bit must be traced in the appropriate trace formats. If PC tracing is disabled, the full PC of the function call (or return) instruction must also be traced.	R/W	Undefined
TLSM	2	Load/Store Miss trace. When set, this indicates to the PDtrace interface that information about data cache misses should be traced. If PC, load/store address, and data tracing are disabled (see the TraceControl2Mode field), the full PC and load/store address are traced for data cache misses.  If load/store data tracing is enabled, the LSM bit must be traced in the appropriate trace format. Note that data cache miss information is only traced if tracing is actually enabled for the current mode.	R/W	Undefined
TIM	1	Trace IM bit. When set, this indicates to the PDtrace interface that the optional IM bit must be traced in the appropriate trace formats. If PC tracing is disabled, the full PC of the instruction that missed in the I-cache must be traced. Note that instruction cache miss information is only traced if tracing is actually enabled in the current mode.	R/W	Undefined
On	0	This is the master trace enable switch in software control. When zero, tracing is always disabled. When set to one, tracing is enabled whenever the other enabling functions are also true.	R/W	0

### 2.2.11.2 Trace Control 2 Register — TraceControl2 (CP0 Register 23, Select 2)

The TraceControl2 register provides additional control and status information. Note that some fields in the TraceControl2 register are read-only, but have a reset state of “Undefined”. This is because these values are loaded from the Trace Control Block (TCB). As such, these fields in the TraceControl2 register will not have valid values until the TCB asserts these values.

This register is only implemented if the MIPS Trace capability is present.

**Figure 2.77 TraceControl2 Register Format**



**Table 2.91 TraceControl2 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:30	Reserved for future use; Must be written as zero; returns zero on read.	R	0
CPUIdV	29	When set, this bit specifies that the VPE defined in CPUId must be traced. Otherwise, instructions from all VPEs are traced when other conditions for tracing are valid. This bit is ignored if TCV is asserted.	R/W	0
CPUId	28:21	This field specifies the number of the VPE to trace when CPUIdV is set.	R/W	0
TCV	20	When set, the TCNum field specifies the number of the TC that must be traced. Otherwise, instructions from all TCs are traced when other conditions for tracing are valid.	R/W	Undefined
TCNum	19:12	Specifies the TC to trace when TCV is set. The right-most bits only are used.	R/W	Undefined
Mode	11:7	When tracing is turned on, these five bits specify what information is to be traced by the core. Each bit turns on tracing of a specific tracing mode when that bit value is a 1. If the corresponding bit is 0, then the corresponding trace (shown in the table below) is not traced by the processor. Each bit in this field is encoded as follows: Bit 7: PC Bit 8: Load address Bit 9: Store address Bit 10: Load data Bit 11: Store data	R/W	Undefined
ValidModes	6:5	This field specifies the subset of tracing that is supported by the processor. This field is encoded as follows: 00: PC tracing only 01: PC and load and store address tracing only 10: PC, load and store address, and load and store data 11: Reserved	R	Preset


**Table 2.91 TraceControl2 Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
TBI	4	This bit indicates how many trace buffers are implemented by the TCB, as follows.  0: Only one trace buffer is implemented, and the TBU bit of this register indicates which trace buffer is implemented.  1: Both on-chip and off-chip trace buffers are implemented by the TCB and the TBU bit of this register indicates to which trace buffer the traces is currently written.	R	Undefined
TBU	3	This bit denotes to which trace buffer the trace is currently being written and is used to select the appropriate interpretation of the <i>TraceControl2<sub>SyP</sub></i> field.  0: Trace data is being sent to an on-chip trace buffer 1: Trace Data is being sent to an off-chip trace buffer  This bit is loaded from <i>TCBCONTROLB<sub>OfC</sub></i> .	R	Undefined
SyP	2:0	The period (in cycles) to which the internal synchronization counter is reset when tracing is started, or when the synchronization counter has overflowed. This field is encoded as follows.  000: 2 <sup>5</sup> 001: 2 <sup>6</sup> 010: 2 <sup>7</sup> 011: 2 <sup>8</sup> 100: 2 <sup>9</sup> 101: 2 <sup>10</sup> 110: 2 <sup>11</sup> 111: 2 <sup>12</sup>  This field is loaded from the <i>PDI_SyncPeriod</i> signal when the <i>PDI_SyncOffEn</i> signal is asserted.	R	Undefined

### 2.2.11.3 Trace Control 3 Register — TraceControl3 (CP0 Register 24, Select 2)

The TraceControl3 register provides additional control and status information. This register is only implemented if the PDtrace capability is present.

**Figure 2.78 TraceControl3 Register Format**



**Table 2.92 TraceControl3 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:14	Reserved. Must be written as zeros; returns zeros on reads.	R	0

**Table 2.92 TraceControl3 Register Field Descriptions (continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
PeCOvf	13	Performance counter overflow. Setting this bit enables the trace control logic to trace a performance counter overflow.	R/W	0
PeCFCR	12	Performance counter function/call return. Setting this bit enables the trace control logic to trace a function call/return condition or an exception handler entry.	R/W	0
PeCBP	11	Performance counter hardware breakpoint. Setting this bit enables the trace control logic to trace a hardware breakpoint condition.	R/W	0
PeCSync	10	Performance counter synchronization counter expiration. Setting this bit enables the trace control logic to trace a synchronization counter expiration condition.	R/W	0
PeCE	9	Performance counter tracing enable. When set to 0, the tracing out of performance counter values as specified is disabled. To enable, this bit must be set to 1. This bit is used under software control. When trace is controlled by an external probe, this enabling is done via $\text{TraceControl3}_{PeCE}$ .	R/W	0
PeC	8	Specifies whether or not Performance Control Tracing is implemented. This bit is always set to 1 in the interAptiv processor.	R	1
0	7:3	Reserved. Must be written as zeros; returns zeros on reads.	R	0
TrIDLE	2	Trace Unit Idle. This bit indicates if the trace hardware is currently idle (not processing any data). This can be useful when switching control of trace from hardware to software and vice versa. The bit is read-only and updated by the trace hardware.	R/W	0
TRPAD	1	Trace RAM Access Disable. Disables program software access to the on-chip trace RAM using load/store instructions. This bit is loaded from $\text{TCBCONTROLB}_{TRPAD}$ .	R/W	0
FDT	0	Filtered data trace mode enable. 0: Filtered data trace mode is disabled. 1: Filtered data trace mode is enabled.	R/W	0

**2.2.11.4 User Trace Data 1 Register — UserTraceData1 (CP0 Register 23, Select 3)**

A software write to any bits in the UserTraceData1 register triggers a trace record to be written with a type indicator TU1.

These register are only implemented if the MIPS Trace capability is present.

**Figure 2.79 User Trace Data 1 Register Format**



**Table 2.93 User Trace Data 1 Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bits			
Data	31:0	Software readable/writable data. When written, this triggers a user format trace record out of the PDtrace interface that transmits the Data field to trace memory.	R/W	0

**2.2.11.5 User Trace Data 2 Register — UserDataTrace2 (CP0 Register 24, Select 3)**

A software write to any bits in the UserTraceData2 register triggers a trace record to be written with a type indicator TU2.

These register are only implemented if the MIPS Trace capability is present.

**Figure 2.80 User Trace Data 2 Register Format**



**Table 2.94 User Trace Data 2 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
Data	31:0	Software readable/writable data. When written, this triggers a user format trace record out of the PDtrace interface that transmits the Data field to trace memory.	R/W	0

**2.2.11.6 Trace Instruction Breakpoint Condition Register — TracelBPC (CP0 Register 23, Select 4)**

The TracelBPC register is used to control start and stop of tracing using an EJTAG Instruction Hardware breakpoint. The Instruction Hardware breakpoint would then be set as a trigger source and optionally also as a Debug exception breakpoint.

This register is only implemented if both Hardware breakpoints and the MIPS Trace capability are present.

**Figure 2.81 TracelBPC Register Format**



**Table 2.95 TracelBPC Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:30	Reserved. Must be written as zeros; returns zeros on reads.	R	0

**Table 2.95 TraceIBPC Register Field Descriptions(continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
PCT	29	Used to specify whether a performance counter trigger signal is generated when an EJTAG instruction breakpoint match occurs. 0: Disables performance counter trigger signal from instruction breakpoints 1: Enables performance trigger signals from instruction breakpoints	R/W	0
IE	28	Used to specify whether or not the trigger signal from EJTAG instruction breakpoint should trigger tracing functions. 0: Disables trigger signals from instruction breakpoints 1: Enables trigger signals from instruction breakpoints	R/W	0
0	27:12	Reserved. Must be written as zeros; returns zeros on reads.	R	0
IBPC3 IBPC2 IBPC1 IBPC0	11:9 9:6 5:3 2:0	The four 3-bit fields are decoded to enable different tracing modes. <a href="#">Table 2.97</a> shows the possible interpretations. Each set of 3 bits represents the encoding for the instruction breakpoint <i>n</i> in the EJTAG implementation, if it exists. If the breakpoint does not exist, then the bits are reserved, read as zero, and writes are ignored.	R/W	0

**2.2.11.7 Trace Data Breakpoint Condition Register — TraceDBPC (CP0 Register 23, Select 5)**

The TraceDBPC register is used to control start and stop of tracing using an EJTAG Data Hardware breakpoint. The Data Hardware breakpoint would then be set as a trigger source and optionally also as a Debug exception breakpoint.

This register is only implemented if both Hardware breakpoints and the MIPS Trace capability are present.

**Figure 2.82 TraceDBPC Register Format**



**Table 2.96 TraceDBPC Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:30	Reserved. Must be written as zeros; returns zeros on reads.	R	0
PCT	29	Used to specify whether a performance counter trigger signal is generated when an EJTAG data breakpoint match occurs. 0: Disables performance counter trigger signal from data breakpoints 1: Enables performance trigger signals from data breakpoints	R/W	0
DE	28	Used to specify whether the trigger signal from EJTAG data breakpoint should trigger tracing functions. 0: Disables trigger signals from data breakpoints 1: Enables trigger signals from data breakpoints	R/W	0
0	27:26	Reserved. Must be written as zeros; returns zeros on reads.	R	0

**Table 2.96 TraceDBPC Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
DBPC0	2:0	The two 3-bit fields are decoded to enable different tracing modes. <a href="#">Table 2.97</a> shows the possible interpretations. Each set of 3 bits represents the encoding for the data breakpoint <i>n</i> in the EJTAG implementation, if it exists. If the breakpoint does not exist then the bits are reserved, read as zero and writes are ignored.	R/W	0
DBPC1	5:3			

**Table 2.97 BreakPoint Control Modes: IBPC and DBPC**

Value	Trigger Action	Description
000	Unconditional Trace Stop	Unconditionally stop tracing if tracing was turned on. If tracing is already off, then there is no effect.
001	Unconditional Trace Start	Unconditionally start tracing if tracing was turned off. If tracing is already turned on, then there is no effect.
010	None	Reserved for future implementations.
011	Unconditional Trace Start (core and CM)	Unconditionally start tracing in both core and coherence manager if tracing was turned off. If tracing is already turned on, then there is no effect.
100	Identical to trigger condition 000, and in addition, dump the full performance counter values into the trace stream	If tracing is currently on, dump the full values of all the implemented performance counters into the trace stream, and turn tracing off. If tracing is already off, then there is no effect.
101	Identical to trigger condition 001, and in addition, also dump the full performance counter values into the trace stream	Unconditionally start tracing if tracing was turned off. If tracing is already turned on, then there is no effect. In both cases, dump the full values of all the implemented performance counters into the trace stream.
110	Not used	Reserved for future implementations.
111	Unconditional Trace Start (core and CM), and in addition, dump the full performance counter values into the trace stream	Unconditionally start tracing in both core and coherence manager if tracing was turned off. If tracing is already turned on, then there is no effect. Dump the full values of all the implemented performance counters into the trace stream.

## 2.2.12 User Mode Support Registers


This section contains the following hardware access registers.

- [Section 2.2.12.1, "YQMask Register \(CP0 Register 1, Select 4\)" on page 172](#)
- [Section 2.2.12.2, "Hardware Enable — HWREna \(CP0 Register 7, Select 0\)" on page 172](#)
- [Section 2.2.12.3, "UserLocal \(CP0 Register 4, Select 2\)" on page 174](#)
- [Section 2.2.12.4, "LLAddr Register \(CP0 Register 17, Select 0\)" on page 174](#)

### 2.2.12.1 YQMask Register (CP0 Register 1, Select 4)

The YQMask register is instantiated per-VPE. The interAptiv core only supports 16 mask bits.

**Figure 2.83 YQMask Register Format**



**Table 2.98 YQMask Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:16	Must be written as zero; return zero on read.	0	0
Mask	15:0	Bit vector which determines which values may be used as external state qualifiers by YIELD instructions.	R/W	0

### 2.2.12.2 Hardware Enable — HWREna (CP0 Register 7, Select 0)

The HWREna register contains a bit mask that determines which hardware registers are accessible via the **rdhwr** instruction when that instruction is executed in user mode.

[Figure 2.84](#) shows the format of the HWREna Register; [Table 2.99](#) describes the HWREna register fields. Refer to Chapter 6 for more information.

**Figure 2.84 HWREna Register Format**



**Table 2.99 HWREna Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:30	Reserved	0	0
Mask	29:0	Each bit in this field enables access by the RDHWR instruction to a particular hardware register (which may not be an actual register). If bit 'n' in this field is a 1, access is enabled to hardware register 'n'. If bit 'n' of this field is a 0, access is disabled.  <a href="#">Table 2.100</a> lists the RDHWR registers, and register number 'n' corresponds to bit 'n' in this field.	R/W	0

The Register Number column below lists the values that are placed into the 5-bit *rd* field of the **RDHWR** instruction. These register numbers correspond to the bits in the HWREna register. For example, when the *rd* field of the **RDHWR**



instruction contains a value of 00000<sub>2</sub>, and bit 0 of the HWREna register is set, the instruction return the value in the EBASE<sub>CPUNum</sub> field and places it into the GPR register specified by the *rt* field of the **RDHWR** instruction.

**Table 2.100 RDHWR Register Numbers**

Register Number (rd Value)	Mnemonic	Description										
0	CPUNum	This register provides read access to the coprocessor 0 EBase <sub>CPUNum</sub> field.										
1	SYNCI_Step	Address step size to be used with the SYNCI instruction. See that instruction's description for the use of this value. In the typical implementation, this value should be zero if there are no caches in the system which must be synchronized (either because there are no caches, or because the instruction cache tracks writes to the data cache). In other cases, the return value should be the smallest line size of the caches that must be synchronized.  For the interAptiv core, the SYNCI_Step value is 32 since the line size is 32 bytes.										
2	CC	High-resolution cycle counter. This register provides read access to the coprocessor 0 Count Register.										
3	CCRes	Resolution of the CC register. This value denotes the number of cycles between update of the register. For example: <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>CCRes Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>CC register increments every cycle</td> </tr> <tr> <td>2</td> <td>CC register increments every second cycle</td> </tr> <tr> <td>3</td> <td>CC register increments every third cycle</td> </tr> <tr> <td colspan="2" style="text-align: center;">etc.</td> </tr> </tbody> </table> In the interAptiv core, the CCRes value is 2 to indicate that the CC register increments every second core cycle.	CCRes Value	Meaning	1	CC register increments every cycle	2	CC register increments every second cycle	3	CC register increments every third cycle	etc.	
CCRes Value	Meaning											
1	CC register increments every cycle											
2	CC register increments every second cycle											
3	CC register increments every third cycle											
etc.												
4-28		These registers numbers are reserved for future architecture use. Access results in a Reserved Instruction Exception.										
29	ULR	User Local Register. This register provides read access to the coprocessor 0 UserLocal register. In some operating environments, the UserLocal register is a pointer to a thread-specific storage block.										
30-31		These register numbers are reserved for future implementation-dependent use. Access results in a Reserved Instruction Exception.										

Using the HWREna register, privileged software may select which of the hardware registers are accessible via the RDHWR instruction. In doing so, a register may be virtualized at the cost of handling a Reserved Instruction Exception, interpreting the instruction, and returning the virtualized value. For example, if it is not desirable to provide direct access to the Count register, access to that register may be individually disabled and the return value can be virtualized by the operating system.

Software may determine which registers are implemented by writing all ones to the HWREna register, then reading the value back. If a bit reads back as a one, the processor implements that hardware register.

### 2.2.12.3 UserLocal (CP0 Register 4, Select 2)

UserLocal is a read-write 32-bit register that is not interpreted by the hardware and conditionally readable by software. This register is suitable for a kernel-maintained ID whose value can be read by user-level code with `rdhwr 29`, as long as `HWRENAUL` is set.

The presence of the UserLocal register is indicated by `Config3ULRI = 1`.

**Figure 2.85 UserLocal Register Format**



**Table 2.101 UserLocal Register Field Description**

Fields		Description	Read / Write	Reset State
Name	Bits			
UserLocal	31:0	Software information that is not interpreted by hardware.	R/W	Undefined


### 2.2.12.4 LLAddr Register (CP0 Register 17, Select 0)

The LLAddr register stores the physical address (to the enclosing 32-byte block) of the target location of any LL/SC sequence. This register is readable purely for diagnostic reasons. This register is used by the hardware to properly handle LL/SC sequences by monitoring if the memory location has potentially been written between the LL and SC instructions.

Stores on this core are checked against all the LLAddr of all TCs and the internal LL bit of a TC will be cleared if a match is found. Similarly, an external intervention indicating that another core is performing a coherent write to the line will be compared to the LLAddr registers and clear the LL bit.

Similarly, an external intervention indicating that another core is performing a coherent write to the line will be compared to the LLAddr registers and clear the LL bit.

**Figure 2.86 LLAddr Register Format**



**Table 2.102 LLAddr Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
PAddr	31:5	Bits [31:5] of address used by last LL instruction.	R	Undefined
0	4:0	Reads as 0.	R	0

## 2.2.13 Kernel Mode Support Registers

This section contains the following hardware access registers.

- [Section 2.2.13.1, "Kernel Scratch Register 1 — KScratch1 \(CP0 Register 31, Select 2\)" on page 175](#)

- Section 2.2.13.2, "Kernel Scratch Register 2 — KScratch2 (CP0 Register 31, Select 3)" on page 175
- Section 2.2.13.3, "Kernel Scratch Register 3 — KScratch3 (CP0 Register 31, Select 4)" on page 175

### 2.2.13.1 Kernel Scratch Register 1 — KScratch1 (CP0 Register 31, Select 2)

KScratch1 is a read-write 32-bit register that is used by the kernel for temporary storage of information .

The presence of the KScratch1 register is indicated by  $\text{Config4}_{KScrExist[2]} = 1'b1$ .

**Figure 2.87 KScratch1 Register Format**



**Table 2.103 KScratch1 Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bits			
KScratch1	31:0	Used by the kernel for temporary storage of information.	R/W	Undefined

### 2.2.13.2 Kernel Scratch Register 2 — KScratch2 (CP0 Register 31, Select 3)

KScratch2 is a read-write 32-bit register that is used by the kernel for temporary storage of information .

The presence of the KScratch2 register is indicated by  $\text{Config4}_{KScrExist[3]} = 1'b1$ .

**Figure 2.88 KScratch2 Register Format**



**Table 2.104 KScratch2 Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bits			
KScratch2	31:0	Used by the kernel for temporary storage of information.	R/W	Undefined

### 2.2.13.3 Kernel Scratch Register 3 — KScratch3 (CP0 Register 31, Select 4)

KScratch3 is a read-write 32-bit register that is used by the kernel for temporary storage of information .

The presence of the KScratch3 register is indicated by  $\text{Config4}_{KScrExist[4]} = 1'b1$ .

**Figure 2.89 KScratch3 Register Format**



**Table 2.105 KScratch3 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
KScratch3	31:0	Used by the kernel for temporary storage of information.	R/W	Undefined

## 2.2.14 Memory Mapped Registers

This section contains the following memory mapped registers.


- [Section 2.2.14.1, "Common Device Memory Map Base Address — CDMMBase \(CP0 Register 15, Select 2\)" on page 176](#)
- [Section 2.2.14.2, "Coherency Manager Global Configuration Register Base Address — CMGCRBase \(CP0 Register 15, Select 3\)" on page 177](#)

### 2.2.14.1 Common Device Memory Map Base Address — CDMMBase (CP0 Register 15, Select 2)

The 32-bit physical base address for the Common Device Memory Map facility is defined by this register. This register only exists if *Config3CDMM* is set to one.

[Figure 2.90](#) shows the format of the CDMMBase register, and [Table 2.106](#) describes the register fields.

**Figure 2.90 CDMMBase Register**



**Table 2.106 CDMMBase Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
U	31:28	Unimplemented physical address bits. Writes are ignored, returns 0 on read	R	0
CDMM_UPPER_ADDR	27:11	Bits 31:15 of the base physical address of the common device memory-mapped registers.	R/W	17'b0_0011_1111_1000_0010
EN	10	Enables the CDMM region.  If this bit is cleared, memory requests to this address region go to regular system memory. If this bit is set, memory requests to this region go to the CDMM logic.  0: CDMM region is disabled. 1: CDMM region is enabled.	R/W	0
CI	9	If set to 1 by hardware, this bit indicates that the first 64-byte Device Register Block (DRB) of the CDMM is reserved for additional registers which manage CDMM region behavior and are not IO device registers.  This bit is always 0 in the interAptiv core since additional I/O device registers are not implemented.	R	0

**Table 2.106 CDMMBase Register Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State		
Name	Bits					
CDMMSize	8:0	This field represents the number of 64-byte Device Register Blocks (DRB) instantiated in the interAptiv core.	R	Preset		
					<b>Encoding</b>	<b>Meaning</b>
					0	1 DRB
					1	2 DRB's
					2	3 DRB's
					...	...
511	512 DRB's					

**2.2.14.2 Coherency Manager Global Configuration Register Base Address — CMGCRBase (CP0 Register 15, Select 3)**

This register is used in a multi-core environment and defines the 36-bit physical base address for the memory-mapped Coherency Manager Global Configuration Register (CMGCR) space. This register only exists if Config3<sub>CMGCR</sub> is set.

Figure 2.91 shows the format of the *CMGCRBase* register, and Table 2.107 describes the register fields.

**Figure 2.91 CMGCRBase Register**



**Table 2.107 CMGCRBase Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
U	31:28	Unimplemented physical address bits. Writes are ignored, returns 0 on read	R	0
CMGCR_BASE_ADDR	27:11	This field contains 31:15 of the base physical address of the memory-mapped Coherency Manager Global Configuration registers. The reset value is set when the core is configured using the Configuration GUI.	R	Preset
0	10:0	Must be written as zero; returns zero on read	R	0

|

# Memory Management Unit

The interAptiv core includes an optional Memory Management Unit (MMU) that translates virtual addresses to physical addresses. The MMU consists of one 4 - 12<sup>1</sup> entry Instruction TLB (ITLB) per core, one 8-entry data TLB (DTLB) per core, and one 16, 32, 48, or 64 dual-entry Joint TLB (JTLB) per VPE. If the functionality of a full MMU is not required, the interAptiv core also provides a Memory Protection Unit (MPU) that controls whether read, write, or execute access is permitted for a given address and causes an exception if a unauthorized access is attempted. For more information, refer to the Memory Protection Unit chapter.

This chapter contains the following sections:

- [Section 3.1, "Introduction" on page 179](#)
- [Section 3.2, "Memory Management Unit Architecture" on page 180](#)
- [Section 3.3, "Overview of Virtual-to-Physical Address Translation" on page 182](#)
- [Section 3.4, "Relationship of TLB Entries and CP0 Registers" on page 186](#)
- [Section 3.5, "Enhanced Virtual Address" on page 192](#)
- [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode" on page 201](#)
- [Section 3.7, "Indexing the JTLB" on page 216](#)
- [Section 3.8, "Hardwiring JTLB Entries" on page 217](#)
- [Section 3.9, "JTLB Random Replacement" on page 217](#)
- [Section 3.10, "TLB Exception Handling" on page 219](#)
- [Section 3.11, "Exception Base Address Relocation" on page 226](#)
- [Section 3.12, "TLB Duplicate Entries" on page 227](#)
- [Section 3.13, "Modes of Operation" on page 227](#)
- [Section 3.14, "TLB Instructions" on page 240](#)

## 3.1 Introduction

The MMU translates a virtual address to a physical address before the request is sent to the cache controllers for tag comparison or to the bus interface unit for an external memory reference. Virtual-to-physical address translation is especially useful for operating systems that must manage physical memory to accommodate multiple tasks active in the same memory, and possibly in the same virtual address space. The MMU also enforces the protection of memory areas and defines the cache protocols.

An Enhanced Virtual Address (EVA) scheme allows physical addresses above 0.5 GB to be translated without using the *HighMem* protocol. The EVA scheme allows for user space segments up to 3.5 GB.

---

1. Three dedicated entries plus one entry per thread context (TC) up to nine threads.

## 3.2 Memory Management Unit Architecture

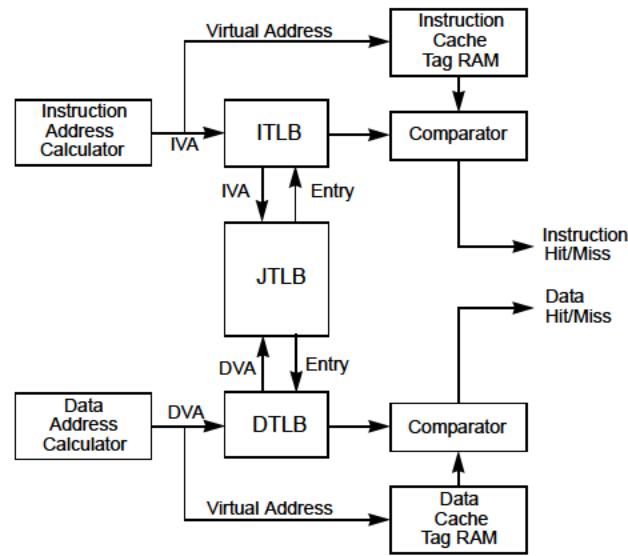
The Memory Management Unit (MMU) in the interAptiv core consists of three address-translation lookaside buffers (TLB):

- 4 - 12 entry Instruction TLB (ITLB)
- 8-entry Data TLB (DTLB)
- 16, 32, 48, or 64 dual-entry Joint Translation Lookaside Buffer (JTLB) per VPE

When an instruction address is to be translated, the ITLB is accessed first. If the translation is not found, the JTLB is accessed. If there is a miss in the JTLB, an exception is taken. Similarly, when a data reference is to be translated, the DTLB is accessed directly. If the address is not present in the DTLB, the JTLB is accessed. If there is a miss in the JTLB, an exception is taken.

Figure 3.1 shows an overview of the interAptiv MMU architecture.

Figure 3.1 Overview of MMU Architecture in the interAptiv Core



### 3.2.1 Instruction TLB (ITLB)

The ITLB is a 4 - 12 entry associative TLB dedicated to performing translations for the instruction stream. The ITLB maps only 4 KB or 1 MB pages (or subpages). Note that all pages copied into the ITLB are on either a 4KB or 1MB boundary. If the page size is less than 1MB, then portions of the page are copied into the ITLB on 4 KB boundaries. If the page size is larger than 1 MB, then portions of the page are copied into the ITLB on 1 MB boundaries. Therefore, the mapping can be for an entire page, or a subpage of a larger mapping.

The ITLB contains three entries shared by all TCs. Each TC also has a private entry. A refill from the JTLB goes into the least recently used of the 3 shared entries. Each TC compares the age of the shared entry that was replaced against their private entry and captures it if the private entry was less recently used.

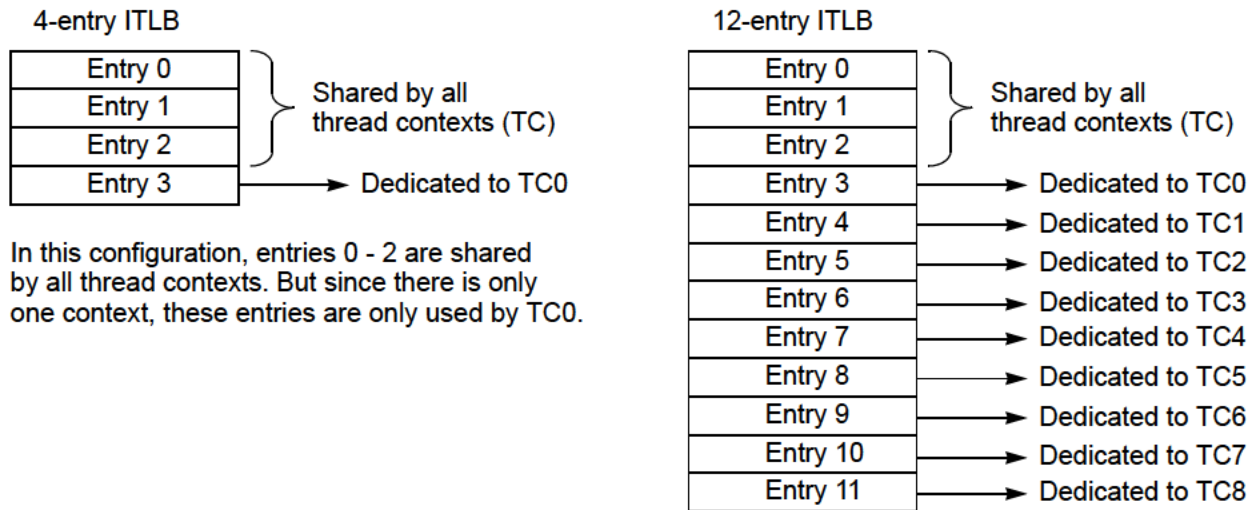
The ITLB is managed by hardware and is transparent to software. The larger JTLB is used as a backup structure for the ITLB. If a fetch address cannot be translated by the ITLB, the JTLB attempts to translate it in the following clock cycle or when available. If successful, the translation information is copied into the ITLB for future use.



The ITLB is functionally invisible to software and its entries are automatically refilled from the JTLB when required, and automatically cleared whenever the associated JTLB is updated.

Figure 3.2 shows the minimum and maximum number of ITLB organization based on the number of thread contexts.

**Figure 3.2 Minimum and Maximum ITLB Organizations**



### 3.2.2 Data TLB (DTLB)

The DTLB is an 8-entry associative TLB dedicated to performing translations for the data stream. The DTLB maps only 4 KB or 1 MB pages.

The DTLB is managed by hardware and is transparent to software. The larger JTLB is used as a backup structure for the DTLB. If a load/store address cannot be translated by the DTLB, the JTLB/SegCtl logic attempts to translate it in the following clock cycle or when available. If successful, the translation information is copied into the DTLB for future use. Note that the JTLB is used during mapped accesses. The SegCtl registers are used during unmapped accesses.

The DTLB is functionally invisible to software and entries are automatically refilled from the JTLB when required, and automatically cleared whenever the associated JTLB is updated.

### 3.2.3 Joint TLB (JTLB)

The 16 - 64 dual-entry, fully associative Joint TLB maps 32 - 128 virtual pages to their corresponding physical addresses. The purpose of the JTLB is to translate virtual addresses and their corresponding ASID into a physical memory address. The translation is performed by comparing the upper bits of the virtual address (along with the ASID bits) against each of the entries in the *tag* portion of the JTLB structure. This structure is used to translate both instruction and data virtual addresses. Note that the JTLB is implemented on a per-VPE basis, so there is one JTLB per VPE.

The JTLB is organized as 64 pairs of even and odd entries. The JTLB implements the following page sizes:

4K, 16K, 64K, 256K, 1M, 4M, 16M, 64M, and 256M

The JTLB is organized in pairs of page entries to minimize its overall size. Each virtual *tag* entry corresponds to two physical data entries, an even page entry and an odd page entry. The highest order virtual address bit not participating in the tag comparison is used to determine which of the two data entries is used. Since page size can vary on a page-

pair basis, the determination of which address bits participate in the comparison and which bit is used to make the even-odd selection must be done dynamically during the TLB lookup.

The *PageMask* register is loaded with the desired page size, which is then entered into the TLB when a new entry is written. Thus, operating systems can provide special-purpose maps. For example, a typical frame buffer can be memory-mapped with only one TLB entry. Software can determine which page sizes are supported by writing all ones to the *PageMask* register, then reading the value back.

The JTLB entries are controlled through select CP0 registers. Refer to [Section 3.4, "Relationship of TLB Entries and CP0 Registers"](#) for more information.

### 3.3 Overview of Virtual-to-Physical Address Translation

Converting a virtual address to a physical address begins by comparing the virtual address from the processor with the virtual addresses in the TLB. There is a match when the VPN of the address is the same as the VPN field of the TLB entry after masking out the bits specified by the entries page size, and either:

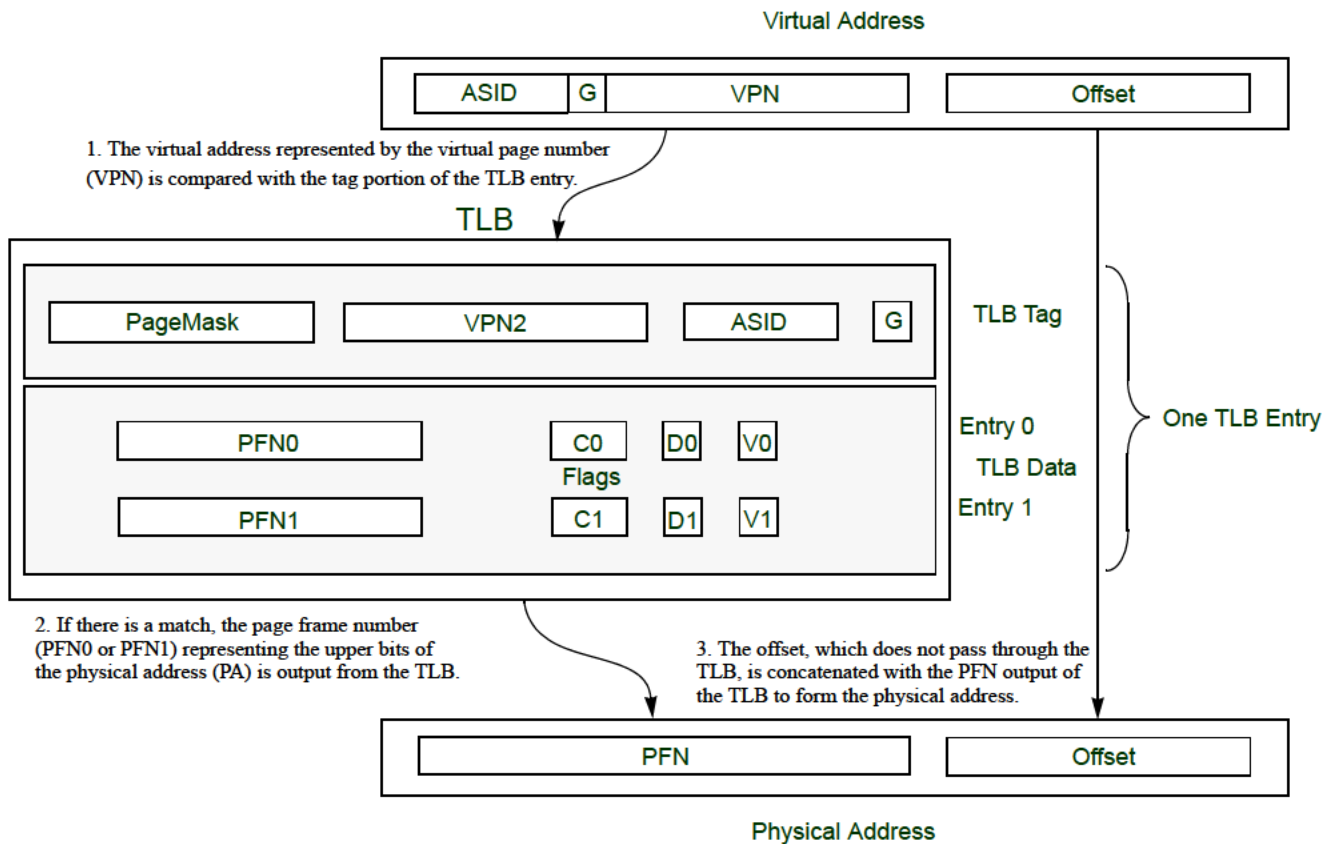
- The Global (G) bit of both the even and odd pages of the TLB entry is set, or
- The Global (G) bit is cleared and the ASID field of the virtual address is the same as the ASID field of the TLB entry

This match is referred to as a TLB *hit*. If there is no match, a TLB *Refill* exception is taken by the processor, and software is allowed to refill the TLB from a page table of virtual/physical addresses in memory.

[Figure 3.3](#) shows the translation of a virtual address into a physical address. In this figure, the virtual address is extended with an 8-bit ASID, which reduces the frequency of TLB flushes during a context switch. This 8-bit ASID contains the number assigned to that process.

Note that the various register fields used during a TLB translation are managed via CP0 registers as described in [Section 3.4, "Relationship of TLB Entries and CP0 Registers"](#).

**Figure 3.3 Overview of Virtual to Physical Address Translation**



If there is a virtual address match in the TLB, the Physical Frame Number (PFN) is output from the TLB and concatenated with the *Offset* to form the physical address. The *Offset* represents an address within the page frame space. As shown in Figure 3.3, the *Offset* does not pass through the TLB. Note that if the G bit is set, the ASID is ignored and the TLB compares only the VPN portion of the virtual address. The G bit is a logical AND of the G bit in the *EntryLo0* and *EntryLo1* registers.

Figure 3.4 shows a flow diagram of the address translation process for a 4 KByte page size. The width of the *Offset* is defined by the page size. The remaining 20 bits of the address represent the virtual page number (VPN).

**Figure 3.4 32-bit Virtual Address Translation — 4 KB Page Size**

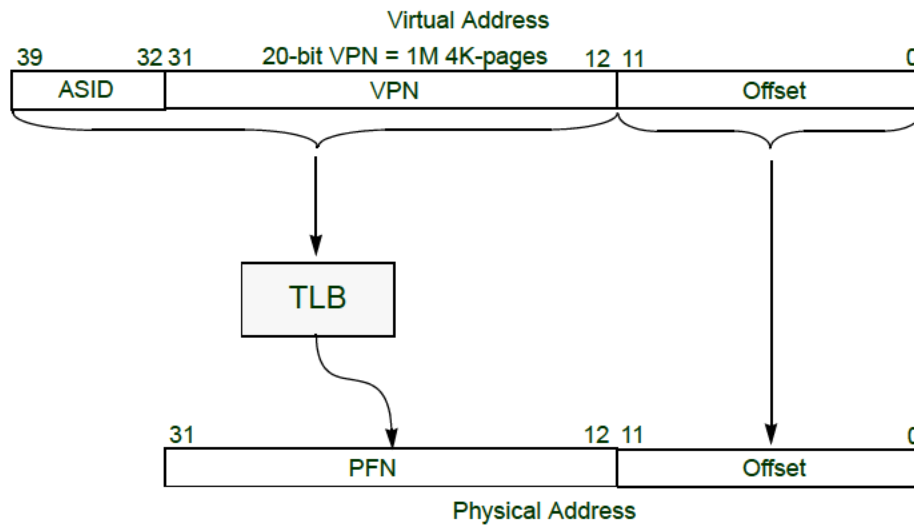



Figure 3.5 shows a flow diagram of the address translation process for a 16 MByte page size. The width of the *Offset* is defined by the page size. The remaining 8 bits of the address represent the virtual page number (VPN).

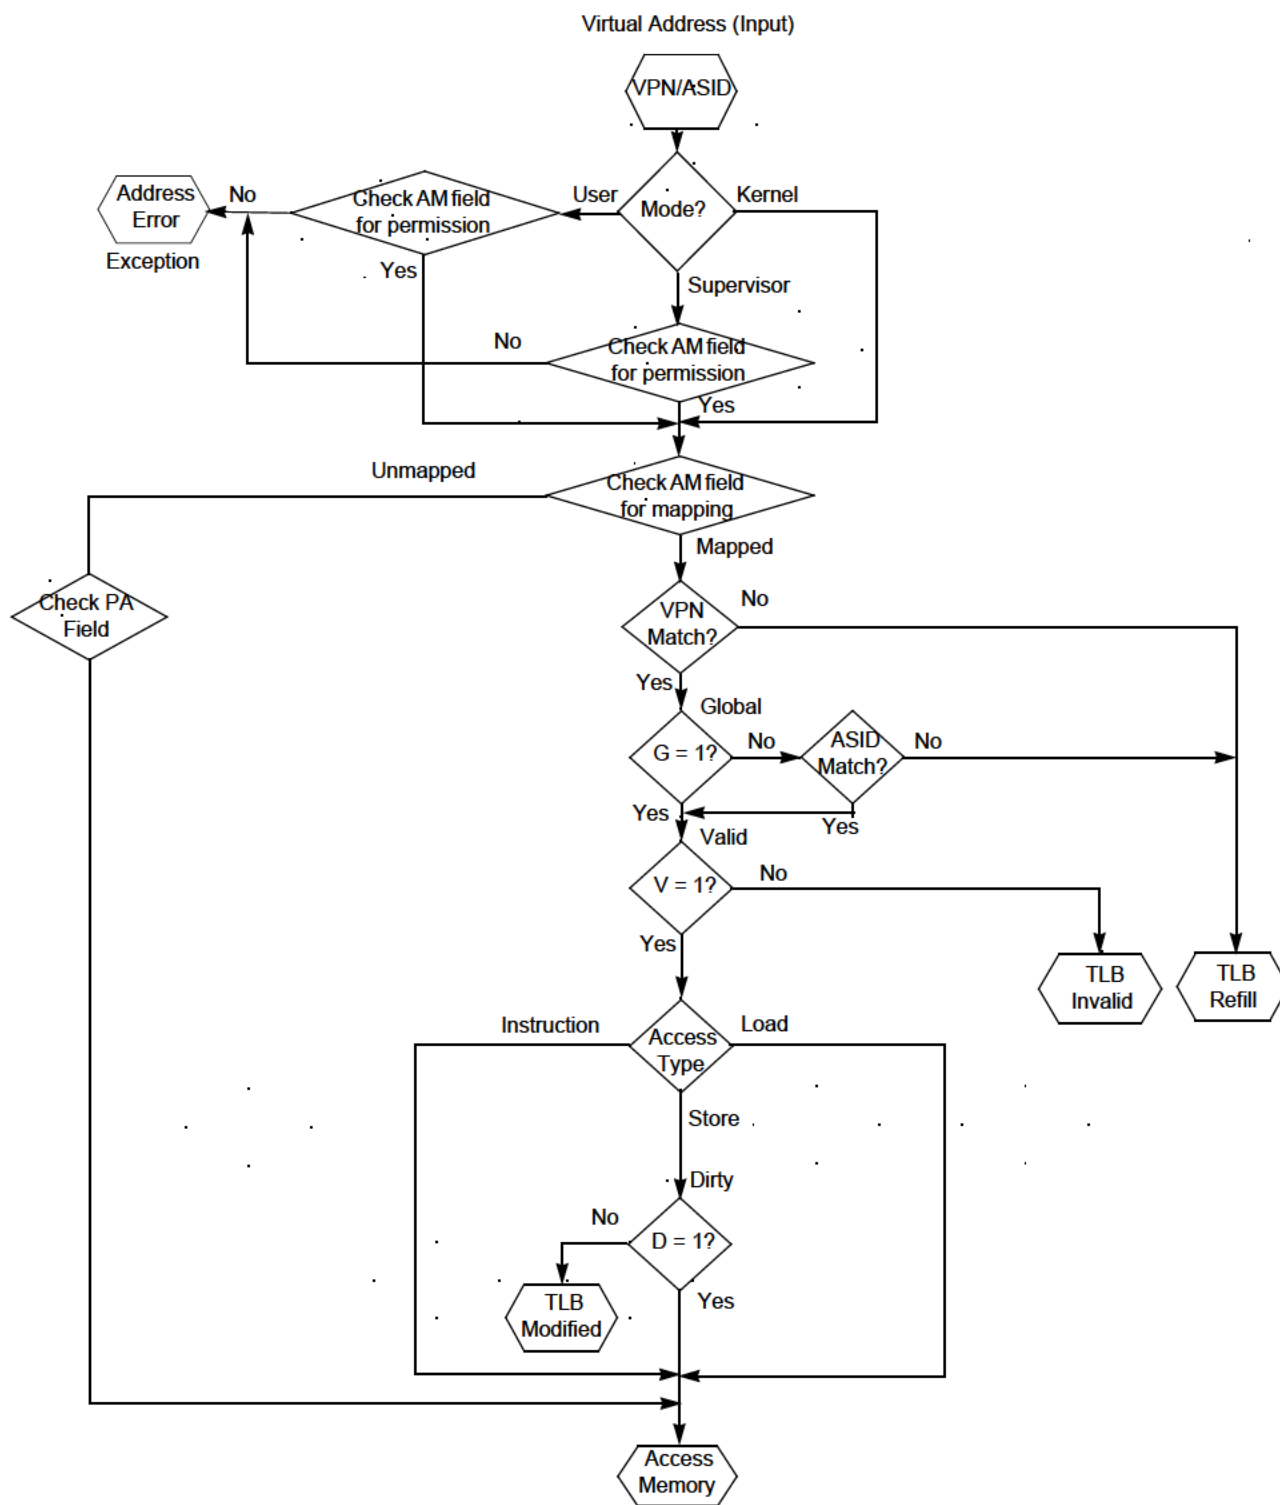
**Figure 3.5 32-bit Virtual Address Translation — 16 MB Page Size**



### 3.3.1 Address Translation Flow

During an address translation, the hardware checks for various conditions such as the addressing mode (user, kernel etc.), access permissions based on the mode, the access type (load/store, etc), and the state of selected bits in the TLB entry. If one or more of the conditions for translation are not met, a TLB exception is taken. This concept is shown in Figure 3.6.

Figure 3.6 Address Translation Flow



### 3.4 Relationship of TLB Entries and CP0 Registers

Each TLB entry in the JTLB consists of a tag portion and dual-data portion as shown in [Figure 3.7](#). In this figure, the following registers are used to manage the TLB entries.

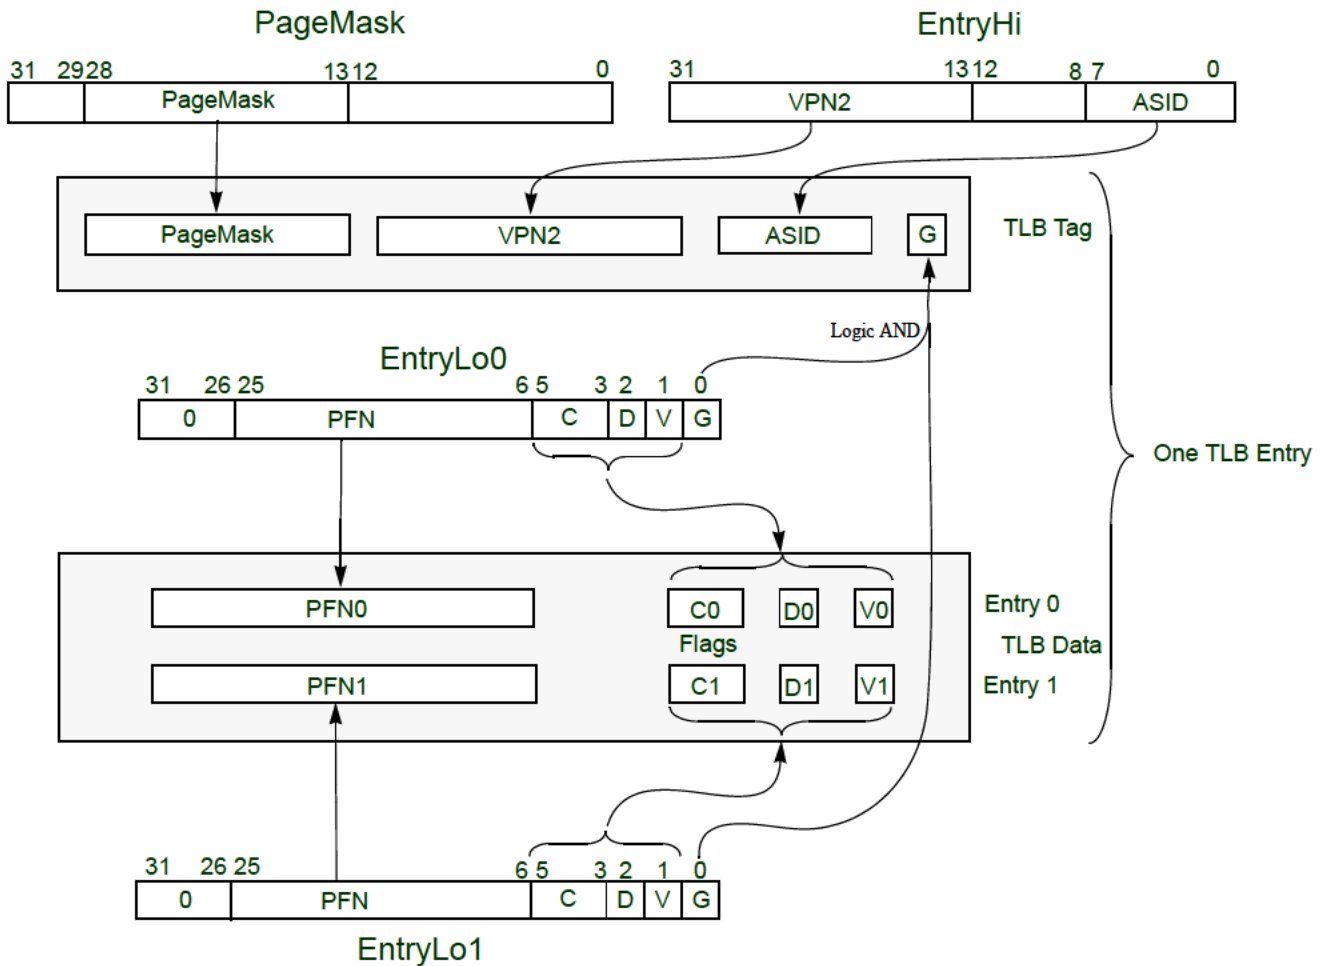
- *EntryLo0* (CP0 Register 2, Select 0)
- *EntryLo1* (CP0 Register 3, Select 0)
- *EntryHi* (CP0 Register 10, Select 0)
- *PageMask* (CP0 Register 5, Select 0)

In order to fill an entry in the JTLB, software executes a **TLBWI** or **TLBWR** instruction (see [Section 3.14](#)). Prior to invoking one of these instructions, the CP0 registers listed above must be updated with the information to be written to the TLB entry:

- PageMask is set in the CP0 *PageMask* register.
- VPN2, and ASID are set in the CP0 *EntryHi* register.
- PFN0, C0, D0, V0, and G bits are set in the CP0 *EntryLo0* register.
- PFN1, C1, D1, V1, and G bits are set in the CP0 *EntryLo1* register.

These register fields and their relationship to a TLB entry is described in the following subsections.

**Figure 3.7 Relationship Between CP0 Registers and TLB Entries**



### 3.4.1 TLB Tag Entry

The tag portion of the TLB entry contains the fields necessary to match an incoming address against that entry. This section describes each field of the TLB tag entry shown in Figure 3.7.

#### 3.4.1.1 VPN2 Field

The virtual page number (VPN) contains the high bits of the program (virtual) address. The 'VPN2' designation indicates that this address is for a double-page-size virtual region which will map to a pair of physical pages. The VPN2 field is generated using the *EntryHi* register.

Note that on a TLB-related exception, the *VPN2* field is automatically set to the virtual address that was being translated when the exception occurred. If the outcome of the exception handler is to find and install the translation to that address, the *VPN2* field will already contain the correct value.

#### 3.4.1.2 ASID Field

The address space identifier (ASID) helps to reduce the frequency of TLB flushing on a context switch. The ASID field extends the virtual address with an 8-bit memory space identifier assigned by the operating system. The ASID allows translations for multiple different applications to co-exist in the TLB (in Linux, for example, each application

has different code and data lying in the same virtual address region). The ASID field is generated using the *EntryHi* register.

### 3.4.1.3 PageMask Field

The size of the tag can be configured using the ‘PageMask’ field. This field determines how many incoming address bits to match. For the TLB, the interAptiv core allows page sizes of 4 Kbytes up to 256 Mbytes in multiples of four. The *PageMask* field is generated using the *PageMask* register.

In the *PageMask* field, a ‘1’ on a given bit means "don't compare this address bit when matching this address". However, only a restricted range of *PageMask* values are legal. The values must start with "1"s filling the *PageMask* field from the low-order bits upward, two at a time. A list of valid 32-bit *PageMask* register values, the corresponding binary value of the PageMask[28:13] field, and the corresponding page size is shown in Table 3.1. For the PageMask[28:13] field, note that the bits are set two at a time from the least significant bit (LSB) to the most significant bit (MSB).

**Table 3.1 PageMask Value and Corresponding Page Size**

32-bit PageMask Register Value	PageMask[28:13]	Page Size	Even/Odd Bank Select Bit
0x0000_0000	0x00_0000_0000_0000_00	4 KBytes	VAddr[12]
0x0000_6000	0x00_0000_0000_0000_11	16 KBytes	VAddr[14]
0x0001_E000	0x00_0000_0000_0011_11	64 KBytes	VAddr[16]
0x0007_E000	0x00_0000_0000_1111_11	256 KBytes	VAddr[18]
0x001F_E000	0x00_0000_1111_1111_11	1 MByte	VAddr[20]
0x007F_E000	0x00_0011_1111_1111_11	4 MBytes	VAddr[22]
0x01FF_E000	0x00_0011_1111_1111_11	16 MBytes	VAddr[24]
0x07FF_E000	0x00_1111_1111_1111_11	64 MBytes	VAddr[26]
0x1FFF_E000	0x11_1111_1111_1111_11	256 MBytes	VAddr[28]

### 3.4.1.4 Global (G) Bit

The ‘G’ (global) bit in the tag entry is a logical AND between the *G* bits of the *EntryLo0* and *EntryLo1* registers. When set, it causes addresses to match regardless of their ASID value, thus defining a part of the address space which will be shared by all applications. For example, Linux applications share some ‘kseg2’ space used for kernel extensions.

Note that since the *G* bit in the TLB tag entry is a logical AND between two *G* bits, software must be sure to set *EntryLo0<sub>G</sub>* and *EntryLo1<sub>G</sub>* to the same value.

## 3.4.2 TLB Data Entry

The data portion of the TLB entry contains the data and associated flag bits for the corresponding tag entry. This section describes each field of the TLB data entry shown in Figure 3.7.

### 3.4.2.1 Page Frame Number (PFN)

The Page Frame Number (PFN) contains the high-order bits of the physical address. For a 4 KByte page size, the 20-bit *PFN*, together with the lower 12 bits of address that are not translated, make up the 32-bit physical address.



### 3.4.2.2 Flag Fields (C, D, V)

These flag bits contain information about the translated address. All of these bits are generated by the *EntryLo0* and *EntryLo1* registers.

*C Field:* This field contains the cacheability attributes for the corresponding TLB entry. It indicates how to cache data for this page. Pages can be marked cacheable, uncacheable, coherent, non-coherent, uncached accelerated, write-back, write-allocate, etc.

*D bit:* The "dirty" flag. Setting this bit indicates that the page has been written, and/or is writable. If this bit is a one, stores to the page are permitted. If this bit is a cleared, stores to the page cause a *TLB Modified* exception. Software can use this bit to track pages that have been written to. When a page is first mapped, this bit should be cleared. It is set on the first write that causes an exception.

*V bit:* The "valid" flag. Indicates that the TLB entry, and thus the virtual page mapping, are valid. If this bit is set, accesses to the page are permitted. If this bit is a zero, accesses to the page cause a *TLB Invalid* exception.

### 3.4.3 Address Translation Examples

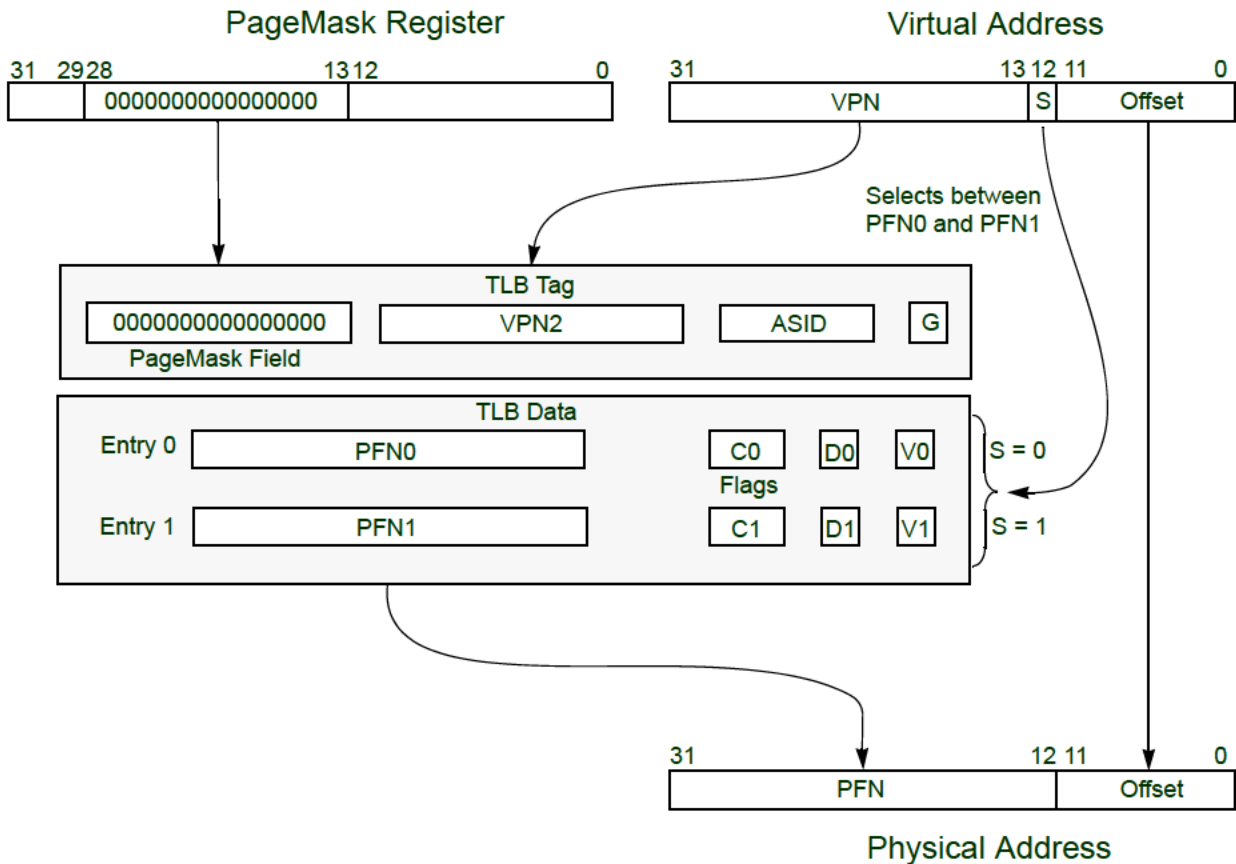
As shown in [Figure 3.7](#), there are two PFN values for each tag match. Which of them is used is determined by the lowest-order bit of the VPN field of the address. So in standard form (using 4KByte pages) each entry translates an 8KByte region of virtual address, but each 4Kbyte page can be mapped onto any physical address (with any permission flag bits). This concept is described in the following subsections.

#### **4 KByte Page Size Example**

In a 4KB page size, 12 address bits are required to select an entry within the page. Therefore, 12 bits of the virtual address are used for the offset into the page. The upper 20 bits of the virtual address are used as a pointer to the page table. With a 4 KByte page size, this allows support for up to 1M page table entries.

The upper 20 bits of virtual address pass through the TLB to generate the corresponding physical address. As described in [Section 3.3](#), the interAptiv core implements a dual-entry JTLB scheme, where each TLB tag corresponds to two data entries. To select between these two entries, hardware reads the low-order bit of the VPN (first bit after the offset, shown as the S bit in the figure below). In a 4 KByte page example, this equates to bit 12. This is shown in [Figure 3.8](#).

**Figure 3.8 Selecting Between PFN0 and PFN1 — 4 KByte Page Size**



As shown in Figure 3.8, the *PageMask* field is derived from the *PageMask* register and is used to determine the page size for the application. Since the interAptiv core supports JTLB page sizes in multiples of four (4 KByte, 16 KByte, 64 KByte, etc. up to 256 MByte), page masking is done in pairs. During translation, hardware checks the VPN against the contents of the *PageMask* field to determine the page size, and therefore how many VPN bits to compare. Refer to Table 3.1 for a list of valid *PageMask* values.

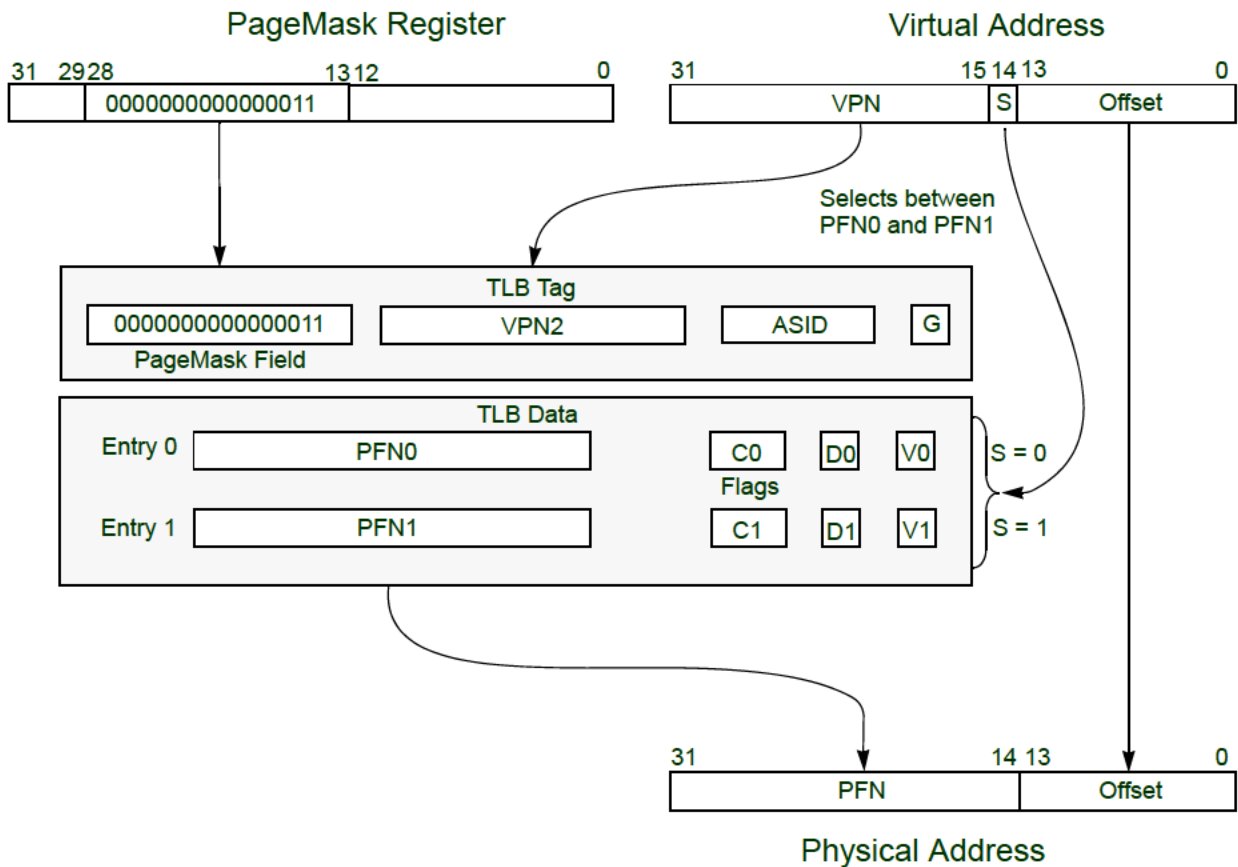
In the above example, all of the *PageMask* field bits are 0, indicating a 4 KByte page size. For a 16 KByte page size, bits 12 and 13 of the *PageMask* field would be set. This concept is described below.

### 16 KByte Page Size Example

In a 16 KByte page size, 14 address bits are required to select an entry within the page. Therefore, 14 bits of the virtual address are used for the offset into the page. The upper 18 bits of the virtual address are used as a pointer to the page table. With a 16 KByte page size, this allows support for up to 256K page table entries.

The upper 18 bits of virtual address pass through the TLB to generate the corresponding physical address. As described in Section 3.3, the interAptiv core implements a dual-entry JTLB scheme, where each TLB tag corresponds to two data entries. To select between these two entries, hardware reads the low-order bit of the VPN (first bit after the offset, shown as the S bit in the figure below). In a 16 KByte page example, this equates to bit 14. This is shown in Figure 3.9.

Figure 3.9 Selecting Between PFN0 and PFN1 — 16 KByte Page Size



As shown in Figure 3.9, the *PageMask* field is used to determine the page size for the application. During translation, hardware checks the VPN against the contents of the *PageMask* field to determine the page size, and therefore how many VPN bits to compare. In the above example, the lower 2 bits of the *PageMask* field bits are 11, indicating a 16 KByte page size. Refer to Table 3.1 for a list of valid *PageMask* values.

### 3.5 Enhanced Virtual Address

Traditional MIPS virtual memory support divides up the virtual address space into fixed size segments, each with fixed attributes and access privileges. Such a scheme limits unmapped kernel access to 512 MBytes, the size of *kseg0/kseg1*. Furthermore, application sizes are growing beyond the 2GB limit imposed by the *useg* user segment.

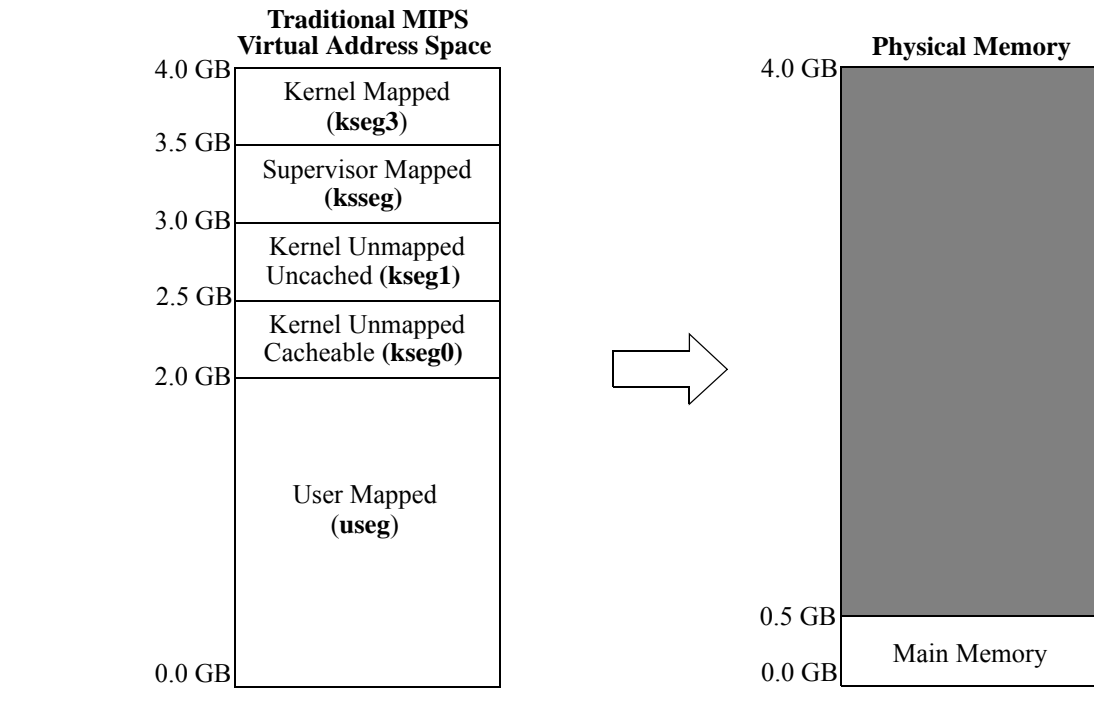
*Programmable Memory Segmentation* relaxes these limitations. The size of virtual address space segments can be programmed, as can their attributes and privilege access. With this ability to overlap access modes, *kseg0* can now be extended up to 3.0GB<sup>2</sup>, leaving at least one 1.0GB segment for mapped kernel accesses. This extended *kseg0* is called *xkseg0*. *xkseg0* overlaps with *useg*, because segments in *xkseg0* are programmed to support mapped user accesses and unmapped kernel accesses. Consequently, user space is equal to the size of *xkseg0*, which can be up to 3.0GB.

To allow for efficient kernel access to user space, new load and store instructions have been defined which allow kernel mapped access to *useg*. The new instructions, along with *Programmable Memory Segmentation*, are requirements for the scheme, called *Enhanced Virtual Address* or EVA, which allows for more efficient use of 32b address space.

#### 3.5.1 Virtual and Physical Address Maps

In previous generation MIPS32 processors, the address map was fixed as shown in [Figure 3.10](#). In this architecture, physical memory is limited by *kseg0* to 0.5GB, the amount of kernel unmapped cached address space. This memory must also be shared by the I/O and kernel, thus in reality less than 0.5GB is available to any user process.

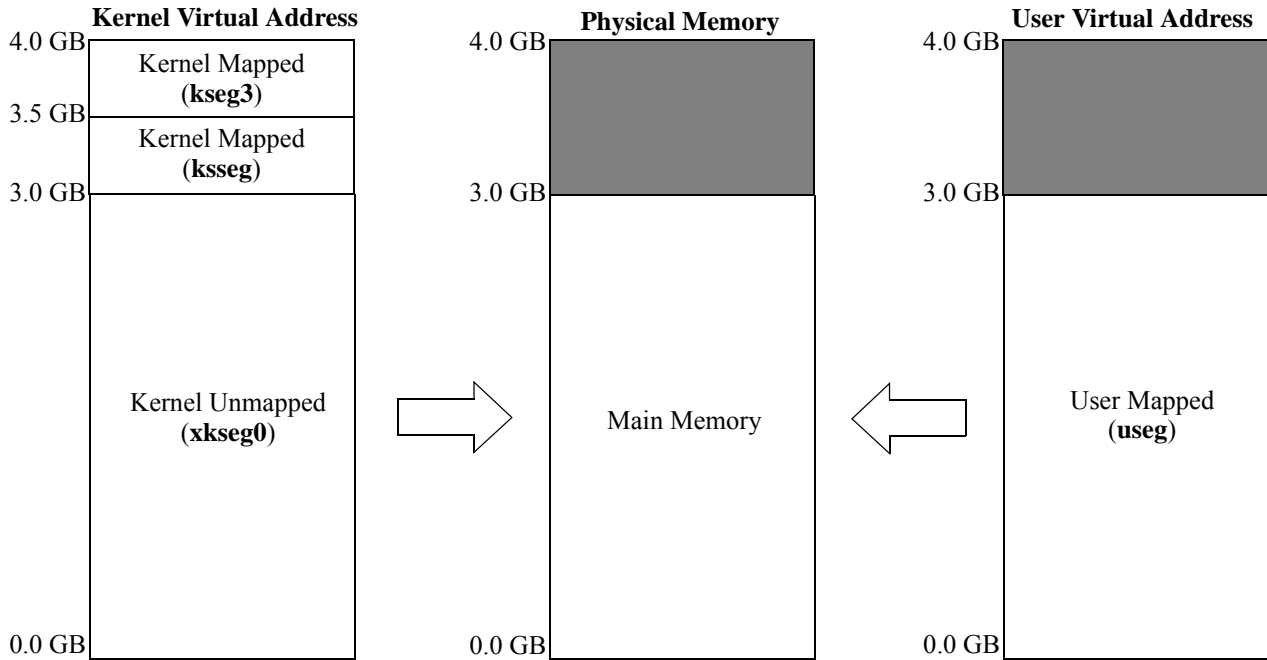
**Figure 3.10 Traditional Virtual Address Mapping in Previous Generation MIPS32 Processors**



2. If necessary, *xkseg0* can be extended to 3.5GB, allowing 0.5GB for Kernel mapped virtual address space (now *kseg2*).

Figure 3.11 shows an example of how the traditional MIPS kernel virtual address space can be remapped using programmable memory segmentation to facilitate the EVA scheme. As a result of defining the larger kernel segment as *xkseg0*, the kernel has unmapped access to the lower 3GB of the virtual address space. The larger user segment could be defined because the address space is not statically partitioned. This allows for a total of 3.5GB of DRAM to be supported in the system.

**Figure 3.11 Example of Remapping Kernel and User Virtual Address Space Using EVA**



Note that *xkseg0* is equivalent to the previous *kseg0* space in that it is a kernel unmapped, cacheable region.

### 3.5.2 Initial EVA Configuration Parameters

During build time, the customer selects the EVA parameters through the GUI. These selections are registered into the CM2, which drives the core EVA pins with the appropriate value. For a listing of EVA related pins and their function, refer to [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#).

### 3.5.3 Programmable Segmentation Control

Programmable segmentation allows for the virtual address space segments to be programmed with different access modes and attributes. Control of the 4GB of virtual address space is divided into six segments that are controlled

using three CP0 registers; *SegCtl0* through *SegCtl2*. Each register has two 16-bit fields. Each field controls one of the six address segments as shown in [Table 3.2](#).

**Table 3.2 Programmable Segmentation Register Interface**

Register	CP0 Location	Memory Segment	Register Bits	Virtual Address Space Controlled	Virtual Address Range
SegCtl2	Register 5 Select 4	CFG5	31:16	0.0 GB to 1.0 GB	0x0000_0000 - 0x3FFF_FFFF
		CFG4	15:0	1.0 GB to 2.0 GB	0x4000_0000 - 0x7FFF_FFFF
SegCtl1	Register 5 Select 3	CFG3	31:16	2.0 GB to 2.5 GB	0x8000_0000 - 0x9FFF_FFFF
		CFG2	15:0	2.5 GB to 3.0 GB	0xA000_0000 - 0xBFFF_FFFF
SegCtl0	Register 5 Select 2	CFG1	31:16	3.0 GB to 3.5 GB	0xC000_0000 - 0xDFFF_FFFF
		CFG0	15:0	3.5 GB to 4.0 GB	0xE000_0000 - 0xFFFF_FFFF

Each 16-bit field listed in the above table contains information on the corresponding memory segment such as address range (for kernel unmapped segments), access mode, and cache coherency attributes. [Table 3.3](#) describes the 16-bit configuration fields (CFG0 - CFG5) defined in the *SegCtl0* - *SegCtl2* registers.

**Table 3.3 CFG (Segment Configuration) Field Descriptions**

CFGn Fields		Description
Name	Bits	
PA	15:9 and 31:25	Physical address bits 31:29 for segment, for use when unmapped. These bits are used when the virtual address space is configured as kernel unmapped to select the segment in memory to be accessed.  For segments 0, 2, and 4, CFG[11:9] correspond to physical address bits 31:29. CFG[15:12] correspond to physical address bits 35:32 in a 36-bit addressing scheme and are reserved for future use. The state of CFG[15:12] are read/write and can be programmed, but these bits are not driven onto the address bus.  For segments 1, 3, and 5, CFG[27:25] correspond to physical address bits 31:29. CFG[31:28] correspond to physical address bits 35:32 in a 36-bit addressing scheme and are reserved for future use.  These bits are not used by the CFG4 and CFG5 spaces listed in <a href="#">Table 3.2</a> above when these segments are programmed to be kernel mapped and the physical address is determined by the TLB. They are also not used for any of the user mapped (useg) region for the same reason.
Reserved	8:7 and 24:23	Reserved.
AM	6:4 and 22:20	Access control mode. See <a href="#">Section 3.5.3.5, "Setting the Access Control Mode"</a> .  For programmable segmentation, these bits are set as shown in <a href="#">Table 3.5</a> .  Bits 6:4 correspond to segments 0, 2, and 4. Bits 22:20 correspond to segments 1, 3, and 5.

**Table 3.3 CFG (Segment Configuration) Field Descriptions**

CFGn Fields		Description
Name	Bits	
EU	3 and 19	Error condition behavior. Segment becomes unmapped and uncached when $Status_{ERL} = 1$ .  Bit 3 corresponds to segments 0, 2, and 4. Bit 19 corresponds to segments 1, 3, and 5.
C	2:0 and 18:16	Cache coherency attribute, for use when unmapped.  For programmable segmentation, these bits are set as shown in <a href="#">Table 3.5</a> .  Bits 2:0 correspond to segments 0, 2, and 4. Bits 18:16 correspond to segments 1, 3, and 5.

### 3.5.3.1 Cache Coherency Attribute Control and the Segmentation Control Registers

The CP0 memory segmentation control registers (*SegCtl0* - *SegCtl2*) are new to the MIPS32 R3 architecture and are used to control the size and function of the various memory map segments in the interAptiv core.

In the previous generation MIPS32 R2 architecture, only the cache coherency attributes of the *kseg0* memory segment could be modified by the user. All other parameters were fixed. In the MIPS32 R3 interAptiv core, each segmentation control register (*SegCtl0* - *SegCtl2*) contains its own cache coherency attribute field to allow for maximum flexibility when assigning cacheability attributes to the memory. However, since existing code will not be aware of the existence of the *SegCtl0* - *SegCtl2* registers, the interAptiv core allows a mechanism for the cache coherency attributes (CCA) of *kseg0* to be set either by the *Config.k0* field, as is done in the MIPS32 R2 architecture, or by the CFG3\_C field (bits 18:16) of the *SegCtl1* register. This allows existing code to configure virtual memory for the legacy setting.

To control where the cache coherency attributes for the memory are taken from, the *Config5.k* bit has been added to the CP0 *Config5* register. If the *Config5.k* bit is cleared, the cache coherency attributes for *kseg0* are derived from the 3-bit *Config.k0* field of the CP0 *Config* register. This can be done when booting the interAptiv core using existing code. If the *Config5.k* bit is set, the cache coherency attributes are derived from the 3-bit *SegCtlx.CFGy\_C* field of the segmentation control registers (where *x* indicates the segmentation control register number 0 - 2, and *y* indicates memory segments 0 - 5). When configured for EVA, each of the six memory segments can be individually defined with its own cache coherency attributes. Refer to [Section 2.2.3, "Memory Segmentation Registers"](#) in the CP0 chapter for more information on the segmentation control registers.

The initial programming of *Config5.k* bit is determined by the state of the *SI\_EVAReset* pin at reset as described in [Section 3.5.3.3, "Setting the Memory Addressing Scheme — SI\\_EVAReset and CONFIG5.K"](#).

### 3.5.3.2 Functions of the Config5.K Bit

The *Config5.k* bit effects the cache coherency attributes, the boot exception vector overlay mechanism, and the location of the exception vector as described below.

When the *Config5.k* bit is cleared, the following events occur:

1. The 3-bit *Config.k0* field is used to set the cache coherency attributes for the *kseg0* region (0x8000\_0000 - 0x9FFF\_FFFF). See [Section 3.5.3.1](#) above for more information.
2. Hardware creates two boot overlay segments, one for *kseg0* and one for *kseg1*. Refer to [Section 3.6.3, "Mapping of the Boot Exception Vector in the Legacy Configuration"](#) for more information.

- Hardware ignores the state of bits 31:30 of the *EBase* register as well as the *SI\_ExceptionBase[31:30]* pins. Instead, hardware forces these bits to a value of 2'b10, causing the vectors to reside in kseg0/kseg1 space. Refer to [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#) for more information.

When the *Config5.K* bit is set, the following events occur:

- The 3-bit *Config.K0* field is ignored and the cache coherency attributes are derived from the *CFGn\_C* fields of the various segmentation control registers (*SegCtl0 - SegCtl2*). Refer to [Section 3.5.3.3, "Setting the Memory Addressing Scheme — SI\\_EVAREset and CONFIG5.K"](#) for more information.
- Hardware creates one boot overlay segment that can reside anywhere in virtual address space. Refer to [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#) for more information.
- The exception vectors are not forced to reside in kseg0/kseg1. Rather, bits 31:30 of the *EBase* register, as well as the *SI\_ExceptionBase[31:30]* signals are used to place the exception vectors anywhere within virtual address space. Refer to [Section 3.11, "Exception Base Address Relocation"](#) for more information.

### 3.5.3.3 Setting the Memory Addressing Scheme — SI\_EVAREset and CONFIG5.K

The *SI\_EVAREset* pin determines the addressing scheme and whether the device boots up in the legacy setting or the EVA setting. The legacy setting is defined as having the traditional MIPS virtual memory map used in previous generation processors. The EVA setting places the device in the enhanced virtual address configuration, where the initial size and function of each segment in the virtual memory map is determined from the segmentation control registers (*SegCtl0 - SegCtl2*).

If the *SI\_EVAREset* pin is deasserted at reset, the interAptiv core comes up in the legacy configuration and hardware takes the following actions:

- The *CONFIG5.K* bit becomes read-write and is programmed by hardware to a value of 0 to indicate the legacy configuration. In this case, the cache coherency attributes for the kseg0 segment are derived from the *Config.K0* field as described in the previous subsection. In addition to selecting the location of the cache coherency attributes, the *CONFIG5.K* bit also causes hardware to generate two boot exception overlay segments, one for kseg0 and one for kseg1, as described in [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#).
- Hardware programs the CP0 memory segmentation registers (*SegCtl0 - SegCtl2*) for the legacy setting. An example of this programming is shown in [Table 3.11](#). Note that these registers are new in the interAptiv core and are not used by legacy software. However, they are used by hardware during normal operation, so their default values should not be changed.


If the *SI\_EVAREset* pin is asserted at reset, the interAptiv core comes up in the EVA configuration (default is *xkseg0* space = 3 GB) and hardware takes the following actions:

- The *CONFIG5.K* bit becomes read-only and is forced to a value of 1 to indicate the EVA configuration. In this case, the *CONFIG.K0* field is ignored and is no longer used to determine the kseg0 cache coherency attributes (CCA). Rather, the values in bits 2:0 (segments 0, 2, and 4) and bits 18:16 (segments 1, 3, and 5) of the *SegCtl0 - SegCtl2* registers are used to define the CCA for each memory segment as shown in [Table 3.3](#). In this case, hardware generates only one BEV overlay segment as described in [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#).
- Hardware sets the CP0 memory segmentation registers (*SegCtl0 - SegCtl2*) for the EVA configuration. An example of this programming is shown in [Table 3.12](#).

These two options are illustrated in [Figure 3.12](#). Refer to [Section 2.2.1.6, "Device Configuration 5 — Config5 \(CP0 Register 16, Select 5\)"](#) in the CP0 register chapter for more information on the *CONFIG5.K* bit.



**Figure 3.12 Relationship Between SI\_EVAREset and CONFIG5.K at Reset**



### 3.5.3.4 Enhanced Virtual Address Detection and Support

As described above, the *SegCtl0* - *SegCtl2* registers are used to control the various memory segments. In addition to these registers, two other configuration registers are also used in enhanced virtual addressing (EVA).

The *EVA* bit in the *Config5* register (*Config5<sub>EVA</sub>*) is used to detect support for the enhanced virtual address scheme. This read-only bit is always 1 to indicate support for EVA.

In addition to the *EVA* bit, the *SC* bit in the *Config3* register (*Config3<sub>SC</sub>*) is used by hardware to detect the presence of the *SegCtl0* - *SegCtl2* registers. This read-only bit is always 1 in the interAptiv core to indicate the presence of these registers. Note that both of these features must be present to configure the virtual address space for EVA.

### 3.5.3.5 Setting the Access Control Mode

In addition to setting the *Config5<sub>EVA</sub>* and *Config3<sub>SC</sub>* bits described above, each memory segment must be set to the programmable segmentation mode. Bits 6:4 (segments 0, 2, and 4) and bits 22:20 (segments 1, 3, and 5) of the *SegCtl0* through *SegCtl2* registers define the access control mode.

To set the programmable segmentation registers to mimic the traditional MIPS32 virtual address mapping shown in [Figure 3.10](#), the *AM* and *C* subfields (defined in [Table 3.3](#)) of each 16-bit *CFG* field of the *SegCtl0* - *SegCtl2* registers should be programmed as shown in [Table 3.4](#).

**Table 3.4 Setting the Access Control Mode for the Legacy Configuration**

<i>SegCtl</i> Register	CFGn	CFGn Subfields		Segment Size	Location in Virtual Memory Map	Description
		AM	C			
0	0 (bits 15:0)	MK (bits 6:4 = 0x1)	0x3 (bits 2:0)	0.5GB	3.5 - 4.0 GB	Mapped kernel region.
0	1 (bits 31:16)	MSK (bits 22:20 = 0x2)	0x3 (bits 18:16)	0.5GB	3.0 - 3.5 GB	Mapped kernel, supervisor region.
1	2 (bits 15:0)	UK (bits 6:4 = 0x0)	0x2 (bits 2:0)	0.5GB	2.5 - 3.0 GB	Kernel unmapped, uncached region.

**Table 3.4 Setting the Access Control Mode for the Legacy Configuration (*continued*)**

SegCtl Register	CFGn	CFGn Subfields		Segment Size	Location in Virtual Memory Map	Description
		AM	C			
1	3 (bits 31:16)	UK (bits 22:20 = 0x0)	0x3 (bits 18:16)	0.5GB	2.0 - 2.5 GB	Kernel unmapped, cached region.
2	4 (bits 15:0)	MUSK (bits 6:4 = 0x3)	0x3 (bits 2:0)	1.0GB	1.0 - 2.0 GB	User, supervisor, and kernel mapped region.
2	5 (bits 31:16)	MUSK (bits 22:20 = 0x3)	0x3 (bits 18:16)	1.0GB	0.0 - 1.0 GB	User, supervisor, and kernel mapped region.

To set the programmable segmentation registers to implement EVA with a 3.0 GB *xkseg0* space as shown in [Figure 3.11](#), the *AM* and *C* subfields (defined in [Table 3.3](#)) of each *CFG* field of the *SegCtl0* - *SegCtl2* registers should be programmed as shown in [Table 3.5](#).

**Table 3.5 Setting the Access Control Mode for the EVA Configuration**

SegCtl Register	CFGn	CFGn Subfields		Segment Size	Location in Virtual Memory Map	Description
		AM	C			
0	0 (bits 15:0)	MK (bits 6:4 = 0x1)	0x3 (bits 2:0)	0.5GB	3.5 - 4.0 GB	Mapped kernel region.
0	1 (bits 31:16)	MK <sup>1</sup> (bits 22:20 = 0x1)	0x3 (bits 18:16)	0.5GB	3.0 - 3.5 GB	Mapped kernel region.
1	2 (bits 15:0)	MUSUK (bits 6:4 = 0x4)	0x3 (bits 2:0)	0.5GB	2.5 - 3.0 GB	Mapped user/supervisor, unmapped kernel region.
1	3 (bits 31:16)	MUSUK (bits 22:20 = 0x4)	0x3 (bits 18:16)	0.5GB	2.0 - 2.5 GB	Mapped user/supervisor, unmapped kernel region.
2	4 (bits 15:0)	MUSUK (bits 6:4 = 0x4)	0x3 (bits 2:0)	1.0GB	1.0 - 2.0 GB	Mapped user/supervisor, unmapped kernel region.
2	5 (bits 31:16)	MUSUK (bits 22:20 = 0x4)	0x3 (bits 18:16)	1.0GB	0.0 - 1.0 GB	Mapped user/supervisor, unmapped kernel region.

1. This segment can also be mapped to MSK (bits 22:20 = 0x2) if supervisor mode is supported.

MUSUK is an acronym for *Mapped User/Supervisor, Unmapped Kernel*. This mode sets the kernel unmapped virtual address space to *xkseg0* as shown in [Figure 3.11](#).

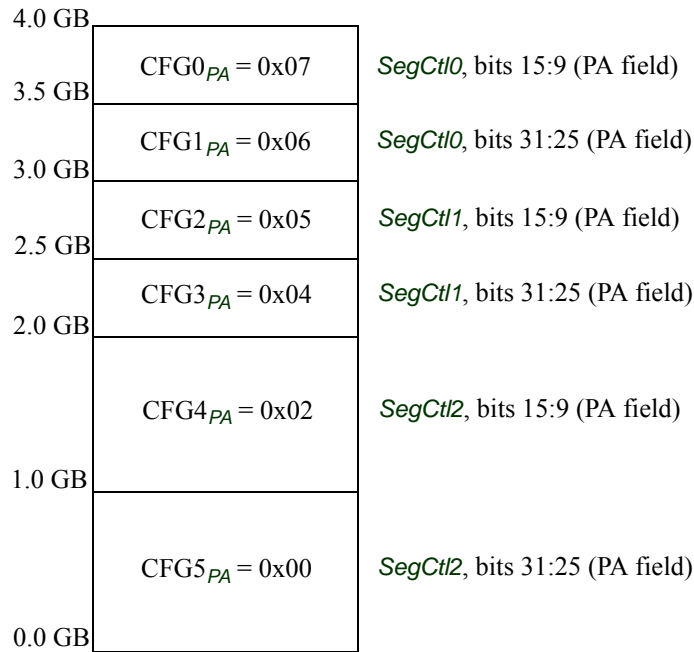
### 3.5.3.6 Defining the Physical Address Range for Each Memory Segment

As shown in [Table 3.2](#), each of the six 16-bit CFGn fields of the *SegCtl0* through *SegCtl2* fields controls a specific portion of the physical address range. Bits 11:9 (segments 0, 2, and 4) and bits 27:25 (segments 1, 3, and 5) of the *SegCtl0* through *SegCtl2* registers represent the state of physical address bits 31:29 and defines the starting address of each segment. These bits control the six segments of the physical address.

Note that bits 31:28 and bits 15:12 are also part of the physical address field, but they are not used in the interAptiv core and are reserved for future use by devices that implement a 36-bit address.

Figure 3.13 below shows an example of how each segment of the physical address can be mapped to the *SegCtl0* through *SegCtl2* registers.

**Figure 3.13 Mapping of SegCtl 0 - 2 Registers to Physical Address Space**



For example, to program the *xkseg0* region to a size of 3.0GB, the PA field of each register would be programmed as follows:

**Table 3.6 Example of a 3.0GB Kernel Unmapped Segment**

Register	CFGn Field	Bits	PA Field	Memory Segment	Virtual Address Range
SegCtl0 <sup>1</sup>	CFG0	15:9	0x07	kseg2	0xE000_0000 - 0xFFFF_FFFF
	CFG1	31:25	0x06		0xC000_0000 - 0xDFFF_FFFF
SegCtl1	CFG2	15:9	0x05		0xA000_0000 - 0xBFFF_FFFF
	CFG3	31:25	0x04		0x8000_0000 - 0x9FFF_FFFF
SegCtl2	CFG4	15:9	0x02	xkseg0	0x4000_0000 - 0x7FFF_FFFF
	CFG5	31:25	0x00		0x0000_0000 - 0x3FFF_FFFF

1. In the 3GB *xkseg0* example, the PA portion of the CFG0 and CFG1 fields are not used because they are associated with kernel mapped address spaces. In this case the PA fields are not required since the physical address is determined by the TLB. In the maximum configuration, *xkseg0* can be extended to 3.5GB. In this case, the CFG1 field of the *SegCtl0* register would become part of the *xkseg0* segment and the PA subfield would be used.

### 3.5.3.7 Enhanced Virtual Address (EVA) Instructions

By default, an implementation that supports EVA requires a number of new load/store instructions that are used when the enhanced virtual address scheme is enabled. These kernel-mode user load/store instructions allow the kernel mapped access to user address space as if it were in user mode.

For example, the kernel can copy data from user address space to kernel physical address space by using such instructions with user virtual addresses. Kernel system-calls from user space can be conveniently changed by replacing normal load/store instructions with these instructions. Switching modes (kernel to user) is an alternative but this is an issue if the same virtual address is being simultaneously used by the kernel. Further, there is a performance penalty in context-switching.

The opcode for these instructions is embedded into bits 2:0 of the instruction, known as the *Type* field. Note that some fields can have the same encoding depending whether the operation is a load or a store. The load/store designation is determined by the *AIU L/S* field, or bits 5:3 of the instruction. [Table 3.7](#) lists the new kernel load/store instructions.

For a complete list of new instructions, refer to [Section 23.7, "New Instructions for the interAptiv™ Core"](#).

**Table 3.7 Load/Store Instructions in Programmable Memory Segmentation Mode**

Instruction Mnemonic	Instruction Name	Description
LBE	Load Byte Kernel	Load byte (as if user from) kernel extended virtual addressing load from user virtual memory while operating in kernel mode.
LBUE	Load Byte Unsigned Kernel	Load byte unsigned (as if user from) kernel.
LHE	Load Halfword Kernel	Load halfword (as if user from) kernel.
LHUE	Load Halfword Unsigned Kernel	Load halfword unsigned (as if user from) kernel.
LWE	Load Word Kernel	Load word (as if user from) kernel.
SBE	Store Byte Kernel	Store byte (as if user from) kernel extended virtual addressing load from user virtual memory while operating in kernel mode.
SHE	Store Halfword Kernel	Store halfword (as if user from) kernel.
SWE	Store Word Kernel	Store word (as if user from) kernel.

## 3.6 Boot Exception Vector Relocation in Kernel Mode

Historically in MIPS processors, the boot exception vector (BEV) has always been at the same location in both virtual and physical memory, being mapped from a virtual address of 0xBFC0\_0000 to a physical address of 0x1FC0\_0000. With the advent of memory segmentation in the interAptiv core, the BEV vector may not always map to a physical address of 0x1FC0\_0000. This can cause a scenario where the boot exception vector resides at two different physical addresses depending on the memory mode. To address this issue, the interAptiv core implements a boot exception vector overlay scheme that allows the BEV to be mapped to a single location in physical memory, regardless of the memory mode.

This section describes how to define the BEV overlay segment and the BEV relocation process for both the legacy setting and the Enhanced Virtual Address (EVA) setting, which is one element of the interAptiv memory segmentation scheme.

Note that boot exception vector relocation is performed only in Kernel mode. For more information on placing the interAptiv core in kernel mode, refer to [Section 3.13.4, "Kernel Mode"](#).

### 3.6.1 Boot Configurations

In kernel mode, the interAptiv core can be powered up in one of two address settings:

- Legacy setting
- Enhanced Virtual Address (EVA) setting

#### ***Legacy Setting***

The legacy setting is the traditional boot mode followed by all MIPS processor prior to interAptiv, where the boot exception vector (BEV) is located at 0xBFC0\_0000 in virtual address space, and maps to 0x1FC0\_0000 in physical address space. An example of legacy mode is described in [Section 3.6.3, "Mapping of the Boot Exception Vector in the Legacy Configuration"](#).

#### ***EVA Setting***

In the EVA setting, the boot exception vector can be located anywhere in virtual address space and mapped to anywhere in physical address space. An example of an EVA configuration is described in [Section 3.6.4, "Example Mapping of the Boot Exception Vector in the EVA Configuration"](#).

For more information on configuring the interAptiv core in the Legacy and EVA settings, refer to [Section 3.6.5.1, "Setting the Type of Memory Addressing Mode"](#).

### 3.6.2 Pins Used to Support Boot Exception Vector Relocation

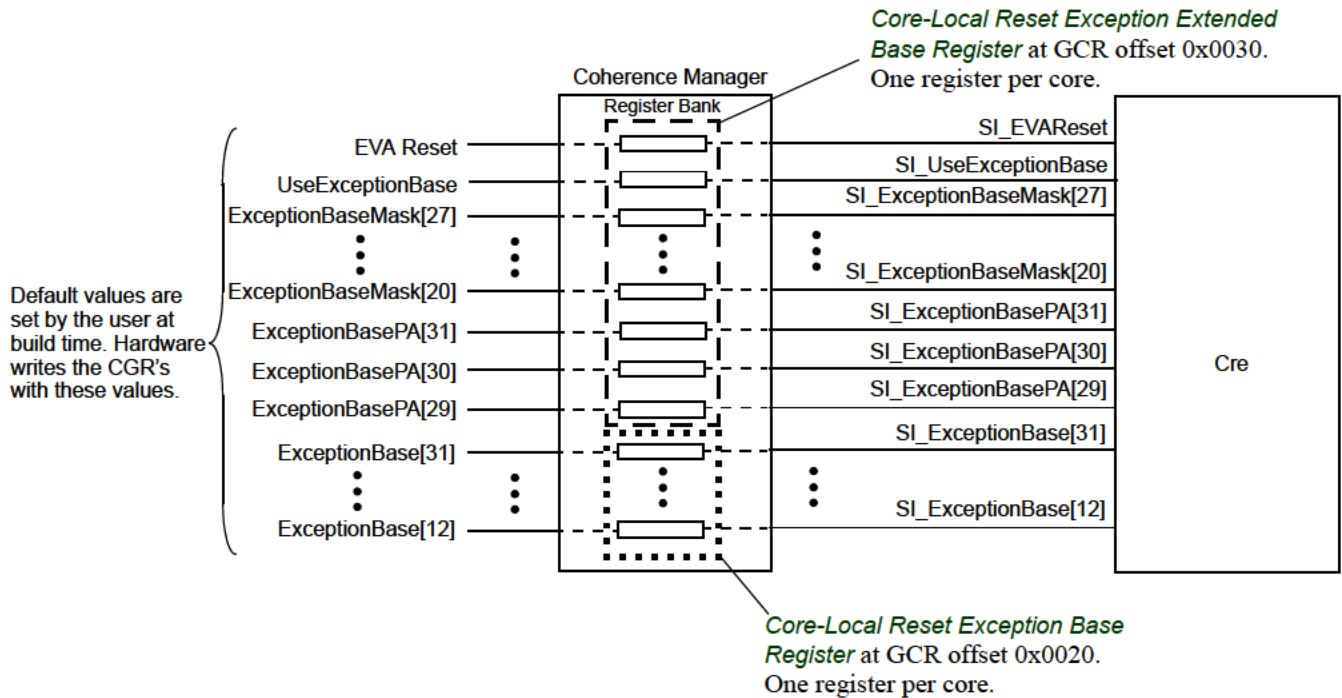
To facilitate the BEV overlay scheme, a number of pins were added to the interAptiv core that allow the user to select the boot overlay parameters at build time. The initial state of the default values selected by the user at build time are registered inside the Coherence Manager (CM2) block using two *Global Configuration Registers* (GCR)

As shown in [Figure 3.14](#), there are two GCR registers used per core. Each core has its own pair of GCR registers and its own set of BEV related pins. This allows each core to be programmed in a different manner and independently from one another.

The CM2 drives these values to the interAptiv cores at reset. Note that the two CGR registers are loaded only on a cold boot and are programmed with the values selected by the user at build time. Each of these pins is described in the following subsections.

Figure 3.14 shows the boot exception vector pins for a single interAptiv core. Each additional core would have an identical set of CM2 registers and set of BEV related pins shown in the figure.

**Figure 3.14 Registered Boot Exception Vector Relocation Pins — One Core**



As noted in the figure above, there is one pair of GCR registers for each core. This allows each interAptiv core to be powered up in a different memory mode and independently from one another.

The boot exception vector relocation pins are described in [Table 3.8](#).

**Table 3.8 interAptiv Core Boot Exception Vector Pins**

Pin Name	Field Size in Bits	CM2 GCR Register Mapping	Description
SI_EVAReset	1	Bit 31 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	<p>If this pin is asserted at reset, the interAptiv core comes up in the EVA configuration. In this case the <i>CONFIG5.K</i> bit becomes read-only with a fixed value of 1 to indicate EVA as the addressing scheme. In addition, the <i>SegCtl0 - SegCtl2</i> registers are configured with values that correspond to the EVA mapping.</p> <p>If this pin is not asserted at reset, the interAptiv core comes up in the legacy setting. In this case the <i>CONFIG5.K</i> bit becomes read-write with an initial value of 0 to indicate legacy mode. This bit is modified by software when switching from legacy mode to EVA mode as described in <a href="#">Section 3.6.6, "Switching the Addressing Scheme from Legacy to EVA After Boot-up"</a>.</p> <p>This pin is used in both the legacy and EVA settings. There is one <i>SI_EVAReset</i> pin per core.</p>
SI_UseExceptionBase	1	Bit 30 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	<p>In the legacy configuration, if the <i>SI_UseExceptionBase</i> pin is not asserted, then the BEV location defaults to 0xBFC0_0000.</p> <p>If the <i>SI_UseExceptionBase</i> pin is asserted, address bits <i>SI_ExceptionBase[31:30]</i> are forced to a value of 2'b10 to force the BEV location into the KSEG0/KSEG1 space.</p> <p>This pin is only used in the legacy configuration. There is one <i>SI_UseExceptionBase</i> pin per core.</p>
SI_ExceptionBaseMask[27:20]	8	Bits 27:20 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	<p>Used to determine the size of the boot exception vector overlay region from 1 MB to 256 MB in powers of two. These pins are used in both the legacy and EVA configurations. There is one set of <i>SI_ExceptionBaseMask</i> pins per core.</p>
SI_ExceptionBasePA[31:29]	3	Bits 3:1 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	<p>Upper physical address bits. The size of the overlay region defined by <i>SI_ExceptionBaseMask[27:20]</i> is remapped to a location in physical address space pointed to by the <i>SI_ExceptionBasePA[31:29]</i> pins. This allows the overlay region to be placed into one of the 512 MB segments in physical memory. These pins are used in both the legacy and EVA configurations. There is one set of <i>SI_ExceptionBasePA</i> pins per core.</p>

**Table 3.8 interAptiv Core Boot Exception Vector Pins (continued)**

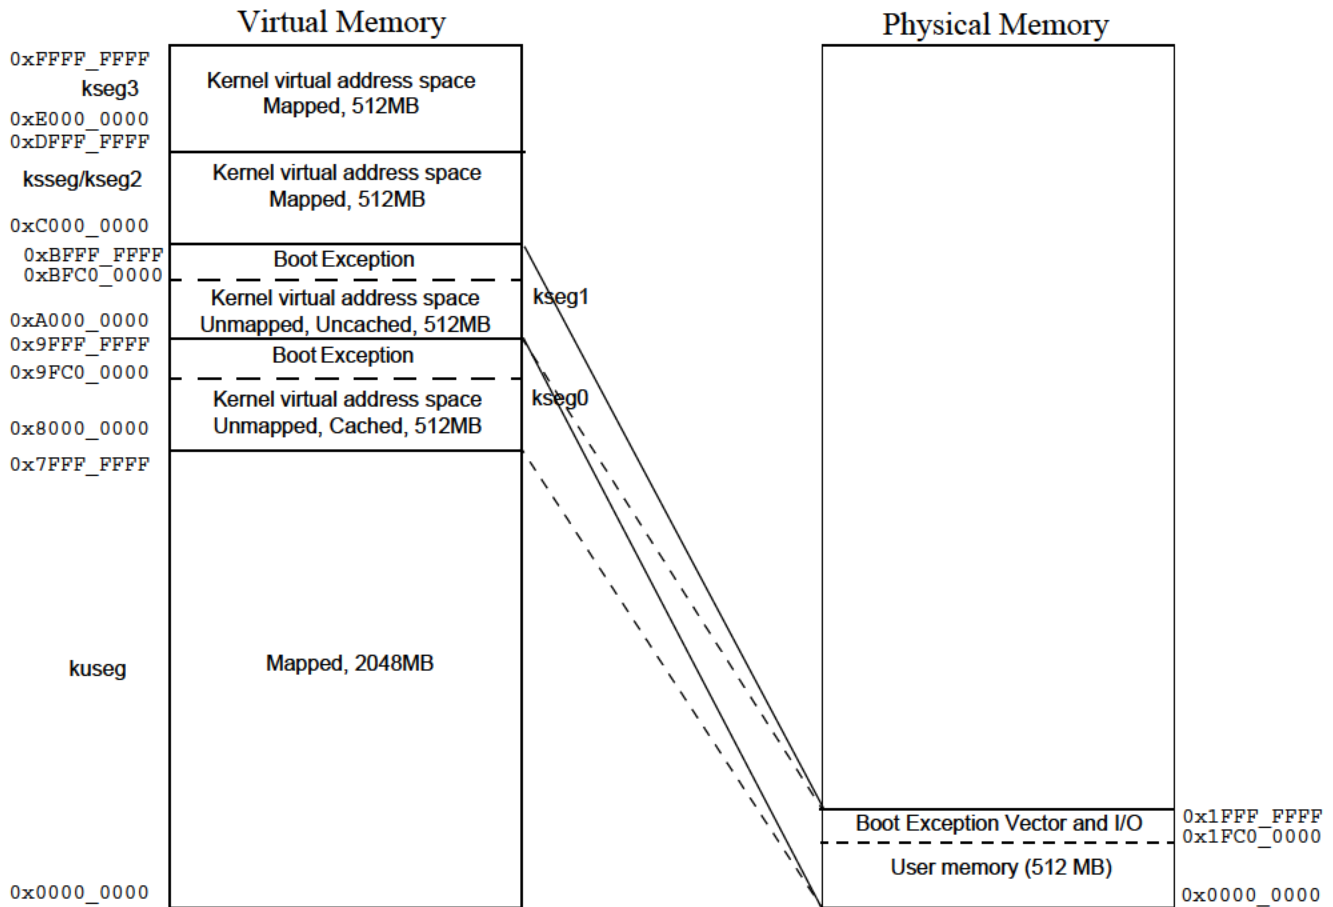
Pin Name	Field Size in Bits	CM2 GCR Register Mapping	Description
SI_ExceptionBase[31:12]	20	Bits 31:12 of the <i>Core-Local Reset Exception Base Register</i> (offset = 0x0020)	<p>The <i>SI_ExceptionBase[31:12]</i> pins define the boot address in virtual address space which is used to define the overlay region. These pins, along with the <i>SI_ExceptionBaseMask[27:20]</i> pins, determine the size and location of the BEV region within virtual address space.</p> <p>Note that the <i>CONFIG5_K</i> CP0 register bit is used to determine which pins of the <i>SI_ExceptionBase[31:12]</i> address are used to calculate the overlay as described in <a href="#">Section 3.6.5.1</a>.</p> <p>These pins are used in the EVA setting and can also be used in the legacy setting. There is one set of <i>SI_ExceptionBase</i> pins per core.</p>

### 3.6.3 Mapping of the Boot Exception Vector in the Legacy Configuration

In all MIPS processors prior to the Aptiv family, the boot exception vector (BEV) was located at a virtual address of 0xBFC0\_0000, and a corresponding physical address of 0x1FC0\_0000. In addition, since both the Kernel Segment 1 (KSEG1) and Kernel Segment 0 (KSEG0) virtual memory spaces mapped to the same physical address space, the contents of the BEV were duplicated at a virtual address of 0x9FC0\_0000. This concept is shown in [Figure 3.15](#).



**Figure 3.15 Mapping of the Boot Exception Vector in the Legacy Configuration**



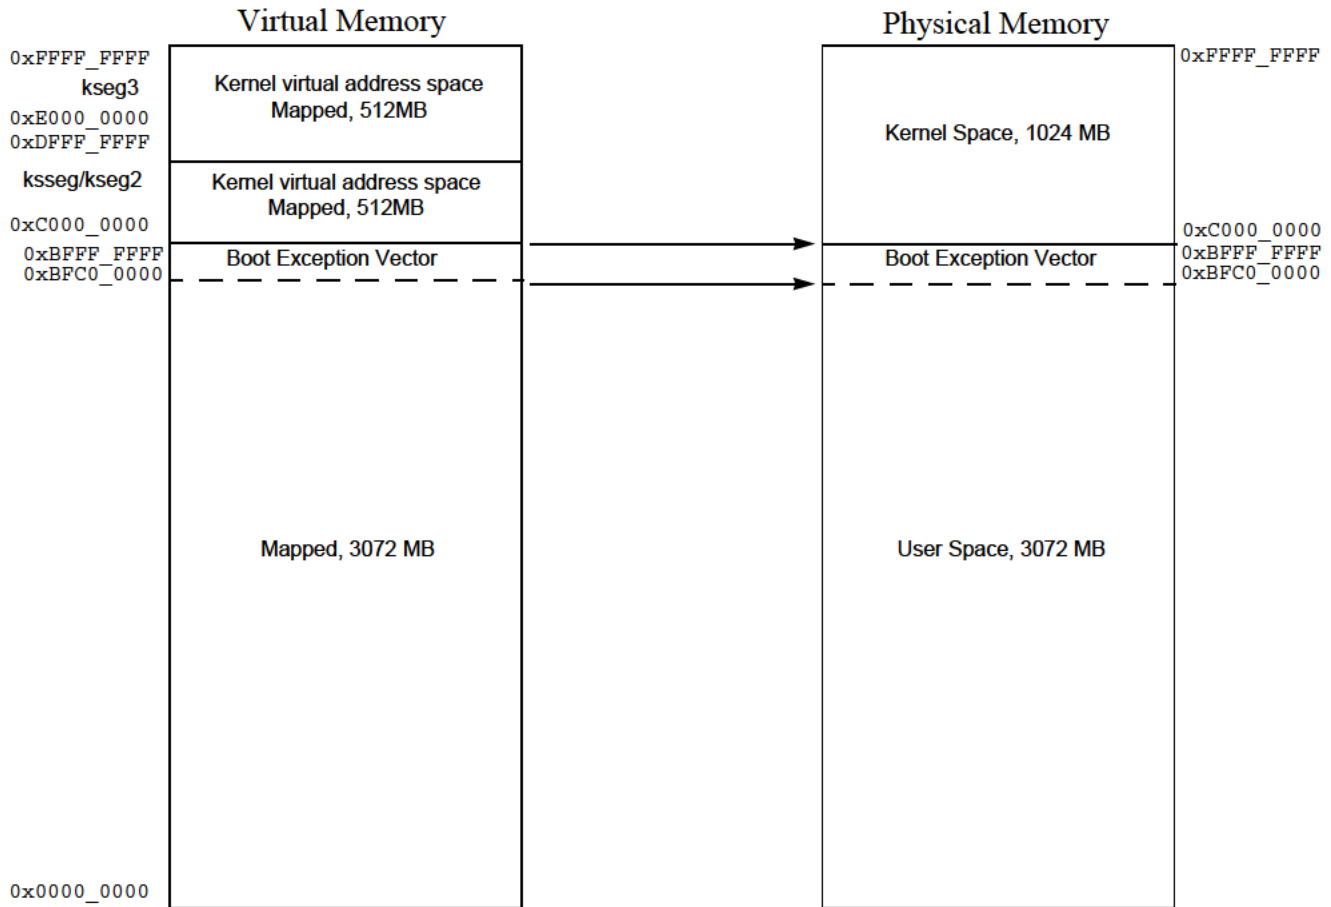
As shown in Figure 3.15 above, the default BEV location is at the top of the first 512 MB physical address space. This space typically includes not only the boot exception vector, but also the I/O memory. It is also used to store the debug exception and non-maskable interrupt (NMI) handlers. In this scheme, the top portion of memory is not available for general storage because of the existence of the boot ROM and I/O devices.

The BEV overlay scheme implemented in the interAptiv core not only ensures that the boot exception vector will be mapped to the same location regardless of the memory mode, but it also eliminates the issue of non-contiguous user memory because the BEV no longer need be in a fixed location of 0x1FC0\_0000.

### 3.6.4 Example Mapping of the Boot Exception Vector in the EVA Configuration

In the interAptiv core, physical memory sizes can be up to 3.5 GB. As described above, in the legacy configuration, the BEV is remapped from 0xBFC0\_0000 in virtual memory to 0x1FC0\_0000 in physical memory. However, in the EVA configuration, if the physical memory size is set to 3.0 GB, no remapping of the BEV is required. In this case, the BEV is remapped from 0xBFC0\_0000 in virtual memory to 0xBFC0\_0000 in physical memory. This relocation of the BEV between the two memory modes creates a conflict if software switches from the legacy configuration to the EVA configuration because the BEV will be mapped to two different locations in physical memory. The remapping of the BEV in the legacy configuration is shown in Figure 3.15 above. An example of mapping of the BEV in the EVA configuration with a 3.0 GB *xkseg0* space is shown in Figure 3.16.

**Figure 3.16 Example of Mapping the Boot Exception Vector in the EVA Configuration**



### 3.6.5 Defining the Boot Exception Vector Overlay Region

To solve the problem of having the boot exception vector residing at different physical address locations based on the memory mode, the interAptiv core defines a boot exception vector overlay region that can be programmed from 1 MB to 256 MB in powers of two using an 8-bit virtual address mask as described below.

- 1 MB, 2 MB, 4 MB, 8 MB, 16 MB, 32 MB, 64 MB, 128 MB, or 256 MB

This space is then mapped to a predetermined location in physical memory, regardless of the memory configuration (legacy or EVA).

The BEV overlay not only allows the location of the boot exception vector to be mapped to an area common to both the Legacy and EVA configurations, but also eliminates the non-contiguous chunk of memory that was created by having the boot exception vector at the top of the first 512M of physical memory space as in previous generation processors.

To set the boot exception overlay region, the following steps are taken. Each of these steps is described in the following subsections:

1. Determine whether the interAptiv core boots up in Legacy mode or EVA mode. This function is described in [Section 3.6.5.1, "Setting the Type of Memory Addressing Mode"](#).

2. Determine which virtual address bits will be used to calculate the boot exception vector base address. This function is described in [Section 3.6.5.2, "Using the SI\\_UseExceptionBase Pin and CONFIG5.K to Determine How to Calculate the BEV Base Address"](#).
3. Determine the size and location of the overlay region in virtual address space. This function is described in [Section 3.6.5.3, "Determining the Size and Location of the Overlay Region in Virtual Address Space"](#).
4. Determine the location of the overlay region in physical address space. This function is described in [Section 3.6.5.4, "Determining the Location of the Overlay Region in Physical Memory"](#).

### 3.6.5.1 Setting the Type of Memory Addressing Mode

The *SI\_EVAReset* pin, along with the *CONFIG5.K* bit, determines whether the addressing scheme is set to legacy or EVA at reset.

Refer to [Section 3.5.3.3, "Setting the Memory Addressing Scheme — SI\\_EVAReset and CONFIG5.K"](#) for more information.

### 3.6.5.2 Using the SI\_UseExceptionBase Pin and CONFIG5.K to Determine How to Calculate the BEV Base Address

The *SI\_UseExceptionBase* pin and the *CONFIG5.K* register bit are also used to determine the addressing scheme and how the location of the boot exception vector will be calculated. The relationship between the *SI\_UseExceptionBase* pin and the *CONFIG5.K* register is shown in [Table 3.9](#). This table shows how to use the various address fields (*SI\_ExceptionBaseMask[27:20]* and *SI\_ExceptionBase[31:12]*) described in [Section 3.6.5.3, "Determining the Size and Location of the Overlay Region in Virtual Address Space"](#).

**Table 3.9 *SI\_UseExceptionBase* Pin and CONFIG5.K Encoding**

CONFIG5.K Bit	<i>SI_UseExceptionBase</i> Pin	Condition	Action
0	0	Legacy Configuration <i>SI_ExceptionBase[31:12]</i> pins are not used.	Use default BEV location of 0xBFC0_0000.
0	1	Legacy Configuration Use only <i>SI_ExceptionBase[29:12]</i> for the BEV base location. Bits 31:30 are forced to a value of 2'b10 to put the BEV vector into KSEG0/KSEG1 virtual address space.	The BEV location is determined as follows: <i>SI_ExceptionBase[31:12]</i> = 2'b10, <i>SI_ExceptionBase[29:12]</i> pins, 12'b0 Bits 31:30 are forced to a value of 2'b10 to put the BEV vector into KSEG0/KSEG1 virtual address space.
1	Don't care	EVA Configuration Use <i>SI_ExceptionBase[31:12]</i> pins.	The <i>SI_ExceptionBase[31:12]</i> pins are used directly to derive the BEV location. The <i>SI_UseExceptionBase</i> pin is ignored.

### 3.6.5.3 Determining the Size and Location of the Overlay Region in Virtual Address Space

The starting location of the overlay region in virtual address space is defined using either the *SI\_ExceptionBase[31:12]* pins, or the *SI\_ExceptionBase[29:12]* pins depending on the state of the

*SI\_UseExceptionBase* pin and *CONFIG5\_K* bit as described in Table 3.9 above. The size of the overlay region where the BEV is located is determined using the *SI\_ExceptionBaseMask[27:20]* pins as shown in Table 3.10.

**Table 3.10 Encoding of *SI\_ExceptionBaseMask[27:20]***

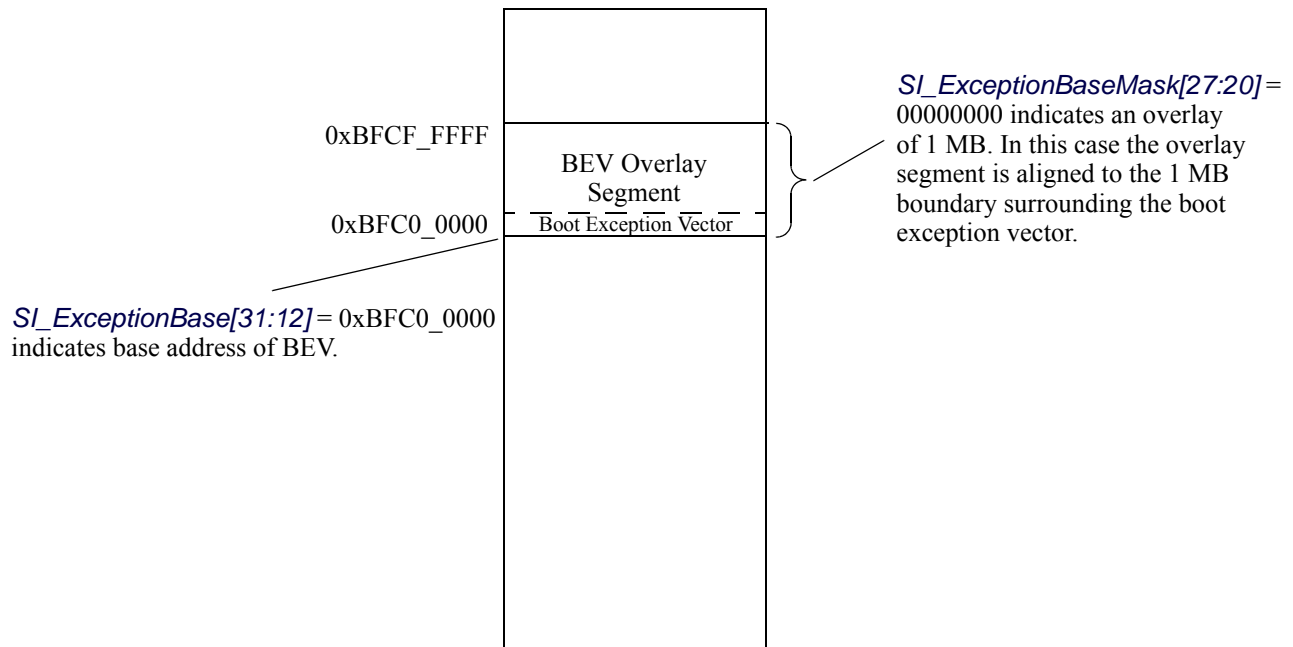
<i>SI_ExceptionBaseMask[27:20]</i>	Segment Size
00000000	1 MB
00000001	2 MB
00000011	4 MB
00000111	8 MB
00001111	16 MB
00011111	32 MB
00111111	64 MB
01111111	128 MB
11111111	256 MB

Consider the following example:

- The location of the BEV is at 0xBFC0\_0000
- The overlay size is 1 MB (*SI\_ExceptionBaseMask[27:20]* = 00000000)
- The *CONFIG5\_K* CP0 register bit is set

In this case the BEV segment would be located in virtual address space as shown in Figure 3.17.

**Figure 3.17 Size and Location of Overlay Region in Virtual Address Space — 1 MB Example**



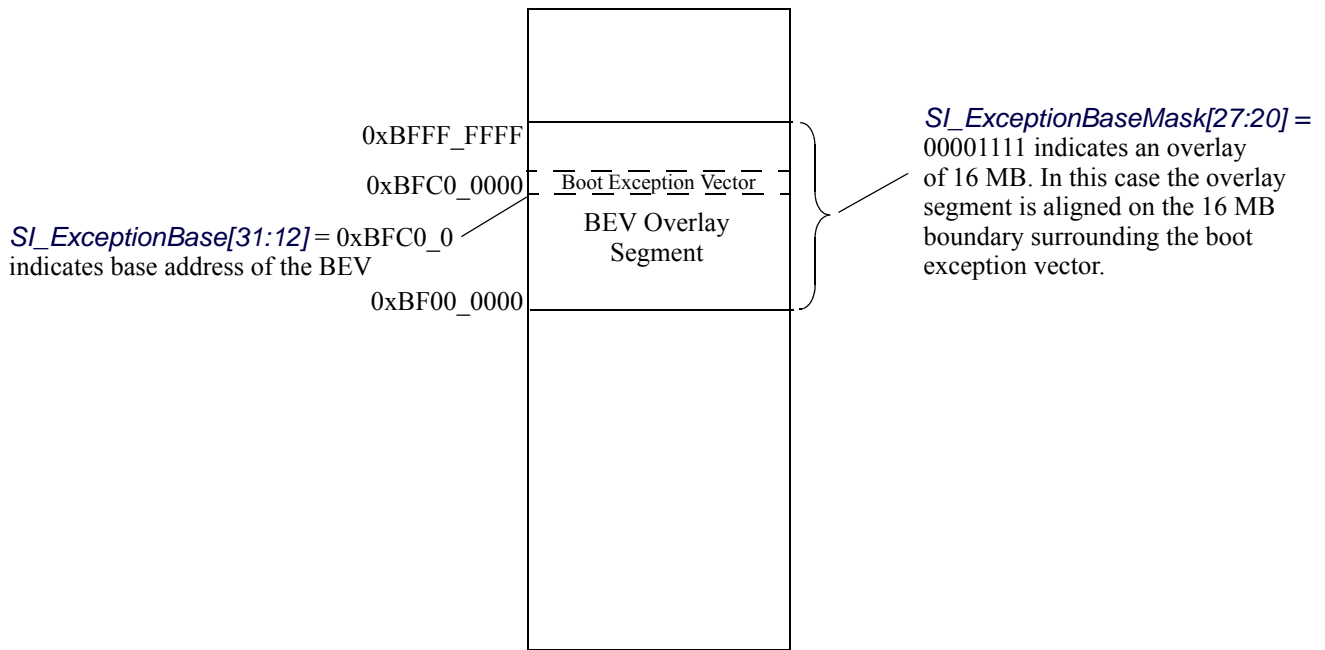
In the above example, the start of the BEV is aligned on a 1 MB boundary and therefore is at the start of the 1MB address space. This may not always be the case depending on the size of the overlay region as shown in Figure 3.18 below.

In another example:

- The location of the BEV is at 0xBFC0\_0000
- The overlay size is 16 MB ( $SI\_ExceptionBaseMask[27:20] = 00001111$ )
- The  $CONFIG5_K$  CP0 register bit is set

In this case the BEV segment would be located in virtual address space as shown in [Figure 3.18](#).

**Figure 3.18 Size and Location of Overlay Region in Virtual Address Space — 16 MB Example**



#### 3.6.5.4 Determining the Location of the Overlay Region in Physical Memory

As described in the previous subsections, the  $SI\_ExceptionBase[31:12]$  and  $SI\_ExceptionBaseMask[27:20]$  fields are used to determine the size and location of the overlay within virtual address space. This segment of virtual memory is then remapped to physical memory at a location determined by the  $SI\_ExceptionBasePA[31:29]$  pins. These pins divide the physical address space into a number of 512 MByte segments. For example, in a 4 GB physical address space, the space can be divided into eight 512 MByte segments. This concept is shown in [Figure 3.19](#).

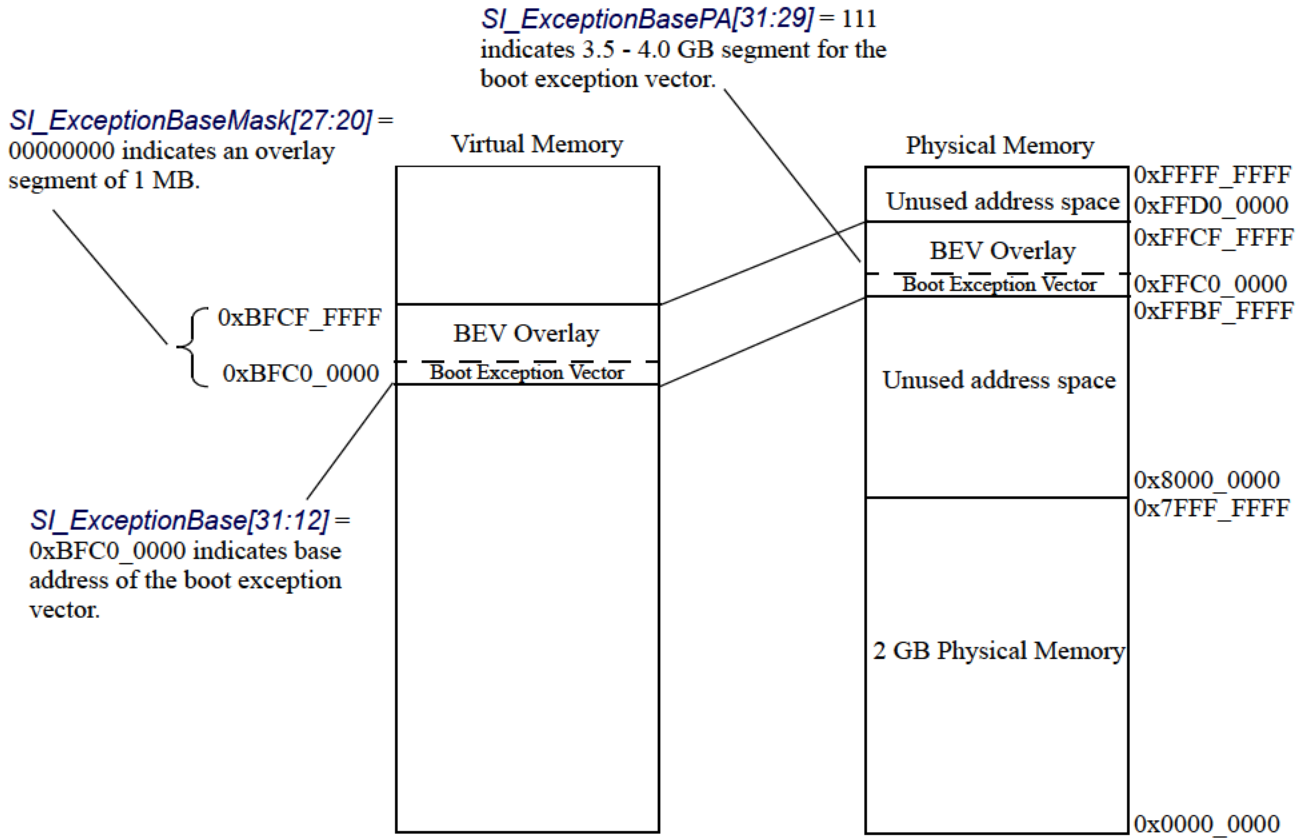
**Figure 3.19 Physical Address Space Segmentation Using *SI\_ExceptionBasePA[31:29]***

	Physical Address
<i>SI_ExceptionBasePA[31:29]</i> = 111	3.5 GB - 4.0 GB
<i>SI_ExceptionBasePA[31:29]</i> = 110	3.0 GB - 3.5 GB
<i>SI_ExceptionBasePA[31:29]</i> = 101	2.5 GB - 3.0 GB
<i>SI_ExceptionBasePA[31:29]</i> = 100	2.0 GB - 2.5 GB
<i>SI_ExceptionBasePA[31:29]</i> = 011	1.5 GB - 2.0 GB
<i>SI_ExceptionBasePA[31:29]</i> = 010	1.0 GB - 1.5 GB
<i>SI_ExceptionBasePA[31:29]</i> = 001	0.5 GB - 1.0 GB
<i>SI_ExceptionBasePA[31:29]</i> = 000	0 - 0.5 GB

For example, assume that the boot exception vector resides at a virtual address of 0xBFC0\_0000, and the size of the segment is 1 MB as determined by the *SI\_ExceptionBaseMask[27:20]* pins. The physical memory size (amount of DRAM) is 2 GB, and the boot ROM that contains the BEV has been relocated to the top 512 MB of the 4 GB physical address space using the *SI\_ExceptionBasePA[31:29]* pins, which selects the segment from 3.5 GB to 4.0 GB. The remapping of the boot exception vector would be as shown in [Figure 3.20](#).

In this example, because the overlay region has been defined, the boot exception vector would be relocated to the same address space, regardless of whether the addressing scheme is legacy or EVA. In addition, the memory space that contains the BEV no longer need be shared with actual physical memory in the first 512 MB of memory space as with previous MIPS processors, thereby allowing for all of the memory to be contiguous and available to the user.

**Figure 3.20 Example of Relocating the Boot Exception Vector**



### 3.6.6 Switching the Addressing Scheme from Legacy to EVA After Boot-up

This section discusses a scenario where the processor is booted-up in the Legacy configuration, and then switches to the EVA configuration. To boot the interAptiv core in the Legacy configuration, the user selects the legacy boot configuration at build time. This causes the CM2 hardware to drive a logic '0' onto the  $SI\_EVAReset$  pin during boot time.

As shown in Figure 3.14, the default values selected by the user during build time are written into the GCR registers of the Coherence Manager (CM2). The CM2 then forwards these values to each interAptiv core in the system. In this case, the  $SI\_EVAReset$  pin is a logic '0' at boot time. This places the cores into the legacy configuration and writes a '0' value into the  $CONFIG5.K$  bit in CP0. The registering of the BEV signals allows software to change the addressing scheme from legacy to EVA when necessary.

Along with clearing the  $CONFIG5.K$  bit in CP0 and making the bit read/write based on the state of  $SI\_EVAReset$  at boot time, hardware also writes the  $SegCtl0 - SegCtl2$  registers with values that mimic the legacy memory map shown in Figure 3.15. In this example, the  $SegCtl0 - SegCtl2$  registers would be programmed with the following values shown in Table 3.11.

In the following tables, note that the PA field includes physical address bits 35:32. These bits are reserved for future expansion purposes and are not available in the interAptiv core. These upper four bits of each PA field are always 0 in the interAptiv core.

**Table 3.11 SegCtl0 - SegCtl2 Register Settings in the Legacy Configuration**

Register	Bits	Segment Size	Name	Definition	Reset value
SegCtl0 (CFG0)	[2:0]	0.5 GB (3.5 - 4.0 GB)	CCA	CFG0 Cache Coherency Attributes	Not defined since this is a mapped kernel region.
SegCtl0 (CFG0)	[3]		EU	CFG0 Error	1'b0: CP0 <i>Status.ERL</i> ignored
SegCtl0 (CFG0)	[6:4]		AM	CFG0 Region Type	3'b001: Mapped only kernel region. This is kseg3.
SegCtl0 (CFG0)	[15:9]		PA	CFG 0 Physical Address Bits [35:29]	Not defined since it is a mapped region
SegCtl0 (CFG1)	[18:16]	0.5 GB (3.0 - 3.5 GB)	CCA	CFG1 Cache Coherency Attributes	Not defined since this is a mapped kernel region.
SegCtl0 (CFG1)	[19]		EU	CFG1 Error	1'b0: CP0 <i>Status.ERL</i> ignored
SegCtl0 (CFG1)	[22:20]		AM	CFG1 Region Type	3'b010: Mapped Kernel/Supervisor region. This is kseg2/ksseg.
SegCtl0 (CFG1)	[31:25]		PA	CFG1 Physical Address Bits [35:29]	Not defined since it is a mapped region
SegCtl1 (CFG2)	[2:0]	0.5 GB (2.5 - 3.0 GB)	CCA	CFG2 Cache Coherency Attributes	0x2: Uncached
SegCtl1 (CFG2)	[3]		EU	CFG2 Error	1'b0: CP0 <i>Status.ERL</i> ignored.
SegCtl1 (CFG2)	[6:4]		AM	CFG2 Region Type	3'b000: Kernel unmapped region. This is kseg1.
SegCtl1 (CFG2)	[15:9]		PA	CFG2 Physical Address Bits [35:29]	0x0: Points to 0.0 - 0.5 GB physical address region.
SegCtl1 (CFG3)	[18:16]	0.5 GB (2.0 - 2.5 GB)	CCA	CFG3 Cache Coherency Attributes	0x3: Cacheable, noncoherent, write-back, write allocate
SegCtl1 (CFG3)	[19]		EU	CFG3 Error	1'b0: CP0 <i>Status.ERL</i> ignored
SegCtl1 (CFG3)	[22:20]		AM	CFG3 Region Type	3'b000: Kernel unmapped region. This is kseg0.
SegCtl1 (CFG3)	[31:25]		PA	CFG3 Physical Address Bits [35:29]	0x0: Points to 0.0 - 0.5 GB physical address region.
SegCtl2 (CFG4)	[2:0]	1.0 GB (1.0 - 2.0 GB)	CCA	CFG4 Cache Coherency Attributes	Not defined since it is a mapped region.
SegCtl2 (CFG4)	[3]		EU	CFG4 Error	1'b1: CP0 <i>Status.ERL</i> bit set.
SegCtl2 (CFG4)	[6:4]		AM	CFG4 Region Type	3'b011: Kernel/Supervisor/User mapped region. This is upper half of kuseg (0x4000_0000 - 0x7FFF_FFFF).
SegCtl2 (CFG4)	[15:9]		PA	CFG4 Physical Address Bits [35:29]	7'b000001x: Points to 1.0 - 2.0 GB region.



**Table 3.11 SegCtl0 - SegCtl2 Register Settings in the Legacy Configuration**

Register	Bits	Segment Size	Name	Definition	Reset value
SegCtl2 (CFG5)	[18:16]	1.0 GB (0.0 - 1.0 GB)	CCA	CFG5 Cache Coherency Attributes	Not defined since it is a mapped region.
SegCtl2 (CFG5)	[19]		EU	CFG5 Error	1'b1: CP0 <i>Status.ERL</i> bit set.
SegCtl2 (CFG5)	[22:20]		AM	CFG5 Region Type	3'b011: Kernel/Supervisor/User mapped region. This is the lower half of kuseg (0x0000_0000 - 0x3FFF_FFFF).
SegCtl2 (CFG5)	[31:25]		PA	CFG5 Physical Address Bits [35:29]	7'b000000x: Points to 0.0 - 1.0 GB region.

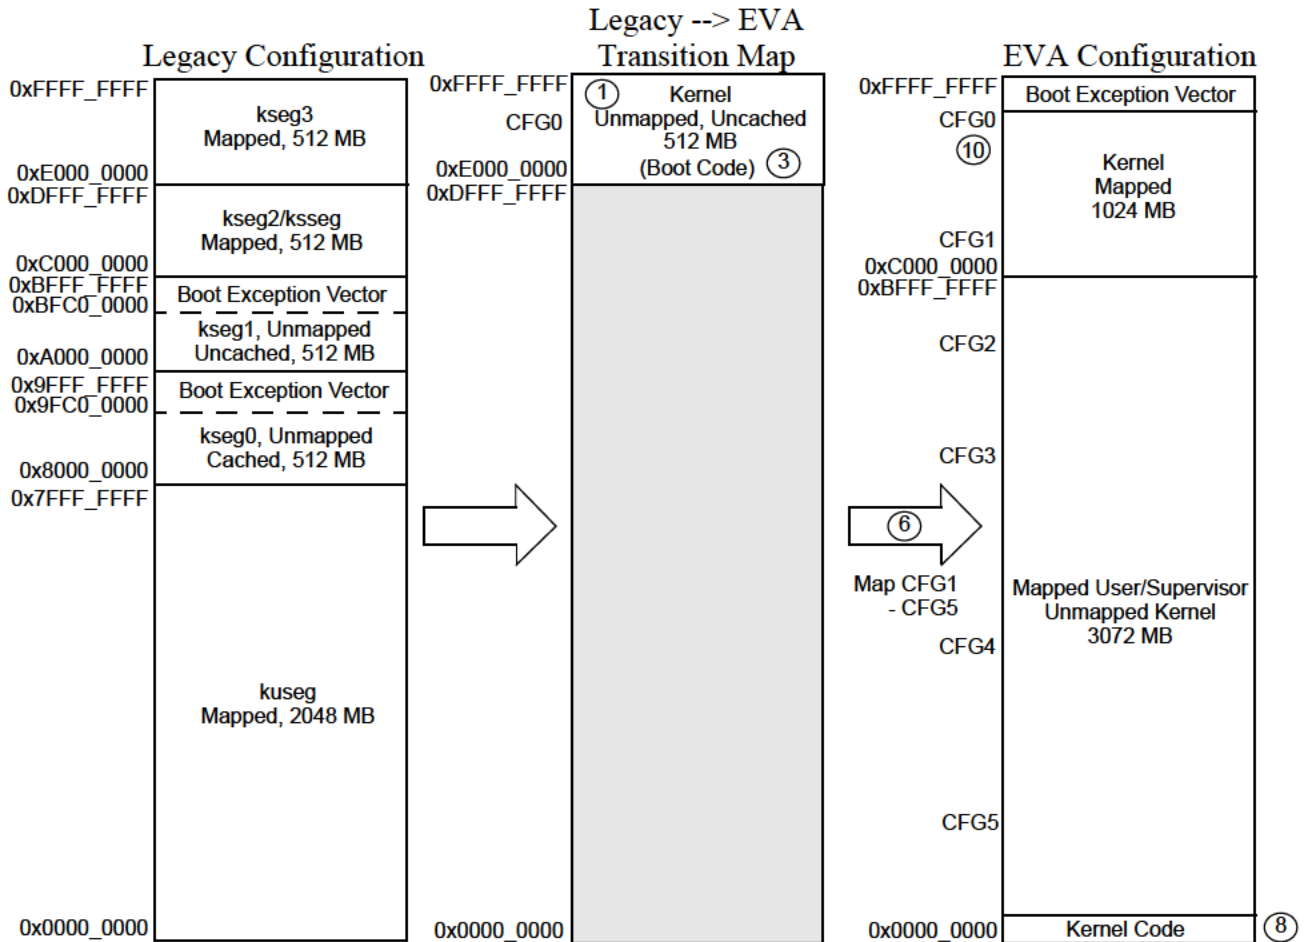
To make the transition from Legacy to EVA, software can execute the following steps. These steps are also called out in [Figure 3.21](#).

1. Temporarily set bits 6:4 of the *SegCtl0* register (CP0 Register 5, Select 2) to a value of 0x0, and bits 2:0 to a value of 0x2 to set the CFG0 region as kernel unmapped, uncached. This is shown as #1 in [Figure 3.21](#).
2. Temporarily set bits 15:9 of the *SegCtl0* register (CP0 Register 5, Select 2) to a value of 0x0 to map the CFG0 segment to a PA (physical address) value of 0x0.
3. Unpack the boot code and copy it to the CFG0 virtual address space (3.5 GB - 4.0 GB). This is shown as #3 in [Figure 3.21](#).
4. Jump to the boot code unpacked in step 3 and continue executing the bring-up code using a PC in the range of 3.5 - 4.0 GB.
5. Initialize the caches.
6. Program all other segments other than segment 0 (since the code is currently executing out of the CFG0 segment) to the following. This is shown as #6 in [Figure 3.21](#).
  - CFG1: AM = 0x1. Mapped Kernel (MK)
  - CFG2: AM = 0x4. Mapped User/Supervisor, Unmapped Kernel (MUSUK)
  - CFG3: AM = 0x4. Mapped User/Supervisor, Unmapped Kernel (MUSUK)
  - CFG4: AM = 0x4. Mapped User/Supervisor, Unmapped Kernel (MUSUK)
  - CFG5: AM = 0x4. Mapped User/Supervisor, Unmapped Kernel (MUSUK)
7. Set the *CONFIG5.K* bit to enable the EVA addressing scheme. Note that the Segment Control registers must be set for EVA as shown in step 6 above before the *CONFIG5.K* bit is set.
8. Unpack the kernel code in the low address space (0.0 - 1.0 GB). This is shown as #8 in [Figure 3.21](#).
9. Jump to the kernel code that was extracted in the previous step.
10. Set the SegCtl0 register to the values shown in [Table 3.12](#) below. This is shown as #10 in [Figure 3.21](#).

When set for EVA, the *SegCtl0* - *SegCtl2* registers should be programmed as shown in [Table 3.12](#) below.

Note that the exact location of the boot exception vector shown mapped to the upper addresses in EVA mode in [Figure 3.21](#) is dependent on the state of the boot overlay pins as described in [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#) and all related subsections.

**Figure 3.21 Mapping the Transition from Legacy Mode to EVA Mode**



**Table 3.12 New SegCtl0 - SegCtl2 Register Settings for the EVA Configuration**

Register	Bits	Segment Size	Name	Definition	Reset value
SegCtl0 (CFG0)	[2:0]	0.5 GB (3.5 - 4.0 GB)	CCA	CFG0 Cache Coherency Attributes	3'bx - Not defined since it is a mapped region.
SegCtl0 (CFG0)	[3]		EU	CFG0 Error	1'b0
SegCtl0 (CFG0)	[6:4]		AM	CFG0 Region Type	3'h1 - Mapped kernel. This is the upper 512 MB of the 1024 MB mapped kernel space (3.5 - 4.0 GB).
SegCtl0 (CFG0)	[15:9]		PA	CFG 0 Physical Address Bits [35:29]	7'hx - Not defined since it is a mapped region
SegCtl0 (CFG1)	[18:16]	0.5 GB (3.0 - 3.5 GB)	CCA	CFG1 Cache Coherency Attributes	3'bx - Not defined since it is a mapped region
SegCtl0 (CFG1)	[19]		EU	CFG1 Error	1'b0
SegCtl0 (CFG1)	[22:20]		AM	CFG1 Region Type	3'h1 - Mapped kernel. This is the lower 512 MB of the 1024 MB mapped kernel space (3.0 - 3.5 GB).
SegCtl0 (CFG1)	[31:25]		PA	CFG1 Physical Address Bits [35:29]	7'hx - Not defined since it is a mapped region
SegCtl1 (CFG2)	[2:0]	0.5 GB (2.5 - 3.0 GB)	CCA	CFG2 Cache Coherency Attributes	3'b3 - WB
SegCtl1 (CFG2)	[3]		EU	CFG2 Error	1'b1
SegCtl1 (CFG2)	[6:4]		AM	CFG2 Region Type	3'b4 - MUSUK. Mapped user/supervisor, unmapped kernel. This is the upper 512 MB of the 3072 MB MUSUK space (2.5 - 3.0 GB).
SegCtl1 (CFG2)	[15:9]		PA	CFG2 Physical Address Bits [35:29]	7'h5 - 0xA000_0000 - 0xBFFF_FFFF (2.5 - 3.0 GB)
SegCtl1 (CFG3)	[18:16]	0.5 GB (2.0 - 2.5 GB)	CCA	CFG3 Cache Coherency Attributes	3'h3 - WB
SegCtl1 (CFG3)	[19]		EU	CFG3 Error	1'b1
SegCtl1 (CFG3)	[22:20]		AM	CFG3 Region Type	3'h4 - MUSUK. Mapped user/supervisor, unmapped kernel. This is the next 512 MB of the 3072 MB MUSUK space (2.0 - 2.5 GB).
SegCtl1 (CFG3)	[31:25]		PA	CFG3 Physical Address Bits [35:29]	7'h4 - 0x8000_0000 - 0x9FFF_FFFF (2.0 - 2.5 GB)

**Table 3.12 New SegCtl0 - SegCtl2 Register Settings for the EVA Configuration (continued)**

Register	Bits	Segment Size	Name	Definition	Reset value
SegCtl2 (CFG4)	[2:0]	1.0 GB (1.0 - 2.0 GB)	CCA	CFG4 Cache Coherency Attributes	3'h3 - WB
SegCtl2 (CFG4)	[3]		EU	CFG4 Error	1'b1
SegCtl2 (CFG4)	[6:4]		AM	CFG4 Region Type	3'h4 - MUSUK. Mapped user/supervisor, unmapped kernel. This is the next 1024 MB of the 3072 MB MUSUK space (1.0 - 2.0 GB).
SegCtl2 (CFG4)	[15:9]		PA	CFG4 Physical Address Bits [35:29]	7'h2 - 0x4000_0000 - 0x7FFF_FFFF (1.0 - 2.0 GB)
SegCtl2 (CFG5)	[18:16]	1.0 GB (0.0 - 1.0 GB)	CCA	CFG5 Cache Coherency Attributes	3'h3 - WB
SegCtl2 (CFG5)	[19]		EU	CFG5 Error	1'b1
SegCtl2 (CFG5)	[22:20]		AM	CFG5 Region Type	3'h4 - MUSUK. Mapped user/supervisor, unmapped kernel. This is the low-order 1024 MB of the 3072 MB MUSUK space (0.0 - 1.0 GB).
SegCtl2 (CFG5)	[31:25]		PA	CFG5 Physical Address Bits [35:29]	7'h0 - 0x0000_0000 - 0x3FFF_FFFF (0.0 - 1.0 GB)

### 3.7 Indexing the JTLB

In the interAptiv core, the JTLB is 64 dual entries. This is shown in [Figure 3.22](#). This value is stored in the *Index* register (CP0 register 0, Select 0).

**Figure 3.22 Index Register Format Depending on TLB Size**



The *Index* register determines which TLB entry is accessed by a **TLBWI** instruction. This register is also used for the result of a **TLBP** instruction (used to determine whether a particular address was successfully translated by the CPU). Note that a **TLBP** instruction which fails to find a match for the specified virtual address sets bit 31 of *Index* register.

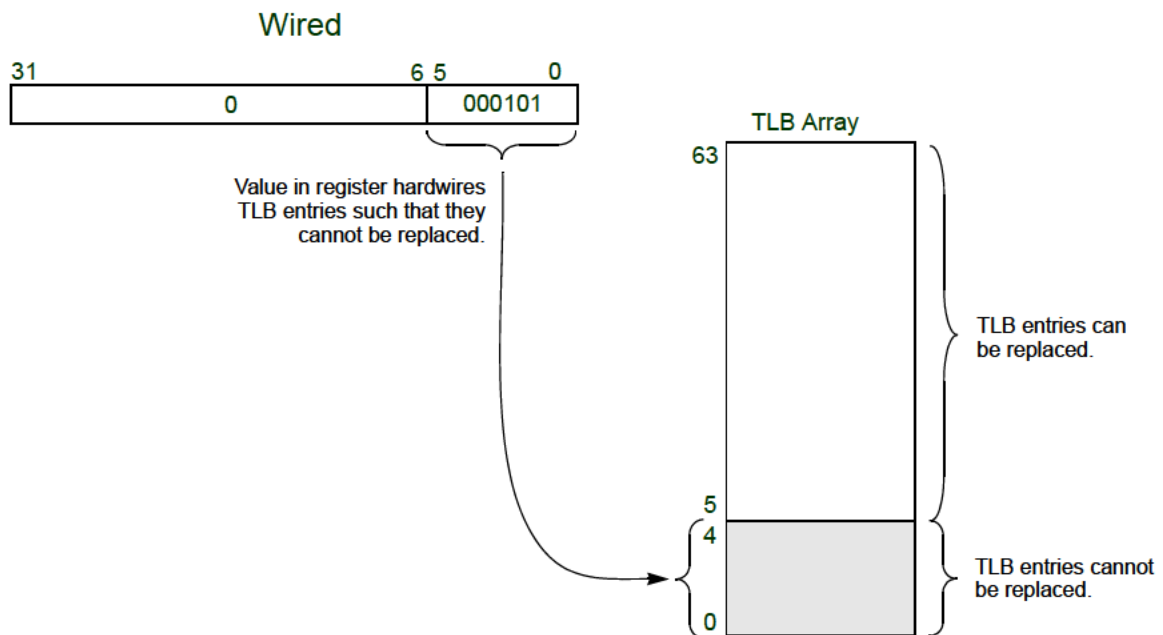
### 3.8 Hardwiring JTLB Entries

The interAptiv core allows up to 63 entries of the JTLB to be hardwired such that they cannot be replaced. This is accomplished using the *Wired* register (CP0 register 6, Select 0). The *Wired* register specifies the boundary between the wired and random entries in the JTLB. Wired entries are fixed, non-replaceable entries that cannot be overwritten by a **TLBWR** instruction. However, wired entries can be overwritten by a **TLBWI** instruction.

Note that wired entries in the JTLB must be contiguous and start from 0. For example, if the *Wired* field of this register contains a value of 5, this indicates that entries 4, 3, 2, 1, and 0 of the TLB are wired. The *Wired* register is reset to zero by a Reset exception. Note that writing to the *Wired* register may cause the *Random* register to change state.

Figure 3.23 shows an example of hardwiring the lower 5 entries of the TLB. A value of 0x0 in the *Wired* register indicates that no entries are hardwired and that all entries are available for replacement.

**Figure 3.23 Hardwiring Entries in the TLB**



### 3.9 JTLB Random Replacement

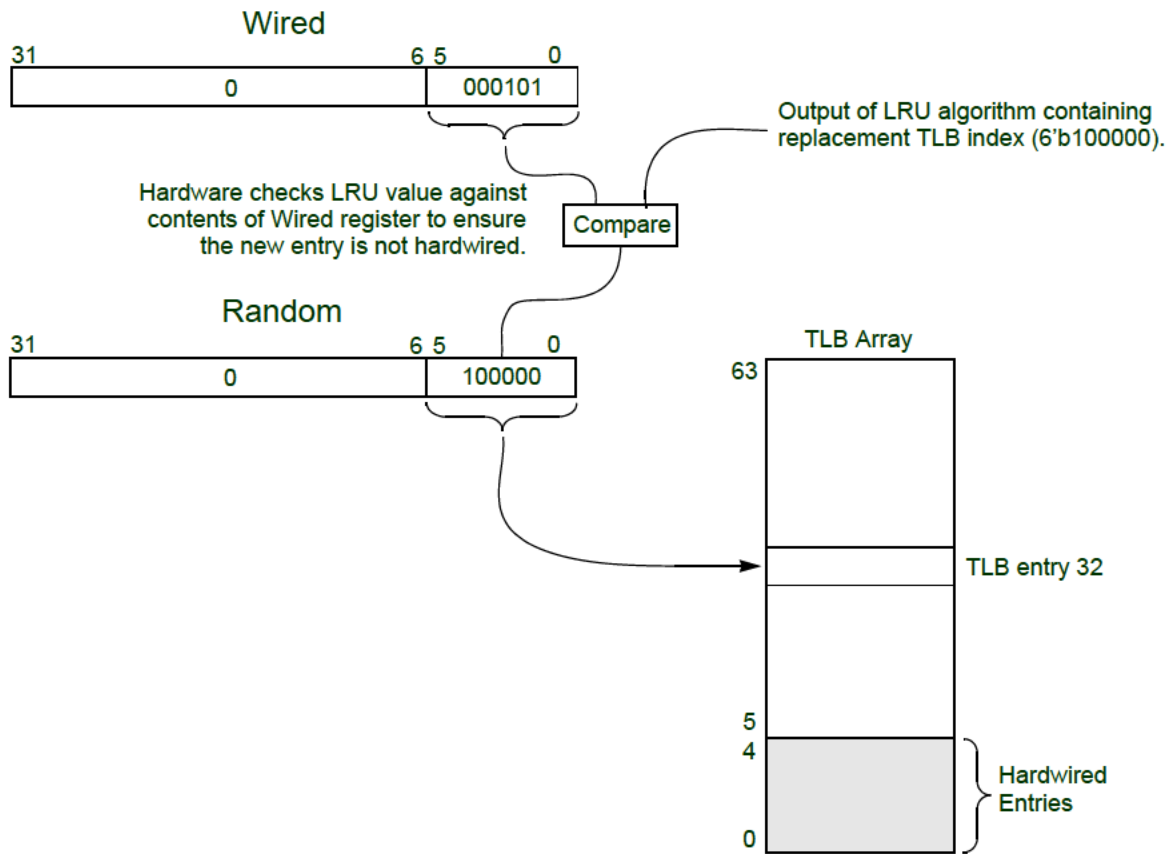
The interAptiv core performs random replacement within the 64 dual-entry JTLB using the CP0 *Random* register (CP0 register 1, Select 0). This read-only register is used to index the TLB during a **TLBWR** instruction. It provides a quick way of replacing a JTLB entry at random.

The *Random* register employs a pseudo-random least-recently-used (LRU) algorithm which ensures that no wired entries are selected. Only those LRU entries that are not in the *Wired* register are targeted for replacement. The contents of the *Random* register are modified after a JTLB write, or on a write to the *Wired* register.

The processor initializes the *Random* register to reflect the maximum number of entries (63) on a Reset exception. Note that the *Random* register is used only for JTLB accesses.

Figure 3.24 shows an example of a random replacement to entry 32 of the JTLB with the lower five entries of the JTLB hardwired.

**Figure 3.24 Random Replacement of a JTLB Entry**



## 3.10 TLB Exception Handling

The interAptiv core allows for the following types of TLB exceptions.

- Address error (AdEL or AdES)
- TLB Refill
- TLB (TLBL, TLBS)
- TLB Modified

The *Address Error* exceptions (AdEL and AdES) are used in both user mode and supervisor mode.

- On a load in user mode, an *AdEL* exception is taken when the user does not have permission for the load address being accessed.
- On a store in user mode, an *AdES* exception is taken when the user does not have permission for the store address being accessed.
- On a load in supervisor mode, an *AdEL* exception is taken when the supervisor does not have permission for the load address being accessed.
- On a store in supervisor mode, an *AdES* exception is taken when the supervisor does not have permission for the store address being accessed.

The *TLB Refill* exception is taken on any TLB miss regardless of the operating mode.

The *TLB* exceptions (TLBL and TLBS) are taken under the following conditions.

- TLBL exception: On a load in any mode, there is a TLB hit, but the valid bit for that TLB entry is not set.
- TLBS exception: On a store in any mode, there is a TLB hit, but the valid bit for that TLB entry is not set.

A *TLB Modified* exception is taken whenever there is a TLB hit and the Dirty bit associated with that entry is not set.

### 3.10.1 Overview of TLB Exception Handling Registers

The interAptiv core uses three CP0 registers to manage TLB exceptions. The exception flow in terms of these registers is described in [Section 3.10.2, "TLB Exception Flow Examples"](#).

- *Context* (CP0 register 4, Select 0): Contains the pointer to an entry in the page table entry (PTE) array.
- *ContextConfig* (CP0 register 4, Select 1): Defines the range of bits used by the *Context* register into which the high order bits of the virtual address causing the TLB exception will be written depending on the page size.
- *BadVAddr* (CP0 register 8, Select 0): Stores the virtual address that caused the exception.


#### 3.10.1.1 Context Register

The *Context* register is a read/write register containing a pointer to an entry in the page table entry (PTE) array. When a TLB exception is taken, hardware performs the bit shifting and manipulation of the value stored in the *BadVAddr* register and places the result into the *BadVPN2* field of the *Context* register. This eliminates software from having to perform this function manually.

A TLB exception causes the virtual address to be written to a variable range of bits, defined as (X-1):Y of the *Context* register. This range corresponds to the contiguous range of set bits in the *ContextConfig* register. Bits 31:X, Y-1:0 are read/write to software and are unaffected by the exception. Software sets the *ContextConfigPTEBase* field to point to the base address of a page table in memory. The *ContextConfigBadVPN2* is derived from the virtual address associated with the exception.

[Figure 3.25](#) shows the format of the *Context* register. Refer to [Section 3.10.2, "TLB Exception Flow Examples"](#) for more information on the usage of this register.

**Figure 3.25 Context Register Format**



#### 3.10.1.2 ContextConfig Register

The *ContextConfig* register defines the bits of the *Context* register into which the high order bits of the virtual address causing a TLB exception will be written (*BadVPN2*), and how many bits of that virtual address will be extracted. In the *Context* register, bits above the selected *BadVPN2* field are read/write to software and serve as the *PTEBase* field. Bits below the selected *BadVPN2* field serve as the *PTEBaseLow* field.


Software writes a set of contiguous ones to the *ContextConfigVirtualIndex* field. Hardware then determines which bits of this register are high and low. The highest order bit that is a logic '1' serves as the MSB of the *BadVPN2* field of the *Context* register. The lowest order bit that is a logic '1' serves as the LSB of the *BadVPN2* field of the *Context* register. A value of all zero's in the *VirtualIndex* field means that the full 32 bits of the *Context* register are R/W for software and are unaffected by TLB exceptions.

A value of all ones in the *ContextConfigVirtualIndex* field means that the full 21 bits of the faulting virtual address will be copied into the context register, making it duplicate the *BadVAddr* register. A value of all zeroes means that the full 32 bits of the *Context* register are R/W for software and unaffected by TLB exceptions.

[Figure 3.26](#) shows the formats of the *ContextConfig* Register. Refer to [Section 3.10.2, "TLB Exception Flow Examples"](#) for more information on use of the this register.



**Figure 3.26 ContextConfig Register Format**



It is permissible to implement a subset of the *ContextConfig* register, in which some number of bits are read-only and set to one or zero as appropriate. It is possible for software to determine which bits are implemented by alternately writing all zeroes and all ones to the register, and reading back the resulting values. Table 3.13 describes some useful *ContextConfig* values. In this table, note that for a page table entry size of 32 bits per page, a total of 64 bits are copied from memory to support the dual-entry structure of the JTLB. In this case, the lower 32 bits would be copied to entry 0 of the dual entry structure, and the upper 32 bits would be copied to entry 1 of the structure. The same is true for a page table with 64 bits per page. In this case, 128 bits would be fetched from memory.

**Table 3.13 Example ContextConfig Values — Single Level Page Table Organization**


Value	Page Table Organization	Page Size	Page Table Entry Size	Memory Structure
0x007F_FFF0	Single Level	4K	64 bits/page	128-bit
0x003F_FFF8	Single Level	4K	32 bits/page	64-bit

### 3.10.1.3 BadVAddr Register

The *BadVAddr* is a 32-bit read-only register which holds the virtual address which caused the last address-related exception. It is set for the exception types shown at the beginning of Section 3.10, "TLB Exception Handling".

Note that the *BadVAddr* register does not capture address information for cache or bus errors, since they are not addressing errors.

**Figure 3.27 BadVAddr Register Format**



## 3.10.2 TLB Exception Flow Examples

The following two examples show the flow of a TLB exception for the single level and dual level page table configurations.

### 3.10.2.1 Single Level Table Configuration

When a JTLB error occurs, hardware writes the most recent virtual address that caused the error into bits 31:0 of the read-only *BadVAddr* register. The number of bits used by hardware to index the page table depends on the page size. For example, with a 4 KByte page size, hardware uses bits 31:13 of the *BadVAddr* register, along with the *PTEBase* field of the *Context* register, to determine the address that caused the exception.

Hardware assembles this information and places the result into the *Context* register. Use of the *Context* and *ContextConfig* registers eliminates software from having to derive the page table index manually. Depending on the page table architecture, software programs the *ContextConfig* register to indicate how many bits of the *BadVAddr* reg-

ister are used by hardware to program the *Context* register. This determines the size of both the *Context<sub>BadVPN2</sub>* and *Context<sub>PTEBase</sub>* fields.

The example shown in [Figure 3.28](#) is for a single level table configuration with a 4 KByte page size and 32 bits per page.

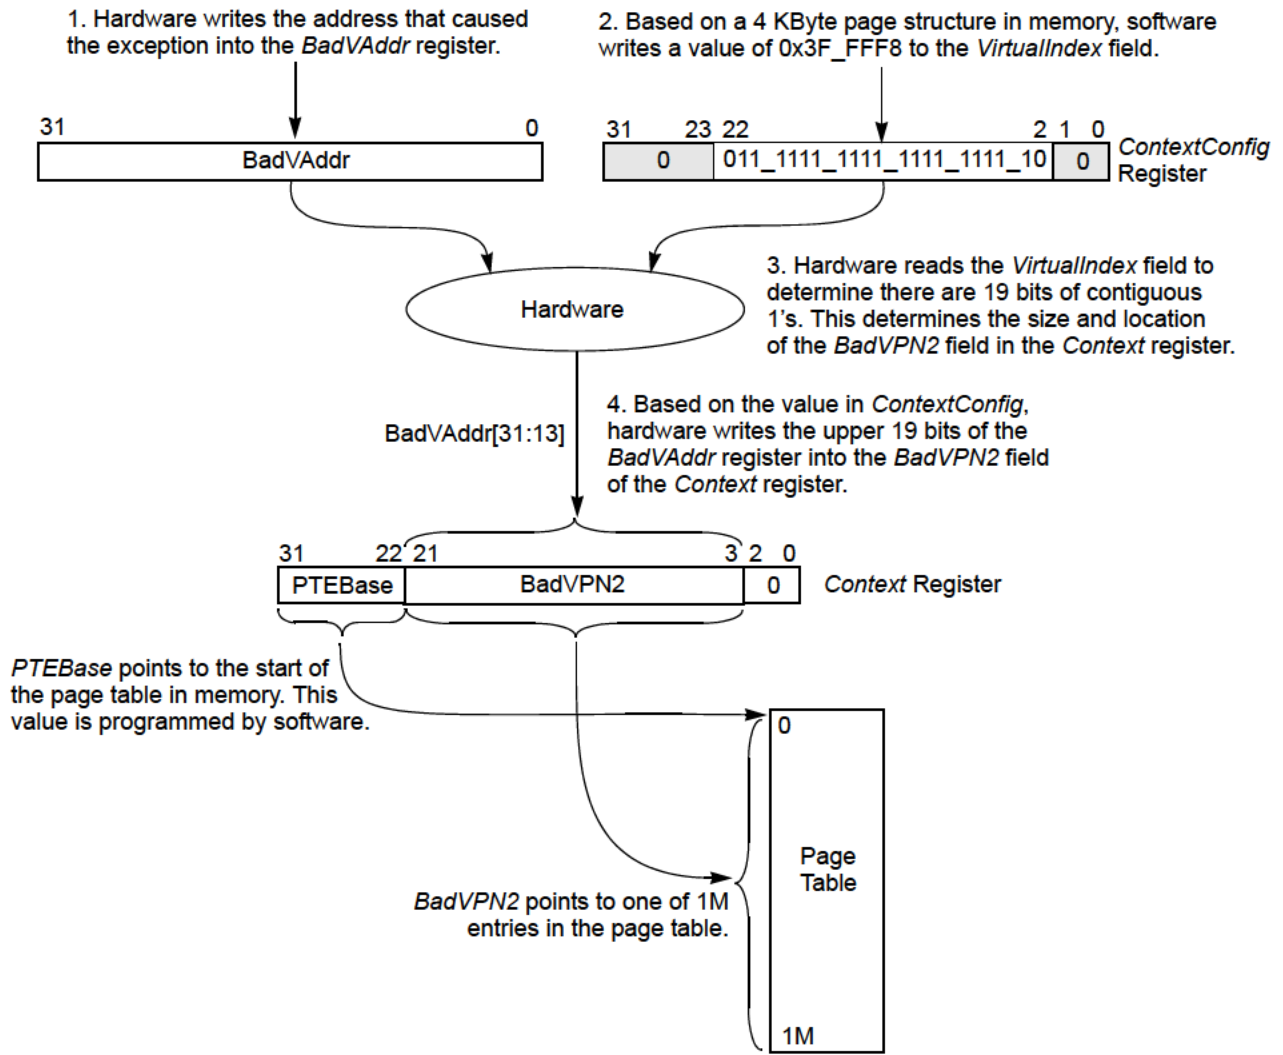
When an exception is taken, hardware writes the address that caused the exception into the *BadVAddr* register. Because the page table is single level and the page size is already known to be 4 KBytes, software programs a value of 0x3F\_FFF8 into the *ContextConfig<sub>VirtualIndex</sub>* field. This value indicates the following information:

- The lower three bits of this value are 0, indicating that a 64-bit memory structure is being accessed. For this 64-bit value, the lower 32 bits are written to the entry 0 of the dual-entry TLB, and the upper 32 bits are written to entry 1 of the same TLB entry. Since the lower 3 bits of this field are zero, bit 3 (the first bit that is set) is used to define the low-order bit of the *BadVPN2* field in the *Context* register.
- The highest-order bit that is 1 in this field is bit 21. This indicates that bit 21 is the last bit of the *BadVPN2* field in the *Context* register. As a result, the *PTEBase* field of the *Context* register occupies bits 31:22.

Based on this information, hardware assembles the value in the *Context* register as follows:

- *Context<sub>PTEBase</sub>* = bits 31:22. Indicates the base address of the page table in memory. This 10-bit value is a pointer to the start of the page table in memory.
- *Context<sub>BadVPN2</sub>* = bits 21:3. Hardware copies bits 31:13 of the *BadVAddr* register into this field. This 19-bit value is a pointer for up to 1M entries in each page table selected by the *Context<sub>PTEBase</sub>* field. Bits 12:0 of the *BadVAddr* register are not used in this case since the page size is 4 KBytes.
- *Context<sub>PTEBaseLow</sub>* = bits 2:0. Indicates access to a 64-bit memory location.

**Figure 3.28 TLB Exception Flow Example — Single Level Table, 4 KB Page Size**



### 3.10.2.2 Dual Level Table Configuration

The TLB exception flow for a dual level page table structure is similar to that of a single level table described in [Section 3.10.2.1, "Single Level Table Configuration"](#). The upper bits of *PTEBase* are used to select the location of the first level table in memory. The *BadVPN2* field of the *Context* register is used to index the first level table and acts as a pointer to each of the second level tables in the page table array.

When a JTLB error occurs, the most recent virtual address that caused the error is stored in bits 31:0 of the read-only *BadVAddr* register. The number of bits in the *BadVAddr* register used by hardware to index the page table depends on the page size and table organization.

Hardware assembles this information and places the result into the *Context* register. Use of the *Context* and *ContextConfig* registers eliminates software from having to derive the page table index manually. Depending on the page table architecture, software programs the *ContextConfig* register to indicate how many bits of the *BadVAddr* register are used by hardware to program the *Context* register. This determines the size of both the *Context<sub>BadVPN2</sub>* and *Context<sub>PTEBase</sub>* fields.

The example shown in [Figure 3.29](#) is for a dual level table configuration with a 4 KByte page size and 32 bits per page.

When an exception is taken, hardware writes the address that caused the exception into the 32-bit *BadVAddr* register. Because each table in this example contains 1K entries, software programs a value of 0x00\_0FFC into the *ContextConfigVirtualIndex* field. This value indicates the following information:

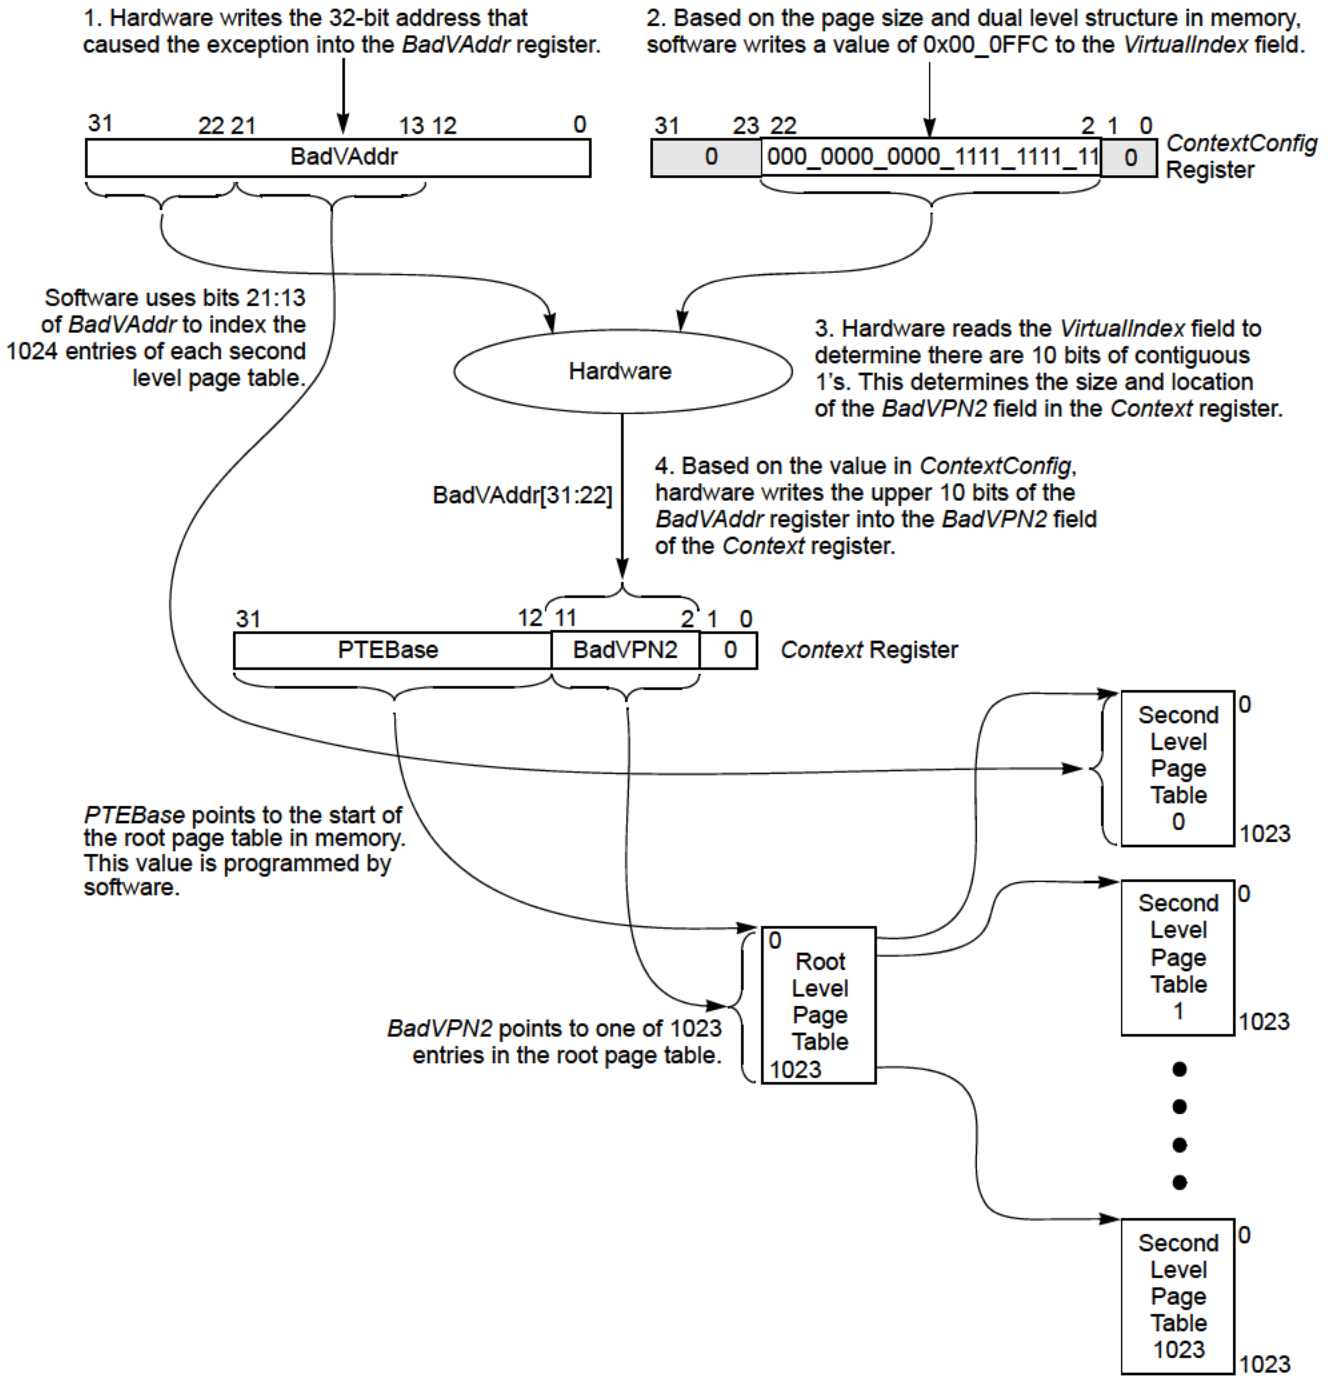
- The lower two bits of this value are 0, indicating that a 32-bit memory structure is being accessed. This also indicates that bit 2 will be the low-order bit for the *ContextBadVPN2* field.
- The highest-order bit that is '1' in the *ContextConfigVirtualIndex* field is bit 11. This indicates that bit 11 will be the highest-order bit of the *ContextBadVPN2* field. As a result, the *ContextPTEBase* field occupies bits 31:12. This field is used to access the location of the root level page table in memory.

Based on this information, hardware assembles the context register as follows:

- *ContextPTEBase* = bits 31:12. Indicates the base address of the page table in memory. This 20-bit value is a pointer to the root page table in memory.
- *ContextBadVPN2* = bits 11:2. Based on the state of the *ContextConfigVirtualIndex* field in this example, hardware copies bits 31:22 of the *BadVAddr* register into this field. This 10-bit value is a pointer to the 1024 entries in the root page table selected by the *ContextPTEBase* field. Bits 12:0 of the *BadVAddr* register are not used in this case since the page size is 4 KBytes.
- *ContextPTEBaseLow* = bits 1:0. Indicates access to a 32-bit memory location.

As stated above, bits 31:22 of the *BadVAddr* register are copied into the *BadVPN2* field of the *Context* register and are used to select one of 1024 entries in the root page table. Each of these entries acts as a pointer to one of the 1024 second level tables. Software uses bits 21:13 of the *BadVAddr* register to index one of 1024 entries in each second level page table. This concept is shown in [Figure 3.29](#).

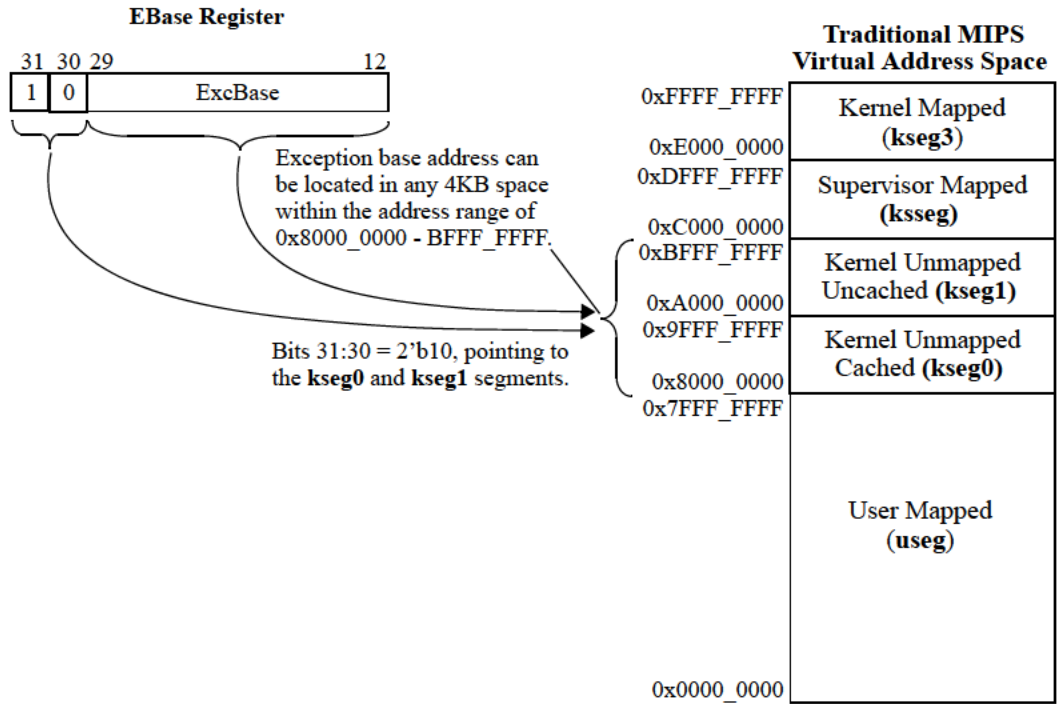
**Figure 3.29 TLB Exception Flow Example — Dual Level Table, 4 KB Page Size**



### 3.11 Exception Base Address Relocation

The interAptiv core allows the base address of an exception vector to be relocated when programmable memory segmentation is enabled. The base address of the exception is stored in the CP0 *EBase* register. In previous generation MIPS32 processors, bits 31:30 of the *EBase* Register were not writeable and had a fixed value of 2'b10 so that the exception handler would be executed from the *kseg0* or *kseg1* segments. This concept is shown in Figure 3.30.

Figure 3.30 Location of Exception Vector Base Address in Traditional MIPS Virtual Address Space



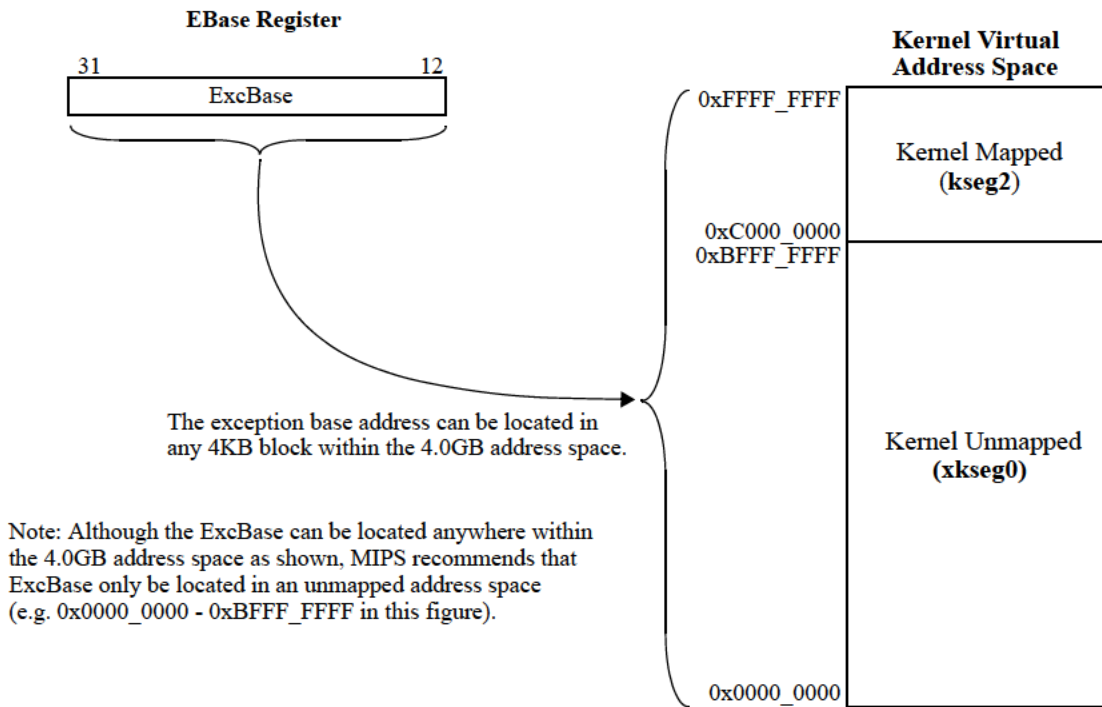
When programmable memory segmentation is enabled, the size of the exception base address is determined by the state of the *WG* bit in the CP0 *EBase* register (CP0 register 15, Select 1). At reset, the *WG* bit is cleared by default and bits 31:30 of the *EBase* Register are forced to a value of 2'b10 by hardware as described above. This is shown in Figure 3.30 above.

When the *WG* bit is set, bits 31:30 of the *ExcBase* field become writeable and are used to relocate the exception base address to other segments after they have been setup using the *SegCtl0* through *SegCtl2* registers. This is shown in Figure 3.31.

Note that if the *WG* bit is set by software (allowing bits 31:30 to become part of the *ExcBase* field) and then cleared, bits 31:30 can no longer be written by software and the state of these bits remains unchanged for any writes after *WG* was cleared. Therefore, it is the responsibility of software to write a value of 2'b10 to bits 31:30 of the *EBase* register prior to clearing the *WG* bit if it wants to ensure that future exceptions will be executed from the *kseg0* or *kseg1* segments.

Note that the *WG* bit is different from the *CV* bit in the *Config5* register. Although their functions are similar, the *CV* bit applies only to cache error exceptions, whereas the *WG* bit applies to all exceptions.

**Figure 3.31 Location of Exception Vector Base Address Using the Enhanced Virtual Addressing Scheme**



### 3.12 TLB Duplicate Entries

The JTLB entries come up in a random state on power-up and must be initialized by hardware before use. Typically, bootstrap software initializes each entry in the TLB. Since the JTLB is a fully-associative array and entries are written by index, it is possible to load duplicate entries, where two or more entries match the same virtual address/ASID.

If duplicate entries are detected on a TLB write, no machine check is generated and the older entries are just invalidated. The new entry gets written. When writing to the TLB, all entries of the JTLB are searched for duplicates.

### 3.13 Modes of Operation

The MMU's virtual-to-physical address translation is determined by the mode in which the processor is operating. The interAptiv core operates in one of four modes:

- User mode
- Supervisor mode
- Kernel mode
- Debug mode

User mode is most often used for application programs. Supervisor mode is an intermediate privilege level with access to an additional region of memory and is only supported with the TLB-based MMU. Kernel mode is typically



used for handling exceptions and privileged operating system functions, including CP0 management and I/O device accesses. Debug mode is used for software debugging and usually occurs within a software development tool.

**Table 3.14 Selecting the Addressing Mode**

Mode	Status			Debug	Description
	EXL	ERL	KSU	DM	
User	0	0	2'b2	0	User addressing mode. In this mode, a TLB miss goes to the TLB Refill Handler.
Supervisor	0	0	2'b1	0	Supervisor addressing mode. In this mode, a TLB miss goes to the TLB Refill Handler.
Kernel	x	x	2'b0	0	Kernel addressing mode. In this mode, a TLB miss goes to the TLB Refill Handler.
	x	1	x	0	Kernel addressing mode. In this mode, a TLB miss goes to the TLB Refill Handler.
	1	x	x	0	Kernel addressing mode. In this mode, a TLB miss goes to the general exception handler as opposed to the TLB Refill handler.
Debug	x	x	x	1	Debug mode.

### 3.13.1 Virtual Memory Segments

The interAptiv core supports the following virtual memory schemes.

- Traditional MIPS32 virtual address space, which contains fixed address ranges for the various user and kernel segments.
- Enhanced Virtual Address (EVA) mode that allows the kernel and user address spaces to be programmed to different sizes depending on the needs of the application.

#### ***MIPS32 Virtual Address Space — Legacy Addressing Scheme***

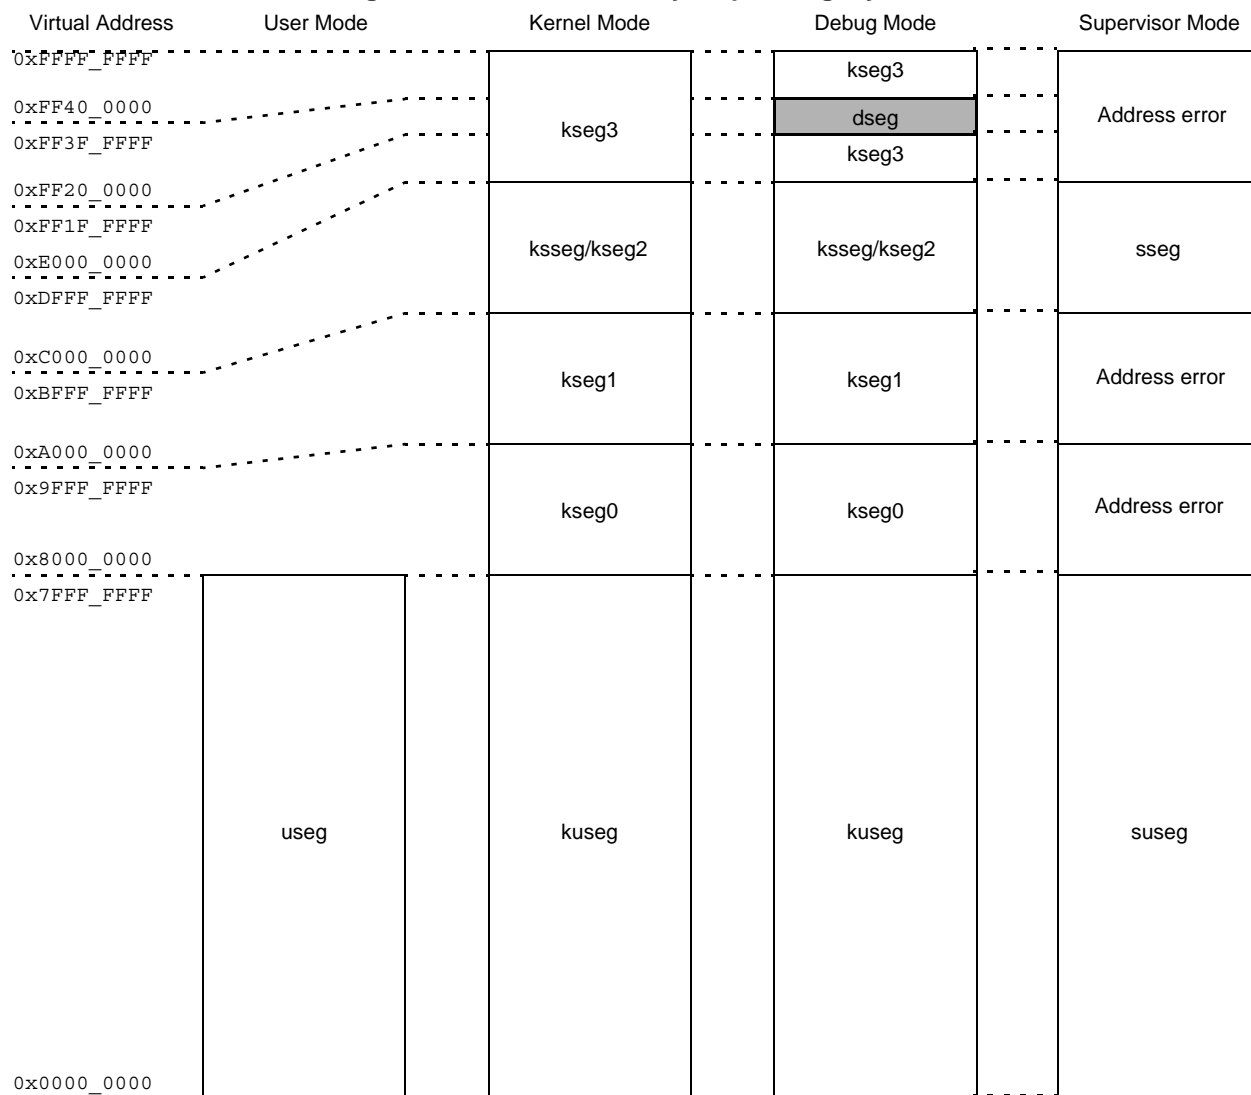
In the legacy mode, the MIPS32 architecture supports a 4 GByte virtual address space that is partitioned into a number of segments, each characterized by a set of attributes defined by hardware and software. The virtual memory segments are different depending on the mode of operation. [Figure 3.32](#) shows the segmentation for the 4 GByte ( $2^{32}$  bytes) virtual memory space, addressed by a 32-bit virtual address, for each of the four modes.

User mode accesses are limited to a subset of the virtual address space (0x0000\_0000 to 0x7FFF\_FFFF) and can be inhibited from accessing CP0 functions. In User mode, virtual addresses 0x8000\_0000 to 0xFFFF\_FFFF are invalid and cause an exception if accessed. Supervisor mode adds access to sseg (0xC000\_0000 to 0xDFFF\_FFFF). kseg0, kseg1, and kseg3 will still cause exceptions if they are accessed. In Kernel mode, software has access to the entire address space, as well as all CP0 registers.

Debug mode is entered on a debug exception. While in Debug mode, the debug software has access to the same address space and CP0 registers as Kernel mode. In addition, while in Debug mode, the CPU has access to the debug segment (dseg). This area overlays part of the kernel segment kseg3. Access to dseg in Debug mode can be turned on or off, allowing full access to the entire kseg3 in Debug mode, if so desired.



**Figure 3.32 Virtual Memory Map — Legacy Mode**

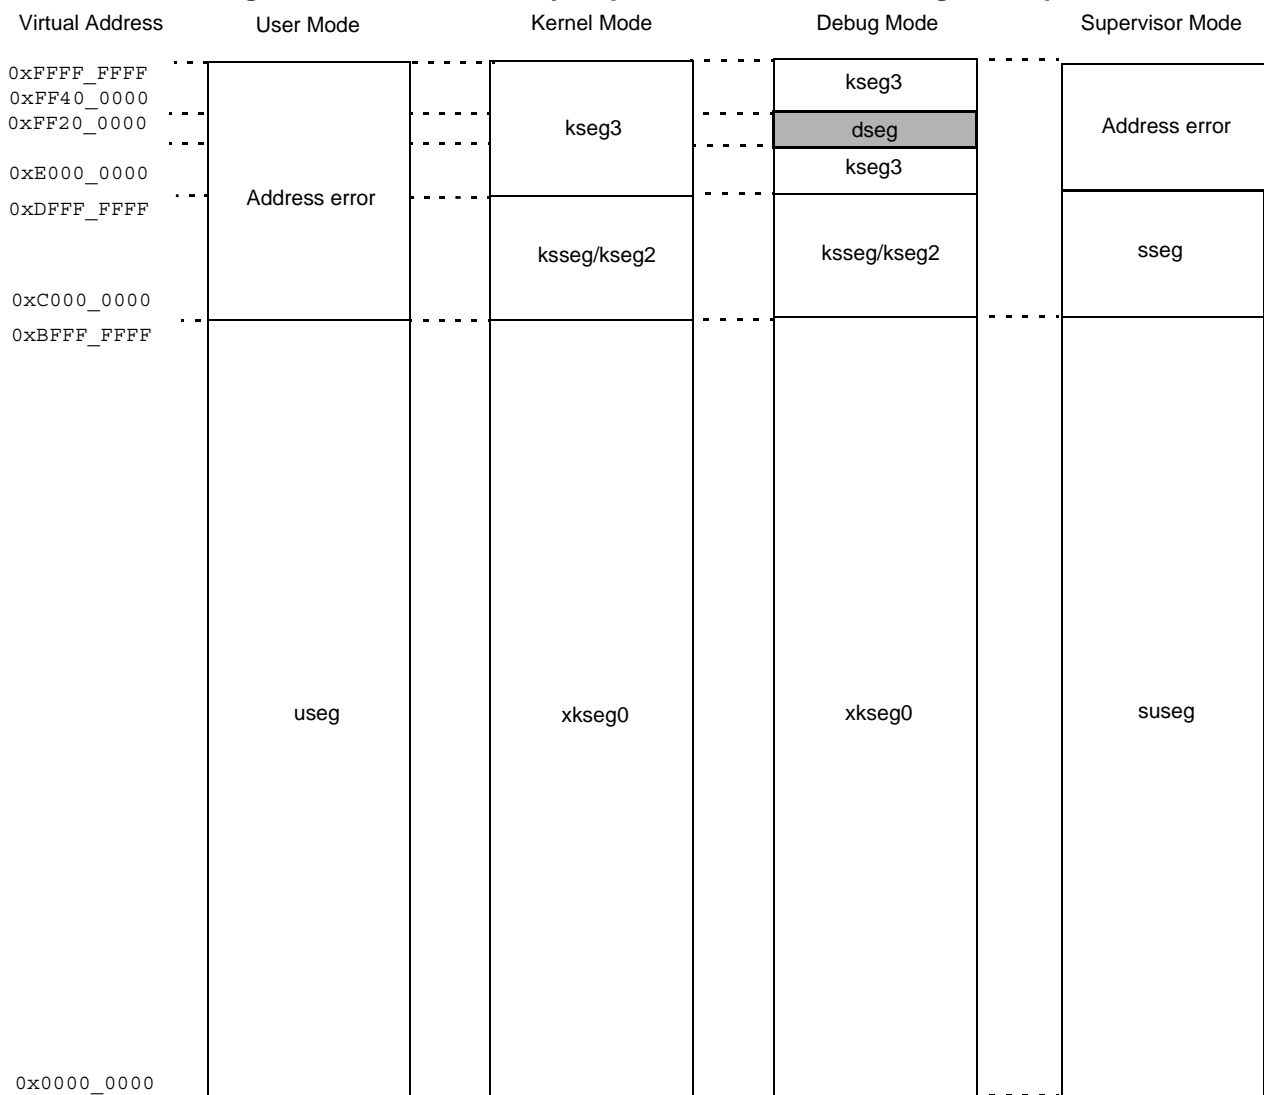


**MIPS32 Virtual Address Space — EVA Addressing Scheme**

In the EVA addressing scheme, the MIPS32 architecture supports a 4 GByte virtual address space that is partitioned into a number of programmable segments using the SegCtl0 through SegCtl2 registers. The EVA scheme is described in [Section 3.5, "Enhanced Virtual Address"](#). The virtual memory segments are different depending on the mode of operation and the programming of these registers.

[Figure 3.33](#) shows an example segmentation for the 4 GByte ( $2^{32}$  bytes) virtual memory space, with the kernel and user segments being defined as 3 GB in size. Note that is only an example and the sizes of these memory segments can be increased or decreased depending on the needs of the application.

**Figure 3.33 Virtual Memory Map — EVA Mode, 3GB xkseg0 Example**



Segments can be mapped or unmapped, as described in the following subsections.

### 3.13.1.1 Unmapped Segments

An unmapped segment does not use the TLB to translate virtual to physical addresses. Especially after reset, it is important to have unmapped memory segments, because the TLB is not yet programmed to perform the translation. Unmapped segments have a simple translation from virtual to physical address.

Except for kseg0, unmapped segments are always uncached. The cacheability of kseg0 is set in the K0 field of the CP0 *Config* register.

### 3.13.1.2 Mapped Segments

A mapped segment uses the TLB to translate from virtual to physical addresses. The translation of mapped segments are handled on a per-page basis. Included in this translation is information defining whether the page is cacheable or not, and the protection attributes that apply to the page.


## 3.13.2 User Mode

In user mode, a single uniform virtual address space, called the user segment (useg), is available. The size of the user segment depends on the virtual addressing mode used.

### 3.13.2.1 User Mode Legacy Configuration

In the legacy mode, the user segment occupies the lower 2 GB of virtual address space. The user segment starts at address 0x0000\_0000 and ends at address 0x7FFF\_FFFF. Accesses to all other addresses cause an address error exception. This is shown in [Figure 3.34](#).

**Figure 3.34 User Mode Virtual Address Space — Legacy Configuration**



The processor operates in User mode when the *Status* register contains the following bit values:


- *KSU* = 0b10
- *EXL* = 0
- *ERL* = 0

In addition to the above values, the *DM* bit in the *Debug* register must be 0.

### 3.13.2.2 User Mode EVA Configuration

In EVA mode, the user segment occupies up to the lower 3.0 GB of virtual address space. The user segment starts at address 0x0000\_0000 and ends at address 0xBFFF\_FFFF. Accesses to all other addresses cause an address error exception. This is shown in [Figure 3.35](#).

**Figure 3.35 User Mode Virtual Address Space — EVA Mode**



All references to useg are mapped through the TLB. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address before translation. Also, bit settings within the TLB entry for the page determine the cacheability of a reference.


### 3.13.3 Supervisor Mode

In supervisor mode, two uniform virtual address spaces are available, legacy and EVA. The size of these spaces depends on the addressing mode used.

#### 3.13.3.1 Supervisor Mode — Legacy Configuration

In the supervisor mode - legacy configuration, the 2GB virtual address space called the supervisor user segment (suseg), and a 512 MByte virtual address space called the supervisor segment (sseg). The supervisor-mode virtual address space is shown in [Figure 3.36](#).

**Figure 3.36 Supervisor Mode Virtual Address Space**



The supervisor user segment begins at address 0x0000\_0000 and ends at address 0x7FFF\_FFFF. The supervisor segment begins at 0xC000\_0000 and ends at 0xDFFF\_FFFF. Accesses to all other addresses in Supervisor mode cause an address error exception.

The processor operates in Supervisor mode when the *Status* register contains the following bit values:

- $KSU = 2'b01$
- $EXL = 0$
- $ERL = 0$

In addition to the above values, the *DM* bit in the *Debug* register must be 0.

Table 3.15 lists the characteristics of the Supervisor mode segments in the legacy.

**Table 3.15 Supervisor Mode Segments — Legacy Configuration**

Address-Bit Value	Status Register				Segment Name	Address Range	Segment Size
	Bit Value						
	EXL	ERL	UM	SM			
32-bit A(31) = 0	0	0	0	1	suseg	0x0000_0000 --> 0x7FFF_FFFF	2 GByte (2 <sup>31</sup> bytes)

**Table 3.15 Supervisor Mode Segments — Legacy Configuration**

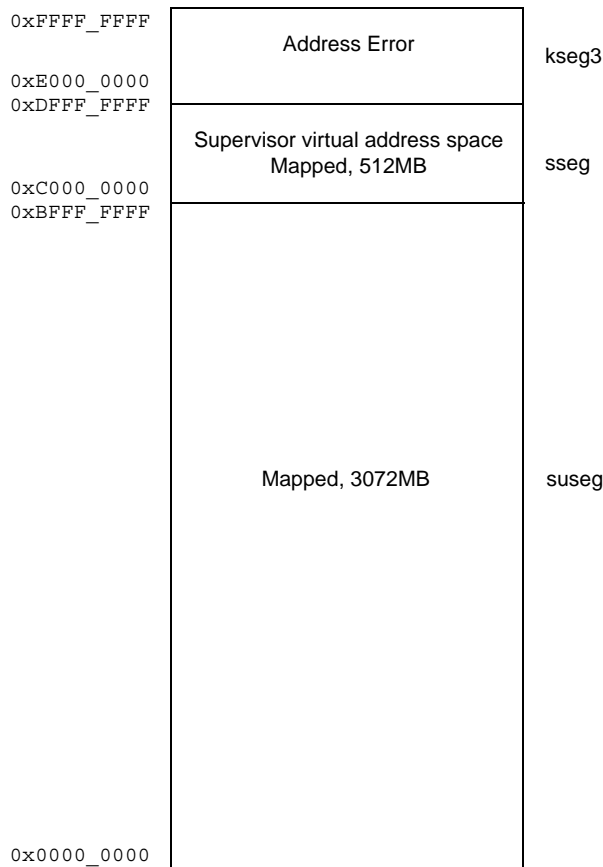
Address-Bit Value	Status Register				Segment Name	Address Range	Segment Size
	Bit Value						
	EXL	ERL	UM	SM			
32-bit A(31:29) = 3'b110	0	0	0	1	sseg	0xC000_0000 -> 0xDFFF_FFFF	512MB (2 <sup>29</sup> bytes)

The system maps all references to *suseg* and *sseg* through the TLB. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address before translation. Also, bit settings within the TLB entry for the page determine the cacheability of a reference.

### 3.13.3.2 Supervisor Mode — EVA Configuration

In the supervisor mode - EVA configuration, the virtual address spaces are called the supervisor user segment (*suseg*), and the supervisor segment (*sseg*). The size of each field depends on the programming of the segment control registers. Figure 3.37 shows an example of a 3GB supervisor user segment (*suseg*), and a 512MB supervisor segment (*sseg*).

**Figure 3.37 Supervisor Mode Virtual Address Space — EVA Configuration**



The supervisor user segment begins at address 0x0000\_0000 and ends at address 0xBFFF\_FFFF. The supervisor segment begins at 0xC000\_0000 and ends at 0xDFFF\_FFFF. Accesses to all other addresses in Supervisor mode cause an address error exception.

Note that the sseg segment is programmed when the bits 17:14 of the *SegCtl0* register contains a value of 0x2. This causes the address range of 0xC000\_0000 to 0xDFFF\_FFFF to be mapped in supervisor space. However, while in supervisor mode, where 0x0000\_0000 - 0xBFFF\_FFFF is defined as the suseg segment, the 0xC000\_0000 to 0xDFFF\_FFFF address range can be configured as kernel mapped. This occurs when 17:14 of the *SegCtl0* register contains a value of 0x1. Refer to [Table 3.4](#) and [Table 3.5](#) for more information.

### 3.13.4 Kernel Mode

In kernel mode, two uniform virtual address spaces are available, legacy and EVA. The size of these spaces depends on the addressing mode used as described below.

#### 3.13.4.1 Kernel Mode — Legacy Configuration

The processor operates in Kernel mode when the *DM* bit in the *Debug* register is 0 and the *Status* register contains one or more of the following values:


- *KSU* = 2'b00
- *ERL* = 1
- *EXL* = 1

When a non-debug exception is detected, *EXL* or *ERL* will be set and the processor will enter Kernel mode. At the end of the exception handler routine, an Exception Return (ERET) instruction is generally executed. The ERET instruction jumps to the Exception PC, clears *ERL*, and clears *EXL* if *ERL*=0. This may return the processor to User mode.

In Kernel mode, a program has access to the entire virtual address space. Kernel mode virtual address space is divided into regions differentiated by the high-order bits of the virtual address, as shown in [Figure 3.38](#). The characteristics of kernel-mode segments are listed in [Table 3.16](#).

The CPU enters Kernel mode both at reset and when an exception is recognized.

**Figure 3.38 Kernel Mode Virtual Address Space — Legacy Configuration**



**Table 3.16 Kernel Mode Segments**

Address-Bit Values	Status Register Is One of These Values			Segment Name	Address Range	Segment Size
	KSU	EXL	ERL			
A(31) = 0	(KSU = 00 <sub>2</sub> or EXL = 1 or ERL = 1) and DM = 0			kuseg	0x0000_0000 through 0x7FFF_FFFF	2 GBytes (2 <sup>31</sup> bytes)
A(31:29) = 3'b100				kseg0	0x8000_0000 through 0x9FFF_FFFF	512 MBytes (2 <sup>29</sup> bytes)
A(31:29) = 3'b101				kseg1	0xA000_0000 through 0xBFFF_FFFF	512 MBytes (2 <sup>29</sup> bytes)
A(31:29) = 3'b110				ksseg/kseg2	0xC000_0000 through 0xDFFF_FFFF	512 MBytes (2 <sup>29</sup> bytes)
A(31:29) = 3'b111				kseg3	0xE000_0000 through 0xFFFF_FFFF	512 MBytes (2 <sup>29</sup> bytes)



### **Kernel Mode, User Space (kuseg)**

In Kernel mode, when the most-significant bit of the virtual address (A31) is cleared, the 32-bit kuseg virtual address space is selected and covers the full  $2^{31}$  bytes (2 GBytes) of the current user address space mapped to addresses 0x0000\_0000 - 0x7FFF\_FFFF. For cores with TLBs, the virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

When  $ERL = 1$  in the *Status* register, the user address region becomes a  $2^{31}$ -byte unmapped and uncached address space. While in this setting, the kuseg virtual address maps directly to the same physical address, and does not include the ASID field.

### **Kernel Mode, Kernel Space 0 (kseg0)**

In Kernel mode, when the most-significant three bits of the virtual address are 3'b100, 32-bit kseg0 virtual address space is selected; it is the  $2^{29}$ -byte (512-MByte) kernel virtual space located at addresses 0x8000\_0000 - 0x9FFF\_FFFF. References to kseg0 are unmapped; the physical address selected is defined by subtracting 0x8000\_0000 from the virtual address. The K0 field of the *Config* register controls cacheability.

### **Kernel Mode, Kernel Space 1 (kseg1)**

In Kernel mode, when the most-significant three bits of the 32-bit virtual address are 3'b101, kseg1 virtual address space is selected. kseg1 is the  $2^{29}$ -byte (512-MByte) kernel virtual space located at addresses 0xA000\_0000 - 0xBFFF\_FFFF. References to kseg1 are unmapped; the physical address selected is defined by subtracting 0xA000\_0000 from the virtual address. Caches are disabled for accesses to these addresses, and physical memory (or memory-mapped I/O device registers) are accessed directly.

### **Kernel Mode, Kernel/Supervisor Space 2 (ksseg/kseg2)**

In Kernel mode, when  $KSU = 2'b00$ ,  $ERL = 1$ , or  $EXL = 1$  in the *Status* register, and  $DM = 0$  in the *Debug* register, and the most-significant three bits of the 32-bit virtual address are 3'b110, 32-bit kseg2 virtual address space is selected. With the FM MMU, this  $2^{29}$ -byte (512-MByte) kernel virtual space is located at physical addresses 0xC000\_0000 - 0xDFFF\_FFFF. Otherwise, this space is mapped through the TLB.

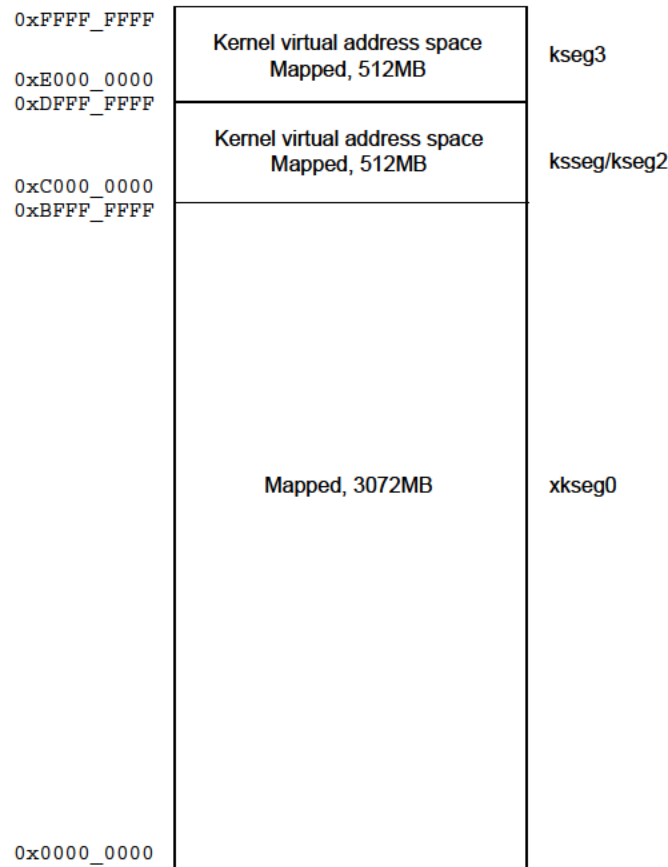
### **Kernel Mode, Kernel Space 3 (kseg3)**

In Kernel mode, when the most-significant three bits of the 32-bit virtual address are 3'b111, the kseg3 virtual address space is selected. With the FM MMU, this  $2^{29}$ -byte (512-MByte) kernel virtual space is located at physical addresses 0xE000\_0000 - 0xFFFF\_FFFF. Otherwise, this space is mapped through the TLB.

#### **3.13.4.2 Kernel Mode — EVA Configuration**

In the kernel mode - EVA configuration, the size of each kernel virtual address segment depends on the programming of the segment control registers. [Figure 3.37](#) shows an example of a 3GB xkseg0 segment (suseg), a 512MB kernel supervisor segment (ksseg), and a 512MB kernel segment 3 (kseg3).

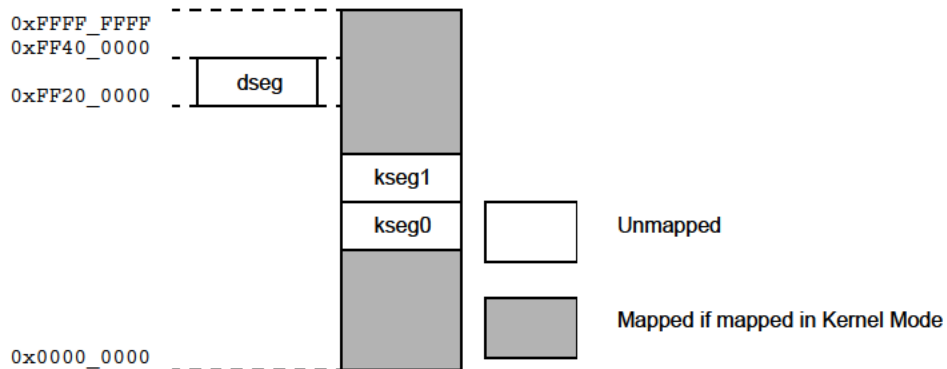
**Figure 3.39 Kernel Mode Virtual Address Space — EVA Mode, 3GB Example**



### 3.13.5 Debug Mode

Except for kseg3, debug-mode address space is identical to kernel-mode address space with respect to mapped and unmapped areas. In kseg3, a debug segment (dseg) coexists in the virtual address range `0xFF20_0000` to `0xFF3F_FFFF`. The layout is shown in Figure 3.40.

**Figure 3.40 Debug Mode Virtual Address Space**



dseg is subdivided into the dmseg segment at 0xFF20\_0000 to 0xFF2F\_FFFF, which is used when the debug probe services the memory segment, and the drseg segment at 0xFF30\_0000 to 0xFF3F\_FFFF, which is used when memory-mapped debug registers are accessed. The subdivision and attributes of the segments are shown in [Table 3.17](#).

Accesses to memory that would normally cause an exception in kernel mode cause the CPU to re-enter debug mode via a debug-mode exception. This includes accesses usually causing a TLB exception, with the result that such accesses are not handled by the usual memory-management routines.

The unmapped kseg0 and kseg1 segments from kernel-mode address space are available in debug mode, which allows the debug handler to be executed from uncached, unmapped memory.

**Table 3.17 Physical Address and Cache Attributes for dseg, dmseg, and drseg**

Segment Name	Sub-Segment Name	Virtual Address	Generates Physical Address	Cache Attribute
dseg	dmseg	0xFF20_0000 through 0xFF2F_FFFF	dmseg maps to addresses 0x0_0000 - 0xF_FFFF in EJTAG probe memory space.	Uncached
	drseg	0xFF30_0000 through 0xFF3F_FFFF	drseg maps to the breakpoint registers 0x0_0000 - 0xF_FFFF	

### 3.13.5.1 Debug Mode, Register (drseg)

The behavior of CPU access to the drseg address range at 0xFF30\_0000 to 0xFF3F\_FFFF is determined as shown in [Table 3.18](#)

**Table 3.18 CPU Access to drseg**

Transaction	LSNM Bit in Debug Register	Access
Load / Store	1	Kernel mode address space (kseg3)
Fetch	Don't care	drseg, see comments below
Load / Store	0	

Debug software is expected to read the *Debug Control* register (*DCR*) to determine which other memory-mapped registers exist in drseg. The value returned in response to a read of any unimplemented memory-mapped register is unpredictable, and writes are ignored to any unimplemented register in drseg. For more information about the *DCR*, refer to [Chapter 16, “EJTAG Debug Support”](#).

The allowed access size is limited for the drseg. Only word-size transactions are allowed. Operation of the processor is undefined for other transaction sizes.

### 3.13.5.2 Debug Mode, Memory (dmseg)

The conditions for CPU accesses to the dmseg address range (0xFF20\_0000 to 0xFF2F\_FFFF) are shown in [Table 3.19](#).

**Table 3.19 CPU Access to dmseg**

Transaction	ProbEn Bit in DCR Register <sup>1</sup>	LSNM Bit in Debug Register	Access
Load / Store	Don't care	1	Kernel mode address space (kseg3)

**Table 3.19 CPU Access to dmseg**

Transaction	ProbEn Bit in DCR Register <sup>1</sup>	LSNM Bit in Debug Register	Access
Fetch	1	Don't care	dmseg
Load / Store	1	0	dmseg
Fetch	0	Don't care	See comments below
Load / Store	0	0	See comments below

1. The NoDCR bit in the CP0 Debug register indicates if the dmseg and drseg address spaces and associated DCR register exists in memory mapped space. The NoDCR bit must be cleared, this DCR register exists. If the bit is set, the register does not exist.

An attempt to access dmseg when the ProbEn bit in the DCR register is 0 should not happen, because debug software is expected to check the state of the ProbEn bit in DCR register before attempting to reference dmseg. If such a reference does occur, the reference hangs until it is satisfied by the probe. The probe must not assume that there will never be a reference to dmseg when the ProbEn bit in the DCR register is 0, because there is an inherent race between the debug software sampling the ProbEn bit as 1, and the probe clearing it to 0.

## 3.14 TLB Instructions

Table 3.20 lists the TLB-related instructions implemented in the interAptiv core. .

**Table 3.20 TLB Instructions**

Mnemonic	Instruction	Description
TLBP	Translation Lookaside Buffer Probe	Used to determine whether a particular address was successfully translated. When a TLBP instruction is executed and fails to find a match for the specified virtual address, hardware sets bit 31 of the <i>Index</i> register.
TLBR	Translation Lookaside Buffer Read	
TLBWI	Translation Lookaside Buffer Write Index	TLB write extended to support invalidation of individual TLB entries.
TLBWR	Translation Lookaside Buffer Write Random	
TLBINV	Translation Lookaside Buffer Invalidate	Added to support set level invalidation of TLB entries.
TLBINVF	Translation Lookaside Buffer Invalidate Flush	Added to support TLB flush based invalidation of TLB entries.

Refer to the Instructions chapter for more information on the TLB instructions.

# Memory Protection Unit

The interAptiv core includes an optional Memory Protection Unit (MPU) for applications that do not require the functionality of a full MMU. The MPU controls whether read, write, or execute access is permitted for a given address and causes an exception if a unauthorized access is attempted. For more information on the MMU, refer to the Memory Management Unit chapter.

This chapter contains the following sections:

- [Section 4.1, "Overview" on page 241](#)
- [Section 4.2, "MPU Memory Mapping" on page 242](#)
- [Section 4.3, "MPU Register Mapping" on page 250](#)
- [Section 4.4, "MPU Registers" on page 251](#)

## 4.1 Overview

The MPU provides for finer-grain protection of memory pages relative to an MMU and also eliminates the address translation mechanism for space-sensitive applications. The MPU provides memory access control at both the segment and the region level. The concept of segments versus regions is described in the following section.

The MPU does not place restrictions on access permissions. All modes (Kernel, User, and Debug) are subject to the same permissions. In addition, all threads are subject to the same permissions. Note that Supervisor mode is not supported when an MPU is implemented.

Even though there are no restrictions on access permissions, Debug memory spaces are still restricted to use only in Debug mode. Additionally, if MPU Segmentation is disabled, User mode access to legacy Kernel segments still causes an Address Error.

The MPU can be enabled or disabled by selecting the option at build time or under software control. Disabling the MPU forces a fixed mapping translation (FMT) address translation mechanism using CCA values from the CP0 Config register. This can be used to leverage legacy software and infrastructure. Refer to the K23 and KU fields of the CP0 Config register (Register 16, Select 0) for more information.

## 4.2 MPU Memory Mapping

The MPU memory map can be configured using the concept of default segments and regions as described in the following subsections.

### 4.2.1 Default Segment Control

The interAptiv core MPU divides the 4 GByte memory space into a series of sixteen 256 MB segments. These fixed size segments can efficiently set the default attributes for each memory address. The default segments and the register fields that control them are shown in [Figure 4.1](#) below. Each segment contains the following programmable elements.

- Cache Coherency Attributes (CCA): Indicates the coherency attributes for the entire 256M segment.
- Read-Inhibit (RI): Determines if data reads are allowed. If the RI bit of the corresponding segment control register is set and a data read is attempted anywhere within that 256 MByte segment, an exception occurs.
- Write-Inhibit (WI): Determines if data writes are allowed. If the WI bit of the corresponding segment control register is set and a data write is attempted anywhere within that 256 MByte segment, an exception occurs.
- Execute-Inhibit (XI): Determines if code fetches are allowed. If the XI bit of the corresponding segment control register is set and a code fetch is attempted anywhere within that 256 MByte segment, an exception occurs.

For example, a 256MB segment could be allocated for I/O devices and should allow data reads and writes but not fetches. In this case, the RI and WI bits for that segment would be programmed with a value of 0, allowing read or write operations to occur. The XI bit would be programmed with a value of 1, indicating that code fetches are not allowed from that memory segment.

Conversely, for a segment configured for only code fetch accesses, the RI and WI bits would be programmed with a value of 1, indicating data reads and writes are not allowed, and the XI bit would be programmed with a value of 0, indicating that code fetches are allowed.

Typically, a combination of default segment mapping and region mapping will be used. The default attributes from the segment mapping will be overridden if the address is programmed as part of a region. Any memory space not specifically defined as a region using the Region Control register uses the default segment mapping.

**Figure 4.1 Default Memory Segments**

4 GByte Address Space		
SegmentControl_3 Register Segment 15, bits 31:24	Segment 15 (256 MByte) CCA, RI, WI, XI	FFFF_FFFF F000_0000
SegmentControl_3 Register Segment 14, bits 23:16	Segment 14 (256 MByte) CCA, RI, WI, XI	FFFF_FFFF E000_0000
SegmentControl_3 Register Segment 13, bits 15:8	Segment 13 (256 MByte) CCA, RI, WI, XI	DFFF_FFFF D000_0000
SegmentControl_3 Register Segment 12, bits 7:0	Segment 12 (256 MByte) CCA, RI, WI, XI	CFFF_FFFF C000_0000
SegmentControl_2 Register Segment 11, bits 31:24	Segment 11 (256 MByte) CCA, RI, WI, XI	BFFF_FFFF B000_0000
SegmentControl_2 Register Segment 10, bits 23:16	Segment 10 (256 MByte) CCA, RI, WI, XI	AFFF_FFFF A000_0000
SegmentControl_2 Register Segment 9, bits 15:8	Segment 9 (256 MByte) CCA, RI, WI, XI	9FFF_FFFF 9000_0000
SegmentControl_2 Register Segment 8, bit 7:0	Segment 8 (256 MByte) CA, RI, WI, XI	8FFF_FFFF 8000_0000
SegmentControl_1 Register Segment 7, bits 31:24	Segment 7 (256 MByte) CCA, RI, WI, XI	7FFF_FFFF 7000_0000
SegmentControl_1 Register Segment 6, bits 23:16	Segment 6 (256 MByte) CCA, RI, WI, XI	6FFF_FFFF 6000_0000
SegmentControl_1 Register Segment 5, bits 15:8	Segment 5 (256 MByte) CCA, RI, WI, XI	5FFF_FFFF 5000_0000
SegmentControl_1 Register Segment 4, bits 7:0	Segment 4 (256 MByte) CCA, RI, WI, XI	4FFF_FFFF 4000_0000
SegmentControl_0 Register Segment 3, bits 31:24	Segment 3 (256 MByte) CCA, RI, WI, XI	3FFF_FFFF 3000_0000
SegmentControl_0 Register Segment 2, bits 23:16	Segment 2 (256 MByte) CCA, RI, WI, XI	2FFF_FFFF 2000_0000
SegmentControl_0 Register Segment 1, bits 15:8	Segment 1 (256 MByte) CCA, RI, WI, XI	1FFF_FFFF 1000_0000
SegmentControl_0 Register Segment 0, bit 7:0	Segment 0 (256 MByte) CCA, RI, WI, XI	0FFF_FFFF 0000_0000

Each segment is set to the default cache coherency attributes (CCA) and associated permissions if no matching region is found. The SegmentControl 0 - 3 registers are used to program the segment attributes. Additionally, the reset values can be configured at build time. Refer to [Section 4.4.3, "Memory Protection Unit Segment Control Registers \(SegmentControl\\_N\) — Byte Offsets 0x10/14/18/1C"](#) for more information.

## 4.2.2 Region Control

For memory systems that require more control and granularity versus the default segment approach described in the previous subsection, a region mapping can be used. Region mapping allows for more dynamic protection, variable stack sizes, etc. and offers more localized protection.

While the default segment mapping has a fixed number of segments (16), each with a fixed size (256 MBytes) as described above, the MPU allows for a variable number of regions: 8/12/16/20/24/28/32. In addition, the size of each region is configurable from 32 bytes to 4 GBytes. The cache coherency attributes (CCA) and the permissions are set per region.

### 4.2.2.1 Super-Region and Subregion Definition

In defining a region, two related concepts are useful to understand; super-region and subregion.

A region must be contained within a super-region. The super-region is constrained to be a power of 2 size and must be naturally aligned to the size (ie. a 1GB super-region must start at a 1GB boundary). A super-region is divided into 16 equally sized subregions. The overall size of the super-region varies depending on the size of each subregion.

A region is the set of enabled subregions. It is programmed by setting the following parameters:

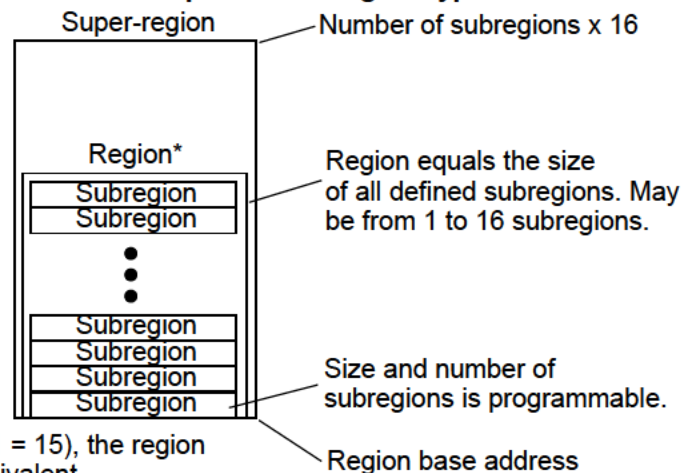
- Subregion size
- Base address, which must be aligned to the subregion size
- Number of subregions to be enabled

For example, if the subregion is defined as 16 KB and there are 5 subregions defined, then the region size is 80 KB (16 KB x 5). However, the super-region size is 256 KB (16 KB/subregion x 16 subregions). The subregion size is programmable using the *RegionCtrl\_N.SIZE* field of the corresponding Region Control register. Note that if all 16 subregions are used, then the region and the super-region would be the same size. A maximum of 32 regions are supported for the MPU. Refer to [Section 4.4.4, "MPU Region Address and Control Registers"](#) for more information.

The combination of these parameters give flexibility to the region definition by allowing non-power of 2 sizes and alignments. Meanwhile, the constraints of the super-region allow hardware optimization, reducing the silicon area required for programming and address matching.

[Figure 4.2](#) shows the relationship between the super-region, the region, and the subregion.

**Figure 4.2 Relationship Between Region Types**



\*If all 16 subregions are used (Count-1 = 15), the region size and the super-region size are equivalent.



### 4.2.2.2 Region Programming Examples

This section describes how to configure the memory regions using the Region Address and Control registers. As described above, the region is defined by its number and size of subregions. The total aggregate size of the subregions, and hence the region, may or may not be aligned on a power of two as is required by the super-region.

The following examples show regions defined as 80 KB and 128 KB.

#### **Example 1: 80 KB Region Aligned to a Power of Two Address Boundary**


In this example the region size is 80 KB and the super-region size is 256 KB. The region is aligned to a power of two address boundary (256 MB).

```
BaseAddr = 0x1000_0000 (256 MB)
Size = 14 (16 KB)
Count-1 = 4
BoundAddr = 0x1001_3FFF Upper address boundary of the region
```

- The *BaseAddr* variable, 0x1000\_0000 in this example, is the base address of the region. This value is programmed in to the *RegionAddr\_N.BaseAddress* field. In this example the base address for the region falls on a power of two boundary.
- The *Size* variable is 14, which equates to 16 KB per subregion. This value is programmed into the *RegionCtrl\_N.Size* field.
- The *Count-1* variable is 4, indicating a total of five (0 - 4) 16 KB subregions, for a total of 80 KB. This value is programmed into the *RegionCtrl\_N.Count* field.

The *BoundAddr* variable, 0x1001\_3FFF in this example, is the upper address bound of the region. This value is shown for completeness and is not programmed into any register. This concept is shown in [Figure 4.3](#).

**Figure 4.3 Region Programming — Example 1**



The following assembly language code shows how to program the Region Control registers for Example 1 above.

```
la r16, MPUBase          # Address of MPU programming registers
la r17, 0x1000_0000      # Base address of region
sw r17, (4*0x8 + 0x20)(r16) # Write to base address 4
li r17, (1 << 15) \ # MPU_En
    | (14 << 10) \ # MPU_Size: 16KB
    | (4 << 6) \ # MPU_Count: 5 subregions
    | (1 << 5) \ # ReadInhibit
    | (1 << 4) \ # WriteInhibit
    | 5 # CCA 5
sw r17, (4*0x8+0x24)(r16) # Write to control 4
```

### **Example 2: 80 KB Region Not Aligned to a Power of Two Address Boundary**

In this example the subregion size and number of subregions is the same as in example 1 above, but the base address for the region is not on a power of two address boundary.

```
BaseAddr = 0x1000_C000
Size = 14 (16K)
Count-1 = 4
BoundAddr = 0x1001_FFFF
```

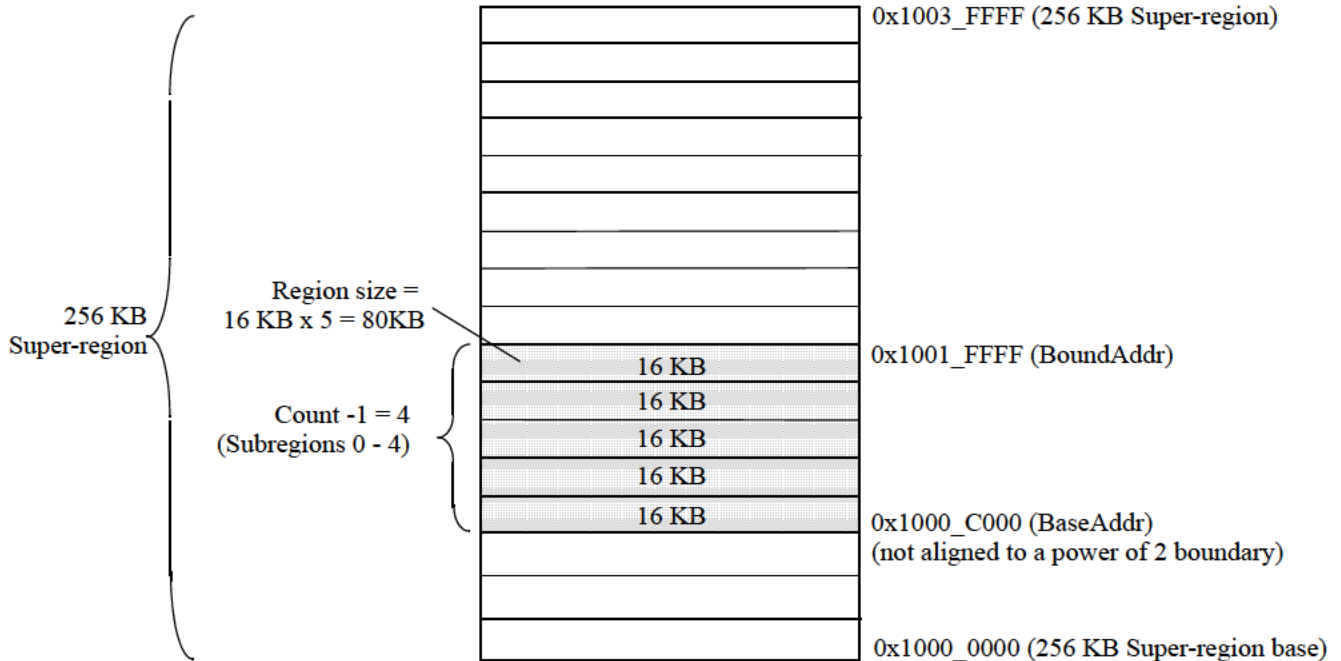
- The *BaseAddr* variable, 0x1000\_C000 in this example, is the base address of the region. This value is programmed into the *RegionAddr\_N.BaseAddress* field.
- The *Size* variable is 14, which equates to 16 KB per subregion. This value is programmed into the *RegionCtrl\_N.Size* field.
- The *Count-1* variable is 4, indicating a total of five (0 - 4) 16 KB subregions, for a total of 80 KB. This value is programmed into the *RegionCtrl\_N.Count* field.

The *BoundAddr* variable, 0x1001\_FFFF in this example, is the upper address bound of the region. This value is shown for completeness and is not programmed into any register. As with example 1, the region shown in this example is encompassed with a 256 KB super-region.

In this example, the base address of the region does not fall on a power of two, but the super-region address does.

This concept is shown in [Figure 4.4](#).

**Figure 4.4 Region Programming — Example 2**



### Example 3: 128 KB Region

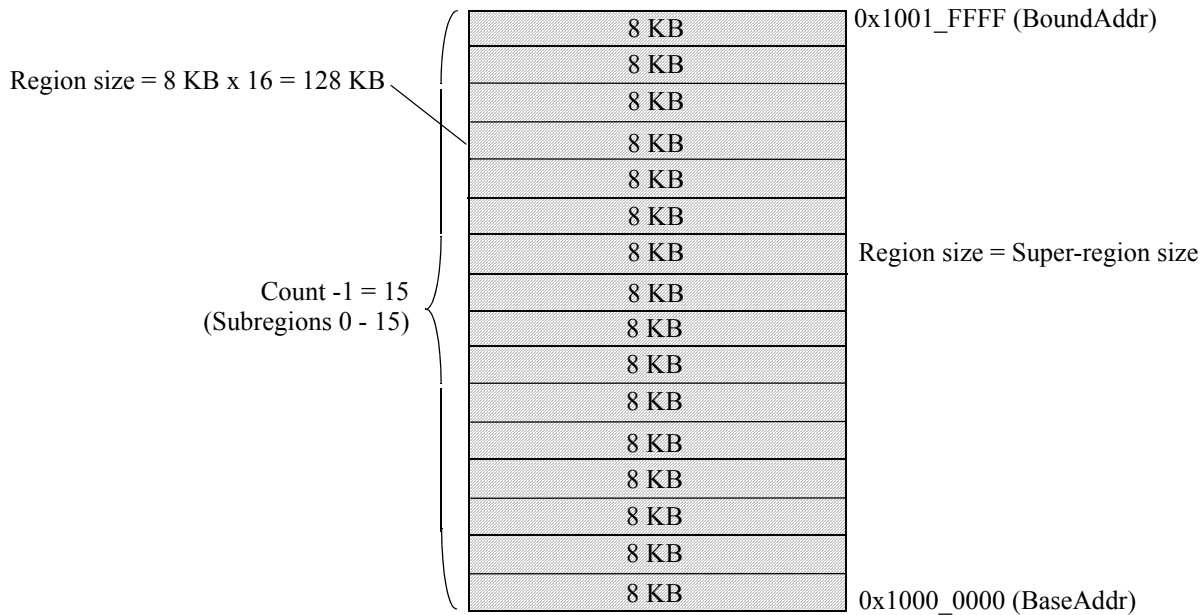
In this example the subregion size and the region size are both 128 KB.

```
BaseAddr = 0x1000_0000
Size = 13 (8K)
Count-1 = 15
BoundAddr = 0x1001_FFFF
```

- The *BaseAddr* variable, 0x1000\_0000 in this example, is the base address of the region. This value is programmed into the *RegionAddr\_N.BaseAddress* field.
- The *Size* variable is 13, which equates to 8 KB per subregion. This value is programmed into the *RegionCtrl\_N.Size* field.
- The *Count-1* variable is 15, indicating a total of sixteen (0 - 15) 8 KB subregions, for a total of 128 KB. This value is programmed into the *RegionCtrl\_N.Count* field.

The *BoundAddr* variable, 0x1001\_FFFF in this example, is the upper address bound of the region. This value is shown for completeness and is not programmed into any register. In this example the region size and the super-region size are one in the same. This concept is shown in [Figure 4.5](#).

**Figure 4.5 Region Programming — Example 3**



### **Subregion Sizing Options**

Regions that are power of two sized and aligned can be equivalently specified using different sizes and number of subregions. The hardware behaves identically for all synonyms and software can use whichever one makes sense. Using the encoding with the fewest subregions may enhance portability as it could be used on a core with fewer subregions. As an example, the 3rd example above (128KB region) could have been encoded as:

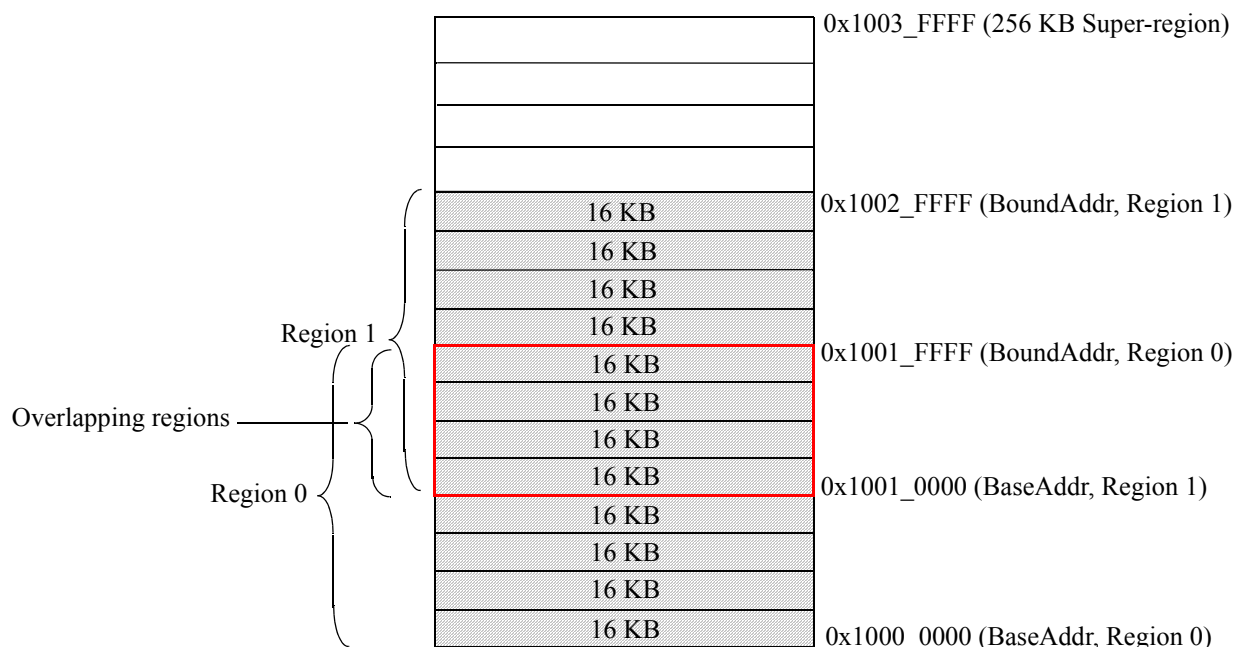
- Size=17, Count-1 = 0 (1x128K)
- Size=16, Count-1 = 1 (2x64K)
- Size=15, Count-1 = 3 (4x32K)
- Size=14, Count-1 = 7 (8x16K)
- Size=13, Count-1 = 15 (16x8K)

### **4.2.2.3 Overlapping Regions**

Regions are implicitly prioritized by entry number. Overlapping regions are allowed, with the permissions being assigned from the highest-numbered entry. For example, if region 3 overlaps region 2, the permissions are assigned to the overlap region as defined in region 3, since it is the highest numbered region. If an access does not match any region, the default segment programming or FMT mapping (if the MPU is disabled at build time) will be used.

The following example shows how two 128 KByte regions are overlapped.

**Figure 4.6 Overlapping Regions**



In this example regions 0 and 1 overlap. Both regions are 128 KBytes. For the 64 KByte section shown in red that overlaps, the permissions and attributes for the overlap area are assigned to region 1, since it is the highest numbered region.

#### 4.2.2.4 Subregion Programming

There can be between 1 and 16 subregions per region. The number of subregions is programmed using the 4-bit *RegionCtrl\_N.COUNT-1* field of the corresponding Region Control register. The COUNT field is encoded as follows:

**Table 4.1 COUNT Field Encoding**

COUNT Field Encoding	Number of Subregions	COUNT Field Encoding	Number of Subregions
0x0	1	0x8	9
0x1	2	0x9	10
0x2	3	0xA	11
0x3	4	0xB	12
0x4	5	0xC	13
0x5	6	0xD	14
0x6	7	0xE	15
0x7	8	0xF	16

The size of each subregion can be between 32 bytes and 256 MBytes. The size of each subregion is programmed using the 5-bit *RegionCtrl\_N.SIZE* field of the corresponding Region Control register. The SIZE field is encoded as follows:

**Table 4.2 SIZE Field Encoding**

SIZE Field Encoding	Size of Subregion	SIZE Field Encoding	Size of Subregion	SIZE Field Encoding	Size of Subregion
0x05	32 bytes	0x0D	8 Kbytes	0x15	2 Mbytes
0x06	64 bytes	0x0E	16 Kbytes	0x16	4 Mbytes
0x07	128 bytes	0x0F	32 Kbytes	0x17	8 Mbytes
0x08	256 bytes	0x10	64 Kbytes	0x18	16 Mbytes
0x09	512 bytes	0x11	128 Kbytes	0x19	32 Mbytes
0x0A	1 Kbytes	0x12	256 Kbytes	0x1A	64 Mbytes
0x0B	2 Kbytes	0x13	512 Kbytes	0x1B	128 Mbytes
0x0C	4 Kbytes	0x14	1 Mbytes	0x1C	256 Mbytes

Refer to [Section 4.4.4, "MPU Region Address and Control Registers"](#) for more information on the COUNT and SIZE fields.

### **Writing to the RegionCtrl Registers**

The MPU is accessed via a set of memory mapped registers. The location and accessibility of these registers can be controlled by software via the CDMMBase CP0 register.

The above tables show the encoding for the SIZE and COUNT fields of the Region Control registers. The following assembly language example performs a write to the Region Address (RegionAddr\_N) and Control (RegionCtrl\_N) registers. Refer to [Section 4.2.2.2, Example 1](#) for more assembly code examples.

```

la r16, MPUBase           # Address of MPU programming register
la r17, <base address>    # Base address of region
sw r17, (N*0x8 + 0x20)(r16) # Write to base address N
li r17, <control val>     # Setup control register
sw r17, (N*0x8+0x24)(r16) # Write to control N

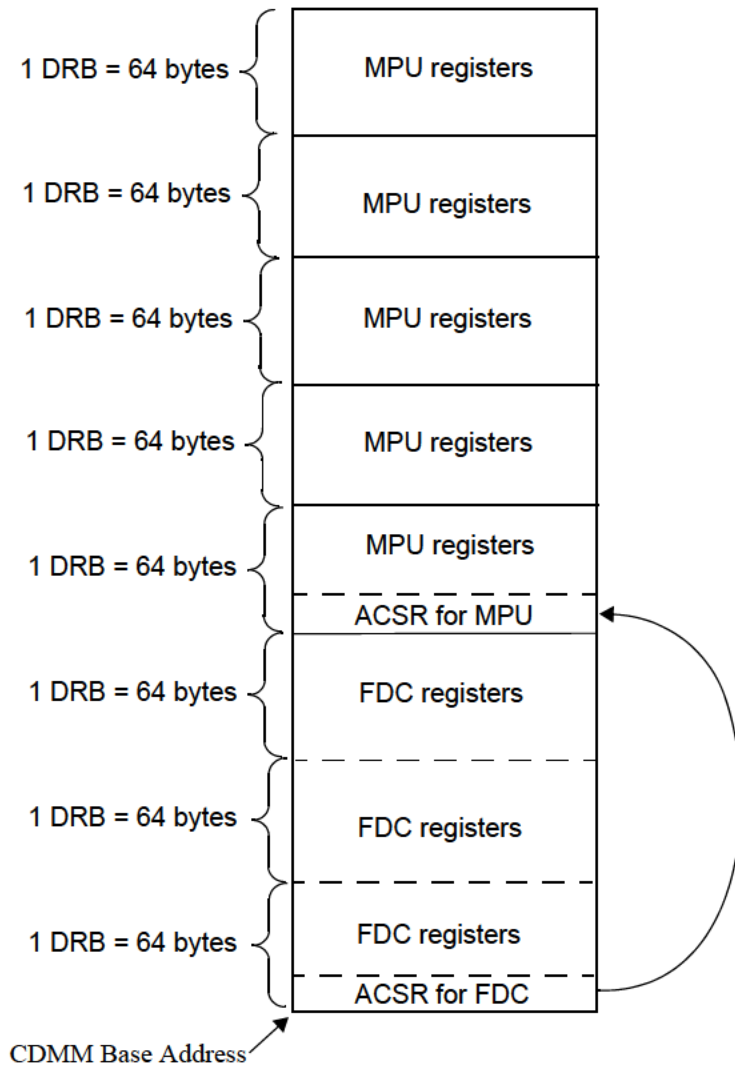
```

## **4.3 MPU Register Mapping**

In the interAptiv core MPU, the segment and region control registers are memory mapped registers that are accessed via the Common Device Memory Map (CDMM) mechanism. These registers are accessed via load and store instructions.

The base address of the location of these registers in memory is programmed into the CP0 CDMMBase register located at CP0 register 15, Select 2. The base address is written into bits 27:11 of this register. This memory space is further divided into a series of 64-byte device register blocks (DRB). The number of blocks is determined by the CDMMSize field in bits 8:0 of the CP0 CDMMBase register. In the interAptiv core, the first three blocks starting at the base address are dedicated to the Fast Debug Channel (FDC). The next five blocks are dedicated to the MPU. This concept is shown in [Figure 4.7](#).

**Figure 4.7 Location of the MPU Registers in the interAptiv Core**




## 4.4 MPU Registers

The following subsections describe the MPU registers. As described in the figure above, the CDMM provides a software discoverable chain of memory mapped devices. The following register addresses are described relative to the device base.

### 4.4.1 Memory Protection Unit Status Register (MPU\_ACSR) — Device Byte Offset 0

The MPU\_ACSR register provides MPU status information such as the device type, size, and revision. It also controls access to the MPU registers by User and Supervisor modes.

**Figure 4.8 MPU\_ACSR Register Format**




**Table 4.3 MPU\_ACSR Register Field Descriptions**

Name	Bits	Description	Read/Write	Reset State
DevType	31:24	Device Type: Set to 0x02 for this MPU.	R	0x02
DevSize	23:16	Device Size: Number of additional 64-byte blocks allocated to this device. Address space is allocated for the maximum 32-entry configuration.	R	0x04
DevRev	15:12	Device Revision.	R	0x0
0	11:4	Reserved. Write as zero.	R	0
UW	3	Allow User mode writes.	R/W	0
UR	2	Allow User mode reads.	R/W	0
SW	1	Allow Supervisor mode writes.	R/W	0
SR	0	Allow Supervisor mode reads.	R/W	0

#### 4.4.2 Memory Protection Unit Configuration Register (MPU\_Config) — Device Byte Offset 0x8

The MPU\_Config register provides MPU configuration information such as the MPU segment control, protection exceptions, and the number of regions.

**Figure 4.9 MPU\_Config Register Format**



**Table 4.4 MPU\_Config Register Field Descriptions**

Name	Bits	Description	Read/Write	Reset State
En	31	Enable MPU segment control. Use FMT mode if not enabled.	R/W	Preset
0	30:20	Reserved. Write as zero.	R	0
ExcR	19	Last Protection exception caused by a load to a RI address.	R	0
ExcW	18	Last Protection exception caused by a load to a WI address.	R	0
ExcX	17	Last Protection exception caused by a fetch to a XI address.	R	0
Exc_Reg_Match	16	Last Protection exception hit in region. This bit is encoded as follows: 0: No Match, exception due to segment register 1: Match, exception due to region register	R	0
0	15:13	Reserved. Write as zero.	R	0
Exc_Reg_Num	12:8	Region number for last protection exception. Undefined if no region match.	R	0
0	7:5	Reserved. Write as zero.	R	0
Num_Regions	4:0	Total number of regions - 1. A maximum of 32 regions are supported.	R	Preset



### 4.4.3 Memory Protection Unit Segment Control Registers (SegmentControl\_N) — Byte Offsets 0x10/14/18/1C

There are four segment control registers that set the default permissions for each of the sixteen 256MB segments when MPU segmentation is enabled. Each segment is controlled by an 8-bit field in the one of the following registers.


**Table 4.5 MPU Segment Control Register Map**

Offset	Segment Control Register Fields							
	31	24	23	16	15	8	7	0
0x10	Segment 3		Segment 2		Segment 1		Segment 0	
0x14	Segment 7		Segment 6		Segment 5		Segment 4	
0x18	Segment 11		Segment 10		Segment 9		Segment 8	
0x1C	Segment 15		Segment 14		Segment 13		Segment 12	

Note that the segment numbering scheme above corresponds to the state of the upper 4 bits of address for each segment. For example, segment 9 corresponds to a starting address of 9000\_0000. Segment 15 corresponds to a starting address of F000\_0000.

Each of the four Segment Control registers have identical bit assignments, except each field represents a different segment.

**Figure 4.10 Segment Control Register Format**



**Table 4.6 Segment Control Register Field Descriptions**

Name	Bits	Description	Read/Write	Reset State
0	31:30	Reserved. Write as zero.	R	0
RI	29	Read inhibit. Trigger exception on data read.	R/W	Preset
WI	28	Write inhibit. Trigger exception on data write.	R/W	Preset
XI	27	Execute inhibit. Trigger exception on instruction fetch.	R/W	Preset
CCA	26:24	Cache Coherency Attributes. Refer to <a href="#">Table 2.19</a> in Chapter 2 for more information.	R/W	Preset
Various	23:16	Same encoding as bits 31:24, except for different segments as shown in <a href="#">Table 4.5</a> above.	R/W	Preset
Various	15:8	Same encoding as bits 31:24, except for different segments as shown in <a href="#">Table 4.5</a> above.	R/W	Preset
Various	7:0	Same encoding as bits 31:24, except for different segments as shown in <a href="#">Table 4.5</a> above.	R/W	Preset

## 4.4.4 MPU Region Address and Control Registers


There are 64 region control registers (two per region) that control up to 32 regions. For each region there is a base address register and a control register that specifies the size and attributes of that region. The following table shows the location of each Region register relative to the Device base address.

**Table 4.7 MPU Region Address and Control Register Map**

Offset	Register	Offset	Register	Offset	Register	Offset	Register
0x20	Region 0 Address	0x60	Region 8 Address	0xA0	Region 16 Address	0xE0	Region 24 Address
0x24	Region 0 Control	0x64	Region 8 Control	0xA4	Region 16 Control	0xE4	Region 24 Control
0x28	Region 1 Address	0x68	Region 9 Address	0xA8	Region 17 Address	0xE8	Region 25 Address
0x2C	Region 1 Control	0x6C	Region 9 Control	0xAC	Region 17 Control	0xEC	Region 25 Control
0x30	Region 2 Address	0x70	Region 10 Address	0xB0	Region 18 Address	0xF0	Region 26 Address
0x34	Region 2 Control	0x74	Region 10 Control	0xB4	Region 18 Control	0xF4	Region 26 Control
0x38	Region 3 Address	0x78	Region 11 Address	0xB8	Region 19 Address	0xF8	Region 27 Address
0x3C	Region 3 Control	0x7C	Region 11 Control	0xBC	Region 19 Control	0xFC	Region 27 Control
0x40	Region 4 Address	0x80	Region 12 Address	0xC0	Region 20 Address	0x100	Region 28 Address
0x44	Region 4 Control	0x84	Region 12 Control	0xC4	Region 20 Control	0x104	Region 28 Control
0x48	Region 5 Address	0x88	Region 13 Address	0xC8	Region 21 Address	0x108	Region 29 Address
0x4C	Region 5 Control	0x8C	Region 13 Control	0xCC	Region 21 Control	0x10C	Region 29 Control
0x50	Region 6 Address	0x90	Region 14 Address	0xD0	Region 22 Address	0x110	Region 30 Address
0x54	Region 6 Control	0x94	Region 14 Control	0xD4	Region 22 Control	0x114	Region 30 Control
0x58	Region 7 Address	0x98	Region 15 Address	0xD8	Region 23 Address	0x118	Region 31 Address
0x5C	Region 7 Control	0x9C	Region 15 Control	0xDC	Region 23 Control	0x11C	Region 31 Control

Each of the 32 Region control and 32 Region address registers have identical bit assignments, except each register represents a different region.


**Figure 4.11 Region Base Address Register Format (RegionAddr\_N)**



**Table 4.8 Region Base Address Register Field Descriptions**

Name	Bits	Description	Read/Write	Reset State
Base_Address	31:5	Base address of the range. Must be aligned to the subregion size.	R/W	Undefined
0	4:0	The lower five bits of the base address are always zero.	R	0

**Figure 4.12 Region Control Register Format (RegionCtrl\_N)**



**Table 4.9 Region Control Register Field Descriptions**

<b>Name</b>	<b>Bits</b>	<b>Description</b>	<b>Read/Write</b>	<b>Reset State</b>
0	31:16	Reserved. Write as zero.	R	Undefined
En	15	Enable bit for this region. This bit must be set in order to enable a given region.	R/W	0
Size	14:10	Size of a subregion in powers of 2. 5: 32B - minimum size 28: 256MB - maximum size (covers 4GB if all 16 subregions are enabled)	R/W	Undefined
Count	9:6	Number of enabled subregions -1 Illegal to extend beyond region size (possible when BaseAddress is not region aligned).	R/W	Undefined
RI	5	Read inhibit. Trigger exception on data read.	R/W	Undefined
WI	4	Write inhibit. Trigger exception on data write.	R/W	Undefined
XI	3	Execute inhibit. Trigger exception on instruction fetch.	R/W	Undefined
CCA	2:0	Cache Coherency Attributes	R/W	Undefined



# Caches

This chapter describes the caches present in an interAptiv core and contains the following sections:

- [Section 5.1 “Cache Configurations”](#)
- [Section 5.2 “L1 Instruction Cache”](#)
- [Section 5.3 “L1 Data Cache”](#)
- [Section 5.4 “L1 Instruction and Data Cache Software Testing”](#)
- [Section 5.5 “L2 Cache”](#)
- [Section 5.6 “The CACHE Instruction”](#)
- [Section 5.7 “Write Back Buffer”](#)

## 5.1 Cache Configurations

The interAptiv core contains three caches; L1 instruction, L1 data, and shared L2. These caches are non-optional in the interAptiv architecture and are always present. The size of each cache can be configured as shown in [Table 5.1](#).

**Table 5.1 interAptiv Cache Configurations**

Attribute	L1 Instruction Cache	L1 Data Cache	L2 Cache
Size <sup>1</sup>	4 KB, 8 KB, 16 KB, 32 KB or 64 KB	4 KB, 8 KB, 16 KB, 32 KB or 64 KB	0 KB, 32 KB, 64 KB, 256 KB, 512 KB, 1 MB, 2 MB, 4 MB, or 8 MB
Line Size	32-byte	32-byte	32-byte or 64-byte
Number of Cache Sets	32, 64, 128, 256, or 512	32, 64, 128, 256, or 512	64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, or 32768
Associativity	4 way	4 way	8 way

1. For Linux-based applications, MIPS recommends an optimum L1 cache size of 64 KB, and a minimum L1 cache size of 32 KB.

The L1 instruction cache is attached to the Instruction Fetch Unit (IFU) via a 64-bit data path. The L1 data cache also contains a single 64-bit data path. The L2 cache is embedded within the Coherence Manager (CM2) and communicates with external memory via a configurable 64-bit, 128-bit, or 256-bit OCP interface.

For more information on the L1 instruction cache, refer to [Section 5.2 “L1 Instruction Cache”](#).

For more information on the L1 data cache, refer to [Section 5.3 “L1 Data Cache”](#).

For more information on the L2 cache, refer to [Section 5.5 “L2 Cache”](#).

## 5.1.1 Cacheability Attributes

The interAptiv core supports the following cacheability attributes:

- *Uncached (code #2)*: Addresses in a memory area indicated as uncached are not read from the cache. Stores to such addresses are written directly to main memory, without changing cache contents.
- *Non-coherent Writeback With Write Allocation (code #3)*: Loads and instruction fetches first search the cache, reading main memory only if the desired data does not reside in the cache. On data store operations, the cache is first searched to see if the target address is in the cache. If it is, the cache contents are updated, but main memory is not written. If the cache lookup misses on a store, main memory is read to bring the line into the cache and merge it with the new store data. Hence, the allocation policy on a cache miss is read- or write-allocate. Data stores will update the appropriate dirty bit in the ‘dirty’ array to indicate that the line contains modified data. When a line with dirty data is displaced from the cache, it is written back to memory.
- *Coherent Write-back With Write Allocation, Exclusive (code #4)*: This attribute is similar to code #5 described below, except that load misses bring data into the cache in the exclusive state rather than the shared state. This can be used if data is not shared and will eventually be written. This can reduce bus traffic, because the line does not have to be refetched in an exclusive state when a store is done.
- *Coherent Write-back With Write Allocation, Exclusive on Write (code #5)*: Use coherent data. Load misses will bring the data into the cache in a shared state. Multiple caches can contain data in the shared state. Stores will bring data into the cache in an exclusive state - no other caches can contain that same line. If a store hits on a shared line in the cache, the line will be invalidated and brought back into the cache in an exclusive state.
- *Uncached Accelerated (code #7)*: Uncached stores are gathered together for more efficient bus utilization.

## 5.2 L1 Instruction Cache

The L1 instruction cache contains three arrays: tag, data, and way-select. The L1 instruction cache is virtually indexed, since a virtual address is used to select the appropriate line within each of the three arrays. The caches are physically tagged, as the tag array contains a physical, not virtual, address.

The tag and data arrays hold 4 ways of information per set, corresponding to the 4-way set associativity of the cache. The way-select array holds information to choose the way to be filled, as well as dirty bits in the case of the data cache.

An instruction cache tag entry consists of the upper bits of the physical address bits, one valid bit for the line, and a lock bit. An instruction cache data entry contains four, 64-bit doublewords in the line, for a total of 32 bytes. All four words in the line are present or not in the data array together, hence the single valid bit stored with the tag. The number of upper address bits depends on the cache size as shown below.

- bits [31:12] for 64KB, 32KB, and 16KB caches
- bits[31:11] for 8KB cache
- bits[31:10] for 4KB cache

A way-select entry holds bits choosing the way to be replaced according to a Least Recently Used (LRU) algorithm. The LRU information applies to all the ways and there is one way-select entry for all the ways in the set. The instruction cache only supports reads, hence only LRU entries are stored in the instruction way-select array.

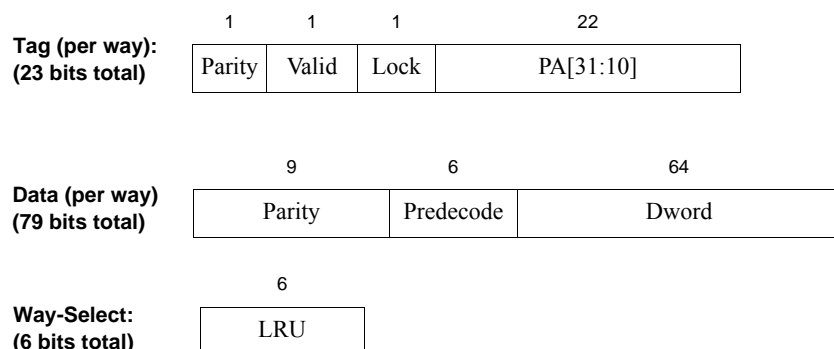
Table 5.2 shows the key characteristics of the L1 instruction cache. Figure 5.1 shows the format of an entry in the three arrays comprising the instruction cache: data, tag, and way-select.

**Table 5.2 L1 Instruction Cache Attributes**

Attribute	With Parity
Size <sup>1</sup>	0 KB, 4 KB, 8 KB, 16 KB, 32 KB or 64 KB
Line Size	32-byte
Number of Cache Sets	32, 64, 128, 256, or 512
Associativity	4-way
Replacement	LRU
Cache Locking	per line
<b>Data Array</b>	
Read Unit	70b x 4 (no parity) 79b x 4 (parity)
Write Unit	70b (no parity) 79b (parity)
<b>Tag Array</b>	
Read Unit	24b x 4 (no parity) 25b x 4 (parity)
Write Unit	24b (no parity) 25b (parity)
<b>Way-Select Array</b>	
Read Unit	6b
Write Unit	1-6b

1. For Linux based applications, MIPS recommends a 64 KB L1 instruction cache size, with a minimum size of 32 KB.

**Figure 5.1 L1 Instruction Cache Organization**



### 5.2.1 L1 Instruction Cache Virtual Aliasing

The instruction cache on the interAptiv core is virtually indexed and physically tagged. The lower bits of the virtual address are used to access the cache arrays and the physical address is used in the tags. Because the way size can be larger than the minimum TLB page size, there is a potential for virtual aliasing. This means that one physical address can exist in multiple indices within the cache, if it is accessed with different virtual addresses. Virtual aliasing comes into effect only for cache sizes that are larger than 16 KB.

In the interAptiv core, the **Config7<sub>IAV</sub>** bit is always set to indicate the existence of instruction cache virtual aliasing hardware. The core allows a physical address to reside at multiple indices if accessed with different virtual addresses. When an invalidate request is made due to the CACHE or SYNCI instructions, the core will serially check each possible alias location for the given physical address.

The hardware can be enabled and disabled using the **Config7<sub>IVAD</sub>** bit. When this bit is cleared, the hardware used to remove instruction cache virtual aliasing is enabled. In this case the virtual aliasing is managed in hardware. No software interaction is required. When the **Config7<sub>IVAD</sub>** bit is set, the virtual aliasing hardware is disabled. This can be done when software ensures that no cache aliases are possible, for example when using a minimum TLB page size of 16KB. In cases where the TLB page size is less than 16 KB, it is up to software to manage virtual aliasing within the instruction cache.

### 5.2.2 L1 Instruction Cache Precode Bits

In order for the fetch unit to quickly detect branches and jumps when executing code, the instruction cache array contains some additional precode bits. These bits indicate the type and location of branch or jump instructions within a 64b fetch bundle. These precode bits are not used when executing MIPS16e code.

### 5.2.3 L1 Instruction Cache Parity

The instruction cache contains 9 parity bits — one for each byte of the 64 bits of data plus 1 bit for the 6-bit precode. The tag array has 1 parity bits for each tag.



## 5.2.4 L1 Instruction Cache Replacement Policy

The L1 instruction cache replacement policy refers to how a way is chosen to hold an incoming cache line on a miss which will result in a cache fill. The replacement policy is least-recently used (LRU), but excluding any locked ways. The LRU bit(s) in the way-select array encode the order in which ways on that line have been accessed.

On a cache miss, the lock and LRU bits for the tag and way-select entries of the selected line may be used to determine the way which will be chosen.

The LRU field in the way select array is updated as follows:

- On a cache hit, the associated way is updated to be the most recently used. The order of the other ways relative to each another is unchanged.
- On a cache refill, the filled way is updated to be the most recently used.
- On CACHE instructions, the update of the LRU bits depends on the type of operation to be performed:
  - **Index (Writeback) Invalidate:** Least-recently used.
  - **Index Load Tag:** No update.
  - **Index Store Tag,  $WST = 0$ :** Most-recently used if valid bit is set in *TagLo* CP0 register. Least-recently used if valid bit is cleared in *TagLo* CP0 register.
  - **Index Store Tag,  $WST = 1$ :** Update the field with the contents of the *TagLo* CP0 register (refer to Table for the valid values of this field).
  - **Index Store Data:** No update.
  - **Hit Invalidate:** Least-recently used if a hit is generated, otherwise unchanged.
  - **Fill:** Most-recently used.
  - **Hit Writeback:** No update.
  - **Fetch and Lock:** For instruction cache, no update. For data cache, most-recently used.

If all ways are valid, then any locked ways will be excluded from consideration for replacement. For the unlocked ways, the LRU bits are used to identify the way which has been used least-recently, and that way is selected for replacement.

## 5.2.5 L1 Instruction Cache Line Locking

The interAptiv core does not support the locking of all 4 ways of either cache at a particular index. If all 4 ways of the cache at a given index are locked by either Fetch and Lock or Index Store Tag CACHE instructions, subsequent cache misses at that cache index will displace one of the locked lines.

Locking lines in the caches is somewhat counter to the idea of coherence. If a line is locked into a particular cache, it is expected that any processes utilizing that data will be locked to that processor and coherence is not needed. Based on this usage model, locking coherent lines into the cache is not recommended. However, should this occur, the CPU adheres to the following rules:

- SYNCI instructions are user-mode instructions. Since locking is a kernel mode feature (requires the CACHE instruction), SYNCI is not allowed to unlock cache lines. This applies to both local and globalized SYNCI instructions.

- Locking overrides coherence. Intervention requests from other CPUs and I/O devices that match on a locked line are treated as misses.
- Self-intervention requests for globalized CACHE instructions are allowed to affect a locked line. This is done primarily for handling lock and unlock requests for kseg0 addresses when kseg0 is being treated coherently.

## 5.2.6 L1 Instruction Cache Memory Coherence Issues

The interAptiv core supports software cache coherency in a multi-CPU cluster. Software must explicitly manage instruction cache coherence via the CACHE or SYNCI instructions to invalidate a line and pick up new data from L2 cache or main memory. These operations are globalized — if a the address used in the operation has a coherent CCA, the request will be sent to all instruction caches in the cluster.

In the interAptiv core, the hardware does not automatically keep the instruction caches coherent with the data caches. Doing so requires many additional cache lookups and would likely require the instruction cache tag array to be duplicated as well. For many types of code, this would be of small benefit, and the added area and power costs would not make sense. Further, the existing non-coherent cores from MIPS do not keep the I-Cache coherent with the D-Cache, so the code already exists for software I-Cache coherence where it is required. Globalized CACHE and SYNCI instructions ease the task of software I-Cache coherence. Existing, single-CPU routines that push dirty data out of the data cache and invalidate stale instruction cache lines using hit-type CACHE or SYNCI instructions can be globalized, and the coherence can be handled for all of the instruction caches in parallel.

## 5.2.7 Software I-Cache Coherence (JVM, Self-modifying Code)

The CPU does not support hardware I-Cache coherence, so code that modifies the instruction stream must clean up the instruction cache. This is equivalent to what is currently required on uniprocessor systems that also do not have a coherent I-Cache. The recommended SYNCI sequence shown below will also work for coherent addresses:

```
SW instn_address
SYNCI instn_address
SYNC
JR.HB instn_address
NOP
```

## 5.2.8 Instruction Cache Way Prediction


Instruction cache way prediction is a power saving feature that takes advantage of sequential read access to the instruction tag and data caches within a cache line and turns off unused datapath logic. I-cache way prediction aims to minimize number of read accesses and data size of each read access to the instruction caches when sequential read access to the same cacheline is encountered.

The instruction cache way prediction scheme is transparent to software. However, it is still possible to disable the way prediction by setting *Config7.ICWP*. Refer to [Section 5.2.9, "L1 Instruction Software Cache Management"](#).

## 5.2.9 L1 Instruction Software Cache Management

The L1 instruction cache is not fully “coherent” and requires OS intervention at times. The CACHE instruction is the building block of such OS interventions, and is required for correct handling of DMA data and for cache initialization. Historically, the CACHE instruction also had a role when writing instructions. Unless the programmer takes the appropriate action, those instructions may only be in the D-cache and would need them to be fetched through the I-cache at the appropriate time. Wherever possible, use the SYNCI instruction for this purpose, as described in [Section 5.2.12 “Cache Management When Writing Instructions - the “SYNCI” Instruction”](#).

A cache operation instruction is written `cache op, addr` where `addr` is just an address format, written as for a load/store instruction. Cache operations are privileged and can only run in kernel mode (SYNCI works in user mode, though).



**Figure 5.2 Fields in the Encoding of a CACHE Instruction**


The `op` field packs together a 5-bit field. The lower 2 bits of this field (17:16) select which cache to work on:

- 00 L1 I-cache
- 01 L1 D-cache
- 10 reserved
- 11 L2 cache

The upper 3-bits of the OP field encodes a command to be carried out on the line the instruction selects.

The CACHE instruction come in three varieties which differ in how they pick the cache entry (the “cache line”) they will work on:

- *Hit-type cache operation*: presents an address (just like a load/store), which is looked up in the cache. If this location is in the cache (it “hits”) the cache operation is carried out on the enclosing line. If this location is not in the cache, nothing happens.
- *Address-type cache operation*: presents an address of some memory data, which is processed just like a cached access - if the cache was previously invalid the data is fetched from memory.
- *Index-type cache operation*: as many low bits of the address as are required are used to select the byte within the cache line, then the cache line address inside one of the four cache ways, and then the way. The size of the cache (contained within the *Config1* register) to know exactly where the field boundaries are located. The address is used as follows:



Note that the MIPS32 specification allows the CPU designer to select whether to derive the index from the virtual or physical address. For index-type operations, MIPS recommends using a `kseg0` address, so that the virtual and physical address are the same. This also avoids a potential of cache aliasing.

## 5.2.10 L1 Instruction Cache CP0 Register Interface

The interAptiv core uses different CP0 registers for instruction cache operations.

**Table 5.3 Instruction Cache CP0 Register Interface**

CP0 Registers	CP0 number
<i>Config1</i>	16.1
<i>CacheErr</i>	27.0
<i>ITagLo</i>	28.0
<i>ITagHi</i>	29.0
<i>IDataLo</i>	28.1
<i>IDataHi</i>	29.1

### 5.2.10.1 Config1 Register (CP0 register 16, Select 1)

The *Config1*.*IS* field (bits 24:22) indicates the number of sets per way in the instruction cache. The interAptiv L1 instruction cache supports 256 sets per way, which is used to configure a 32 KB cache, or 512 sets per way, which is used to configure a 64 KB cache.

The *Config1*.*IL* field (bits 21:19) indicates the line size for the instruction cache. The interAptiv L1 instruction cache supports a fixed line size of 32 bytes as indicated by a default value of 4 for this field.

The *Config1*.*IA* field (bits 18:16) indicates the set associativity for the instruction cache. The interAptiv L1 instruction cache is fixed at 4-way set associative as indicated by a default value of 3 for this field.

For more information, refer to [Section 2.2.1.2, "Device Configuration 1 — Config1 \(CP0 Register 16, Select 1\)"](#).

### 5.2.10.2 CacheErr Register (CP0 register 27, Select 0)

The *CacheErr* register is a read-only register used to determine the status of a cache error. The upper two bits of this register (*CacheErr*.*EREC*) indicate whether the contents of the register pertain to an L1 instruction cache error, an L1 data cache error, a TLB error, or an external error. This register provides information such as:

- L1 data versus L2 data cache error
- Tag RAM versus Data RAM error
- External snoop request indication in multi-core systems
- Indicates coherent L1 cache error in another CPU in a multi-core system
- Fatal/non-fatal error indication
- Indicates if the error affects the Scratchpad RAM
- Indicates the cache index or Scratchpad RAM index of the double word entry where the error occurred

For more information, refer to [Section 2.2.6.11, "Cache Error — CacheErr \(CP0 Register 27, Select 0\)"](#).

### 5.2.10.3 L1 Instruction Cache TagLo Register (CP0 register 28, Select 0)

These registers are a staging location for cache tag information being read/written with **cache** load-tag/store-tag operations.

The interpretation of this register changes depending on the setting of the *ErrCtl<sub>WST</sub>* and *ErrCtl<sub>SPR</sub>* bits.

- Default cache interface mode (*ErrCtl<sub>WST</sub>* = 0, *ErrCtl<sub>SPR</sub>* = 0)
- Diagnostic "way select test mode" (*ErrCtl<sub>WST</sub>* = 1, *ErrCtl<sub>SPR</sub>* = 0)
- For scratchpad memory setup (*ErrCtl<sub>WST</sub>* = 0, *ErrCtl<sub>SPR</sub>* = 1)

For more information, refer to [Section 2.2.6.1, "Level 1 Instruction Cache Tag Low — ITagLo \(CP0 Register 28, Select 0\)"](#).

#### 5.2.10.4 L1 Instruction Cache DataLo Register (CP0 register 28, Select 1)

Staging registers for special **cache** instruction which loads or stores data from or to the cache line. Two registers (*IDataHi*, *IDataLo*) are needed, because the interAptiv core loads I-cache data at least 64 bits at a time. This register stores the lower 32 bits of the load data.

For more information, refer to [Section 2.2.6.2, "Level 1 Instruction Cache Data Low — IDataLo \(CP0 Register 28, Select 1\)"](#).

#### 5.2.10.5 L1 Instruction Cache DataHi Register (CP0 register 29, Select 1)

Staging registers for special **cache** instruction which loads or stores data from or to the cache line. Two registers (*IDataHi*, *IDataLo*) are needed, because the interAptiv core loads I-cache data at least 64 bits at a time. This register stores the upper 32 bits of the load data.

For more information, refer to [Section 2.2.6.3, "Level 1 Instruction Cache Data High — IDataHi \(CP0 Register 29, Select 1\)"](#).

### 5.2.11 L1 Instruction Cache Initialization

The L1 instruction cache must be initialized during power-up or reset in order to place the lines of the cache in a known state. This is accomplished via the cache initialization routine, which is normally part of the boot code. For experienced user's, a sample boot code is shown in the following subsection.

#### 5.2.11.1 L1 Instruction Cache Initialization Routine

The following assembly provides an example initialization routine for the instruction cache.

```
/******  
init_icache invalidates all Instruction cache entries  
*****/  
  
LEAF(init_icache)  
  
    // For this Core there is always an instruction cache  
    // The IS field determines how many sets there are:  
    // IS = 2 there are 256 sets  
    // IS = 3 there are 512 sets  
    // $11 set to line size, will be used to increment through the cache tags  
  
    li    $11, 32          # Line size is always 32 bytes.  
    mfc0  $10, $16, 1     # Read C0_Config1  
    ext   $12, $10, 22, 3 # Extract IS  
    li    $14, 2          # Used to test against
```

```

    beq    $14, $12, Isets_done# if IS = 2
    li     $12, 256          # sets = 256
    li     $12, 512          # else sets = 512 Skipped if branch taken
Isets_done:
    lui    $14, 0x8000       # Get a KSeg0 address for cacheops
    // clear the lock bit, valid bit, and the LRF bit
    mtc0   $0, $28          # Clear C0_ITagLo to invalidate entry

next_icache_tag:
    cache  0x8, 0($14)      # Index Store tag Cache opt
    add    $12, -1          # Decrement set counter
    bne    $12, $0, next_icache_tag # Done yet?
    add    $14, $11         # Increment line address by line size
done_icache:

    ins    r31_return_addr, $0, 29, 1
    jr     r31_return_addr
    nop
END(init_icache)

```

### 5.2.11.2 L1 Instruction Cache Initialization Routine Details

This section provides a detailed description of each line of code in the L1 instruction cache initialization routine described above. Note that this code represents an example of an implementation specific cache initialization. The code is used in specific cache sizes of 32K or 64K, is always part of the MPS, and will always have the L2 cache present. The code example is written with those parameters in mind.

Before use, the cache must be initialized to a known state; that is, all cache entries must be invalidated. This code example initializes the cache, finds the total number of cache sets, then loops through the cache sets using the cache instruction to invalidate each cache set.

LEAF (init\_icache)

```

// For this Core there is always an L1 instuction cache
// The IS field determines how many sets there are
// IS = 2 there are 256 sets
// IS = 3 there are 512 sets
// $11 set to line size, will be used to increment through the cache tags

```

```

    li     $11, 32          # Line size is always 32 bytes.

```

This instruction cache always has a line size of 32 bytes, 4 ways and can have a size of either 32 KB or 64 KB. The IS field (sets per way) of the *Config1* register will be used to determine the size of the cache. This field can have one of two values. A value of 0x2 indicates a 32 KB cache and a value of 0x3 indicates a 64 KB cache.

```

    mfc0   $10, $16, 1      # Read C0_Config1
    ext    $12, $10, 22, 3  # Extract IS
    li     $14, 2          # Used to test against

```

If the check is true, the code uses the branch delay slot (which is always executed) to set the set iteration value to 256 for a 32 KB cache and then branches ahead to *Isets\_done*. If the check is false, the code assumes that the size of the cache is 64 KB. At this point, the code still sets the iteration value to 256 in the branch delay slot, but then falls through and sets it again to 512 for a 64 KB cache.

```

beq  $14, $12, Isets_done # if IS = 2
li   $12, 256             # sets = 256
li   $12, 512             # else sets = 512 Skipped if branch taken

```

#### Isets\_done:

GPR 14 will be used as an index into the cache. It will be set to a virtual address, and then translated to a physical address. Since the address 0x8000\_0000 is in kseg0, the CPU will ignore the top bit, so virtual 0x8000\_0000 will become physical address 0x0000\_0000. Since the cache is physically indexed, the first time through the loop, the cache instruction will write the tag to way 0 index line 0.

The **lui** instruction loads 0x8000 into the upper 16 bits and clears the lower 16 bits of the GPR14 register.

```

lui  $14, 0x8000          # Get a KSeg0 address for cacheops

```

Clearing the tag registers performs two important functions: it sets the Physical Tag address called PTagLo to 0, which ensures the upper physical address bits are zeroed out, and it also clears the valid bit for the set, which ensures that the set is free and may be filled as needed.

The code uses the Move to Coprocessor Zero (MTC0) instruction to move the general purpose register zero, which always contains a zero, to the tag register.

```

// clear the lock bit, valid bit, and the LRF bit

```

```

mtc0 $0, $28             # Clear C0_ITagLo to invalidate entry

```

The **Cache** instruction uses the **Index Store Tag** operation on the Level 1 instruction cache so the op field is coded with a value of 0x8. The first two bits are 2'b00 for the L1 instruction cache, and the operation code for **Index Store tag** is encoded as 3'b010 in bits two, three and four.

#### next\_icache\_tag:


```

cache 0x8, 0($14) # Index Store tag Cache op

```

The index type of operation can be used to address a byte in the cache in a specific way of the cache. This is done by breaking down the virtual address argument stored in the base register of the **Cache** instruction into several fields.

#### Bits 14:0 of the Cache Instruction



The size of the index field varies according to the size of a cache way. The larger the way, the larger the index. In the table above, the combined byte and page index is 13 bits because each way of the cache is 8K. The way number is always the next two bits following the index.

The code does not explicitly set the way bits. Instead it just increments the virtual address by the cache lines size so the next time through the loop the **Cache** instruction will initialize the next set in the cache. Eventually this increment has the effect of setting the cache to index 0 of the next way in the cache because it overflows into the way bits.

At this point all the code needs to do is loop maintenance. First decrement the loop counter (12/t4).

```

add  $12, -1              # Decrement set counter

```

Then test it to see if it has gotten to zero and if it has not branch back to label one.

```

bne  $12, $0, next_icache_tag # Done yet?

```

The instruction in the branch delay slot, which is always executed, is used to increment the virtual address (14/t6) to the next set in the cache. (11/t3) holds the line size in bytes.

```
add    $14, $11          # Increment line address by line size
```

From this point on, the code can be executed from a cached address. This is easily done by changing the return address from a KSEG1 address to a KSEG0 address by simply inserting a 0 into bit 29 of the address. However, during debugging, this operation will confuse the debugger and you will no longer be able to do source-level debugging. That is why it is commented out here. Once the code has been debugged, the "ins" line can be uncommented.

**done\_icache:**

```
    // Modify return address to kseg0 which is cacheable
    // (for code linked in kseg1.)
    // However it makes it easier to debug if this is not done. So while
    // debugging, this should be commented out.

    ins    r31_return_addr, $0, 29, 1
    jr     r31_return_addr
    nop

END (init_icache)
```

## 5.2.12 Cache Management When Writing Instructions - the “SYNCI” Instruction

The **synci** instruction (new to the MIPS32 Release 2 update) provides a mechanism available to user-level code for ensuring that previously written instructions are correctly presented for execution (it combines a D-cache writeback with an I-cache invalidate). Use of the **synci** instruction is preferred to the traditional alternative of a D-cache writeback followed by an I-cache invalidate.



## 5.3 L1 Data Cache

The L1 data cache is similar to the instruction cache, with a few key differences;

- The dirty bit is part of the way-select RAM.
- The data cache does not contain any precode information.
- To handle store bytes, the data array is byte-accessible, and the data parity is 1 bit per byte.
- ECC code is generated across a 32b word. Reads and writes are 32 or 64b. Sub-word stores are handled by doing a read-modify-write sequence.

Like the L1 instruction cache, the L1 data cache is virtually indexed, since a virtual address is used to select the appropriate line within each of the arrays. The cache is physically tagged, as the tag array contains a physical, not virtual, address.

The tag and data arrays hold 4 ways of information per set, corresponding to the 4-way set associativity of the cache. The way-select array holds information to choose the way to be filled, as well as dirty bits in the case of the data cache.

A tag entry consists of the upper bits of the physical address bits [31:10], a valid bit, and a lock bit. A data entry contains the four, 64-bit doublewords in the line, for a total of 32 bytes. All four words in the line are present or not in the data array together, hence the single valid bit stored with the tag. Once a valid line is resident in the cache, byte, half-word, triple-byte, word, or doubleword stores can update all or a portion of the words in that line. The tag and data entries are repeated for each of the 4 lines in the set.

A way-select entry holds bits choosing the way to be replaced according to a Least Recently Used (LRU) algorithm. The LRU information applies to all the ways and there is one way-select entry for all the ways in the set.

Table 5.4 shows the key characteristics of the data cache.

**Table 5.4 interAptiv L1 Data Cache Organization**

Attribute	Without Parity	With Parity	With ECC
Size <sup>1</sup>	0 KB, 4 KB, 8 KB, 16 KB, 32 KB or 64 KB	0 KB, 4 KB, 8 KB, 16 KB, 32 KB or 64 KB	0 KB, 4 KB, 8 KB, 16 KB, 32 KB or 64 KB
Line Size	32-byte	32-byte	32-byte
Number of Cache Sets	32, 64, 128, 256, or 512	32, 64, 128, 256, or 512	32, 64, 128, 256, or 512
Associativity	4-way	4-way	4-way
Replacement	LRU	LRU	LRU
Cache Locking	per line	per line	per line
<b>Data Array</b>			
Read Unit	64b x 4	72b x 4	78b x 4
Write Unit	8b	9b	39b
<b>Tag Array</b>			
Read Unit	24b x 4	25b x 4	31b x 4
Write Unit	24b	25b	31b
<b>Way-Select Array</b>			
Read Unit	10b	14b	22b
Write Unit	1-10b	1-14b	1-22b

1. For Linux based applications, MIPS recommends a 64 KB L1 data cache size, with a minimum size of 32 KB.

**Figure 5.3 L1 Data Cache Organization**

**Tag w/ parity (per way):**  
(25 bits total)

1	1	1	22
Parity	Valid	Lock	PA31:10

**Tag w/ ECC (per way):**  
(31 bits total)

7	1	1	22
ECC	Valid	Lock	PA31:10

**Data (per way): Parity**  
(72 bits total)

1	8	Bytes 1 - 6. (each byte contains one parity bit)	1	8
Parity7	Data7	...	Parity0	Data0

**Data (per way): ECC**  
(78 bits total)

7	32	7	32
ECC	Data1	ECC	Data0

**Way-Select: Parity**  
(14 bits total)

4	4	6
Dirty Parity	Dirty	LRU

**Way-Select: ECC**  
(22 bits total)

12	4	6
ECC	Dirty	LRU

### 5.3.1 L1 Data Cache Virtual Aliasing

The data cache on the interAptiv core is virtually indexed and physically tagged. The lower bits of the virtual address are used to access the cache arrays and the physical address is used in the tags. Because the way size can be larger than the minimum TLB page size, there is a potential for virtual aliasing. This means that one physical address can exist in multiple indices within the cache, if it is accessed with different virtual addresses.

The following table indicates the conditions under which virtual aliasing can occur.

**Table 5.5 L1 Data Cache Virtual Aliasing Conditions**

Cache Size	MMU Page Size	Way Size	Aliasing Can Occur	Hardware Aliasing Fix Required
32 KB	4 KB	8 K	Yes	Yes
64 KB	4 KB	16 K	Yes	Yes
32 KB	>= 16 KB	8 K	No	No
64 KB	>= 16 KB	16 K	No	No

In the interAptiv core, the read-only `Config7.AR` bit determines whether the data cache virtual aliasing hardware is enabled based on the build-time configuration. Note that for some of the configuration options in the table above, the hardware aliasing fix (HWAF) is required. As such, it is incumbent upon the designer to select the HWAF option at build time. The selection of this option causes hardware to set the `Config7.AR` bit.

### 5.3.2 L1 Data Cache Parity (Optional)

The L1 cache data parity option provides one parity bit for each byte, corresponding to the minimum number of bytes for a store. The tag array has a single parity bit for each tag. The way-select array has separate parity bits to cover each dirty bit, but the LRU bits are not covered by parity. When the interAptiv core is configured with data cache parity, instruction cache parity is also present.

### 5.3.3 Level 1 Data Cache Error Detection and Correction (EDC)

The interAptiv core includes a configuration option for error detection and correction (EDC) on the Level 1 Data Cache and the Data Scratch Pad Memory. When selected, error correction codes (ECCs) are added to information stored in data-cache and data-scratch-pad RAMs. The error detection and correction logic protects against data corruption caused by single-bit transient errors that may occur while data is stored in RAM. The error codes allow for single-bit error correction and double-bit error detection. ECC generation and checking and error handling is done in the Load/Store Unit (LSU).

#### 5.3.3.1 L1 Data Cache Organization

As shown in the [Figure 5.3](#) above, the interAptiv core level 1 data cache comprises three logical RAM arrays: a tag array, a data array and a way select array. With error detection and correction, a seven-bit ECC is added to each 24-bit tag stored in the tags array; a seven-bit ECC is added to each 32-bit data word stored in the data array; and a three-bit ECC is added to each individual dirty bit stored in the way select array. Multi-processor capable cores also have a duplicate tags array for handling interventions. As with the primary tags array, a seven-bit ECC is added to each 24-bit tag stored in the duplicate tags array.

In addition to the dirty bits, the way select array also comprises a least-recently-used (LRU) code for each data cache set. The LRU code has no ECC because an LRU code error has no effect on data integrity and is self-healing. An LRU code error may lead to a suboptimal cache line replacement, but its impact is fleeting.

#### 5.3.3.2 L1 Data Cache Load/Store Operations

Cacheable loads and stores generate a data cache read to see if the memory operand is in the cache. If an error is detected, incoming loads and stores are halted by hardware and the LSU determines whether an ECC error is uncorrectable or correctable. Uncorrectable errors generate an exception. Correctable errors generate a data cache scrub sequence, during which errors at the affected index in the tags and data arrays are corrected. After scrubbing, the affected access and subsequent operations are replayed. The overall delay for correcting an ECC error is expected to be about 15 clock cycles. Aside from the stall, data cache error corrections are transparent to program execution.

A data cache scrub sequence repairs all ways at a selected index. After the LSU pipeline is stalled, the tags, data and way select arrays are read. A subsequent read is necessary because a write may have occurred between the read on which the error was detected and the stall. In sequence, each of the four ways – including tag and data – are corrected as needed and updated. Duplicate tags are updated as a side-effect of updating primary tags.

#### 5.3.3.3 L1 Data Cache Error Types

L1 data cache ECC errors can be correctable or uncorrectable. Single-bit transient errors are correctable. However, if a single-bit error repeats during replay, it is assumed to be a persistent error and thus uncorrectable. Multiple-bit errors cannot be repaired. Multiple-bit errors in a data word of an invalid cache line are ignored. Keep in mind, a tag

needs to be free of errors to affirm that a line is invalid. Hence, tag errors are processed before processing multiple-bit data errors. A multiple-bit error is uncorrectable if it occurs in (a) a tag, (b) a dirty bit of a valid cache line, (c) a data word in a dirty cache line or (d) a data word in a locked cache line.

#### 5.3.3.4 Store Operations Less than 32-bits

The addition of ECC to the cache data array has special implications for stores into the data cache when the operand is smaller than a single 32-bit word (i.e. either a byte or a half-word). When those partial-word stores hit in the cache, the LSU performs a cache read-modify-write on the affected word because the ECC is a function of the entire 32-bit word.

With ECC, a byte or half-word store generates a read of the target word index in the data array, as well as reading the tags array. In the case of a hit, the store data is merged with the target word read from cache and the entire word is put into the Fill/Store Buffer. Opportunistically, when there is an idle data cache access cycle, the modified word, with an updated ECC, is written into the data array. From the ALU's perspective, there is no difference in the timing of a byte or half-word store between an ECC implementation and a non-ECC implementation.

In non-ECC configurations byte or half-word stores do not require a read-modify-write of the data cache. A cacheable store generates a read of the tags arrays, but not the data arrays. The store operand is put in the Fill-Store Buffer. The data cache has byte write enables. If the store was a cache hit, only the bytes in the store operand written into the cache. The rest of the word is unaffected.

#### 5.3.3.5 Examples of L1 Data Cache ECC Errors

Consider some data cache ECC error scenarios:

##### ***Loads and Stores***

During local CPU loads and stores, single-bit errors in the primary tags array are scrubbed on detection. Multiple-bit errors in the primary tag array generate an exception.

During local CPU loads and stores, single-bit errors in the data array of valid lines are scrubbed on detection. Double-bit data errors generate an exception if a line is valid. The exception handler may invalidate a line with a double-bit data error if it is shared. Data errors (single-bit or double-bit) in invalid lines are ignored.

##### ***Evictions***

During an eviction, correctness of the dirty bit is only pertinent when the evicted line is in the exclusive state. Single-bit dirty bit errors are corrected on the fly. In the event of a multi-bit dirty bit error, the associated cache line is assumed to be dirty.

During eviction of a dirty cache line, single-bit data errors are corrected on the fly as data is written back to the Bus Interface Unit (BIU). Multiple-bit errors in an evicted line are reported as a parity error (i.e. an uncorrectable error) to the BIU and generate an exception.

##### ***Interventions***

During interventions, single-bit errors in the primary tags array are scrubbed on detection. Multiple-bit errors in the primary tags array generate an exception and return an ERROR response for the intervention.

During an intervention, correctness of the dirty bit is pertinent on an intervention hit when the line is in the exclusive state. Single-bit dirty bit errors are corrected on the fly. Multiple-bit dirty bit errors generate an exception and return an ERROR response for the intervention.

During an intervention write-back of an exclusive cache line, single-bit data errors are corrected on the fly as data is forwarded to the BIU. If a copy of the line is retained in the data cache, it is not scrubbed during the intervention

write-back. A subsequent load or store is required to scrub the error. Multiple-bit data errors during an intervention write-back are reported as a parity error (i.e. an uncorrectable error) to the BIU and an exception is generated.

An intervention may also encounter an ECC error in the primary cache arrays. These errors are handled the same as with other cache accesses. The pipeline stalls. An uncorrectable error generates an exception. A correctable error generates a scrub sequence. On completion of the scrub sequence, the intervention is replayed.

The duplicate tags array also have error detection and correction capability. The duplicate tags are identical copies of the primary cache tags. Consequently, any single-bit or multiple-bit error detected in the duplicate tag can be corrected by updating the erroneous tag with the value stored in the primary tags array. When a duplicate tag ECC error is detected during an intervention, the intervention sequencer makes a special request to the cache access control logic to force a scrub of the affected cache index. As the tag in each way of the primary tags array is updated, the copy in the duplicate tags array is also updated. After the scrub, the intervention proceeds normally.

### **Dirty Bit Errors**

In the event of an uncorrectable dirty bit error, the associated cache line is assumed to be dirty. A correctable dirty bit error is corrected on-the-fly before determining whether the line needs to be written back to memory. No scrubbing is done in the way select array.

## **5.3.4 L1 Data Cache Replacement Policy**

The replacement policy refers to how a way is chosen to hold an incoming cache line on a miss which will result in a cache fill. The replacement policy is least-recently used (LRU), but excluding any locked ways. The LRU bit(s) in the way-select array encode the order in which ways on that line have been accessed.

On a cache miss, the lock and LRU bits for the tag and way-select entries of the selected line may be used to determine the way which will be chosen.

The LRU field in the way select array is updated as follows:

- On a cache hit, the associated way is updated to be the most recently used. The order of the other ways relative to each another is unchanged.
- On a cache refill, the filled way is updated to be the most recently used.
- On CACHE instructions, the update of the LRU bits depends on the type of operation to be performed:
  - **Index (Writeback) Invalidate:** Least-recently used.
  - **Index Load Tag:** No update.
  - **Index Store Tag,  $WST = 0$ :** Most-recently used if valid bit is set in *TagLo* CP0 register. Least-recently used if valid bit is cleared in *TagLo* CP0 register.
  - **Index Store Tag,  $WST = 1$ :** Update the field with the contents of the *TagLo* CP0 register.
  - **Index Store Data:** No update.
  - **Hit Invalidate:** Least-recently used if a hit is generated, otherwise unchanged.
  - **Hit (Writeback) Invalidate:** Least-recently used if a hit is generated, otherwise unchanged.
  - **Hit Writeback:** No update.
  - **Fetch and Lock:** For instruction cache, no update. For data cache, most-recently used.

If all ways are valid, then any locked ways will be excluded from consideration for replacement. For the unlocked ways, the LRU bits are used to identify the way which has been used least-recently, and that way is selected for replacement.

If the way selected for replacement has its dirty bit asserted in the way-select array, then that 32-byte line will be written back to memory before the new fill can occur.

### 5.3.5 L1 Data Cache Line Locking

The mechanism for line locking in the L1 data cache is identical to that of the L1 instruction cache. For more information, refer to [Section 5.2.5, "L1 Instruction Cache Line Locking"](#).

### 5.3.6 L1 Data Cache Memory Coherence Protocol

The interAptiv core supports cache coherency in a multi-CPU cluster using Cache Coherence Attributes (CCAs) specified on a per cache-line basis and an Intervention Port containing coherent requests by all CPUs in the system. Each interAptiv core monitors its Intervention Port and updates the state of its cache lines (valid, lock, and dirty tag bits) accordingly.

The L1 data caches utilize a standard MESI protocol. Each cache line will be in one of the following four states:

**Invalid:** The line is not present in this cache.

**Shared:** This cache has a read-only copy of the line. The line may be present in other L1 data caches, also in a Shared state. The line will have the same value as it does in the L2 cache or memory.

**Exclusive:** This cache has a copy of the line with the right to modify. The line is not present in other L1 data caches. The line is still clean - consistent with the value in L2 cache or memory.

**Modified:** This cache has a dirty copy of the line. The line is not present in other L1 data caches. This is the only up-to-date copy of the data in the system (the value in the L2 cache or memory is stale).

The SYNC instruction may also be useful to software in enforcing memory coherence, because it flushes the write buffers.

Some of the basic characteristics of the coherence protocol are summarized below. Coherence can occur on the data cache.

- Writeback cache - Uses a writeback cache to ensure high performance
- Cache-line based - Coherence and ownership is maintained per 32-byte cache line
- Snoopy protocol - Each CPU snoops the stream of transactions and updates its cache state accordingly
- Invalidate - A line is invalidated from the cache (possibly with a writeback to memory) when a store from another processor is seen.

### 5.3.7 L1 Data Cache Initialization

The L1 data cache must be initialized during power-up or reset in order to place the lines of the cache in a known state. This is accomplished via the cache initialization routine, which is normally part of the boot code. For experienced user's, a sample boot code is shown in the following subsection.

### 5.3.7.1 L1 Data Cache Initialization Routine

The following assembly provides an example initialization routine for the data cache.

```

/*****
init_dcache invalidates all data cache entries
*****/

LEAF (init_dcache)

    // For the interAptiv MPS there is always an L1 data cache
    // The ID field determines how many sets there are
    // DS = 2 there are 256 sets
    // DS = 3 there are 512 sets
    // $11 set to line size, will be used to increment through the cache tags

    li    $11, 32          # Line size is always 32 bytes
    mfc0  $10, $16, 1      # Read C0_Config1
    ext   $12, $10, 13, 3 # Extract DS
    li    $14, 2          # Used to test against
    beq   $14, $12, Dsets_done # if DS = 2
    li    $12, 256        # sets = 256
    li    $12, 512        # else sets = 512, skipped if branch taken

Dsets_done:

    lui   $14, 0x8000     # Get a KSeg0 address for cacheops
    // clear the lock bit, valid bit, and the LRF bit
    mtc0  $0, $28, 2      # Clear C0_DTagLo to invalidate entry

next_dcache_tag:

    cache 0x9, 0($14)     # Index Store tag Cache opt
    add   $12, -1         # Decrement set counter
    bne   $12, $0, next_dcache_tag # Done yet?
    add   $14, $11        # Increment line address by line size

done_dcache:

    jr    r31_return_addr
    nop

END (init_dcache)

```

### 5.3.7.2 L1 Data Cache Initialization Routine Details

This section provides a detailed description of each line of code in the initialization routine. The L1 data cache initialization routine is very similar to the L1 instruction cache initialization routine.

```

LEAF(init_dcache)

    // For the interAptiv CPS there is always a L1 data cache
    // The DS field determines how many sets there are
    // DS = 2 there are 256 sets
    // DS = 3 there are 512 sets
    // $11 set to line size, will be used to increment through the cache tags

    li    $11, 32          # Line size is always 32 bytes.

```

The data cache always has a line size of 32 bytes and 4 ways, and can have a size of either 32 KB or 64 KB. The DS field (sets per way) of the *Config1* register is used to determine the size of the cache. This field can have one of two values. A value of 0x2 indicates a 32 KB cache and a value of 0x3 indicates a 64 KB cache.

```
mfc0 $10, $16, 1      # Read C0_Config1
ext  $12, $10, 13, 3  # Extract DS
li   $14, 2           # Used to test against
```

If the check is true, the code uses the branch delay slot (which is always executed) to set the set iteration value to 256 for a 32 KB cache and then branches ahead to **Dsets\_done**. If the check is false, the code assumes that the size of the cache is 64 KB. At this point, the code still sets the iteration value to 256 in the branch delay slot, but then falls through and sets it again to 512 for a 64 KB cache.

```
beq  $14, $12, Dsets_done # if DS = 2
li   $12, 256             # sets = 256
li   $12, 512             # else sets = 512 Skipped if branch taken
```

#### **Dsets\_done:**

GPR 14 will be used as an index into the data cache. It is set to a virtual address and then translated to a physical address. Since the address 0x8000\_0000 is in kseg0, the CPU will ignore the top bit, so virtual 0x8000\_0000 will become physical address 0x0000\_0000. Since the cache is physically indexed, the first time through the loop, the cache instruction will write the tag to way 0 index line 0.

The **lui** instruction loads 0x8000 into the upper 16 bits and clears the lower 16 bits of the GPR14 register.

```
lui  $14, 0x8000        # Get a KSeg0 address for cacheops
```

Clearing the tag registers performs two important functions: it sets the Physical Tag address called PTagLo to 0, which ensures the upper physical address bits are zeroed out, and it also clears the valid bit for the set, which ensures that the set is free and may be filled as needed.

The code uses the Move to Coprocessor zero instruction to move the general purpose register zero, which always contains a zero, to the tag register.

```
// clear the lock bit, valid bit, and the LRF bit
mtc0 $0, $28, 2        # Clear C0_DTagLo to invalidate entry
```

The **Cache** instruction uses the **Index Store Tag** operation on the Level 1 data cache so the op field is coded with a value of 0x9. The first two bits are 2'b01 for the L1 data cache, and the operation code for **Index Store tag** is encoded as 3'b010 in bits two, three and four.

#### **next\_dcache\_tag:**

```
cache 0x9, 0($14) # Index Store tag Cache opt
```

The index type of operation can be used to address a byte in the cache in a specific way of the cache. This is done by breaking down the virtual address argument stored in the base register of the **Cache** instruction into several fields.

#### **Bits 14:0 of the Cache Instruction**



The size of the index field varies according to the size of a cache way. The larger the way, the larger the index. In the table above, the combined byte and page index is 13 bits because each way of the cache is 8K. The way number is always the next two bits following the index.



The code does not explicitly set the way bits. Instead it just increments the virtual address by the cache line size so the next time through the loop the **Cache** instruction will initialize the next set in the cache. Eventually this increment has the effect of setting the cache to index 0 of the next way in the cache because it overflows into the way bits.

At this point all the code needs to do is loop maintenance. First decrement the loop counter (12/t4).

```
add    $12, -1           # Decrement set counter
```

Then test it to see if it has gotten to zero and if not branch back to label one.

```
bne    $12, $0, next_dcachetag # Done yet?
```

The instruction in the branch delay slot, which is always executed, is used to increment the virtual address (14/t6) to the next set in the cache. (11/t3) holds the line size in bytes

```
add    $14, $11         # Increment line address by line size
```

At this point the Dcache initialization is done.

**done\_dcachetag:**

```
jr     r31_return_addr
nop
```

END (init\_dcachetag)

### 5.3.8 Data Cache CP0 Register Interface

The interAptiv core uses the following CP0 registers for data cache operations.

**Table 5.6 Data Cache CP0 Register Interface**

CP0 Registers	CP0 number
<i>Config1</i>	16.1
<i>CacheErr</i>	27.0
<i>DTagLo</i>	28.2
<i>DTagHi</i>	29.2
<i>DDataLo</i>	28.3

#### 5.3.8.1 Config1 Register (CP0 register 16, Select 1)

The *Config1.DS* field (bits 15:13) indicates the number of sets per way in the data cache. The interAptiv L1 data cache supports 256 sets per way, which is used to configure a 32 KB cache, or 512 sets per way, which is used to configure a 64 KB cache.

The *Config1.DL* field (bits 12:10) indicates the line size for the data cache. The interAptiv L1 data cache supports a fixed line size of 32 bytes as indicated by a default value of 4 for this field.

The *Config1.DA* field (bits 9:7) indicates the set associativity for the data cache. The interAptiv L1 data cache is fixed at 4-way set associative as indicated by a default value of 3 for this field.

For more information, refer to [Section 2.2.1.2, "Device Configuration 1 — Config1 \(CP0 Register 16, Select 1\)"](#).

### 5.3.8.2 CacheErr Register (CP0 register 27, Select 0)

The *CacheErr* register is a read-only register used to determine the status of a cache error. The upper two bits of this register (*CacheErr.EREC*) indicate whether the contents of the register pertain to an L1 instruction cache error, an L1 data cache error, a TLB error, or an external error.

For more information, refer to [Section 2.2.6.11, "Cache Error — CacheErr \(CP0 Register 27, Select 0\)"](#).

### 5.3.8.3 L1 Data Cache TagLo Register (CP0 register 28, Select 2)

These registers are a staging location for cache tag information being read/written with **cache** load-tag/store-tag operations.

In a multi-core system, the D-cache has five logical memory arrays associated with this *DTagLo* register. The tag RAM stores tags and other state bits with special attention to the needs of the CPU. The duplicate tag RAM also stores tags and state, but is optimized for the needs of interventions. Both of these arrays are set-associative (4-way). The Dirty RAM and duplicate Dirty RAM store the dirty bits (indicating modified data) for CPU and intervention uses, and each combine their ways together in a single entry per set. The WS RAM combines the dirty and LRU data in a single entry per set. Accessing these arrays for index cache loads and stores is controlled by using three bits in the *ErrCtl* register to create modes that allow the correct access to these arrays.

The interpretation of this register changes depending on the settings of *ErrCtl<sub>WST</sub>*, *ErrCtl<sub>DYT</sub>*, and *ErrCtl<sub>SPR</sub>*.

For more information, refer to [Section 2.2.6.4, "Level 1 Data Cache Tag Low — DTagLo \(CP0 Register 28, Select 2\)"](#).

### 5.3.8.4 L1 Data Cache TagHi Register (CP0 register 29, Select 2)

The *DTagHi* register is used to store ECC error information for the L1 data cache and DSPRAM memories. The bit assignments of the register depends on the type of memory being accessed. On a DSPRAM ECC error, bits 19:0 contain ECC error information. On a L1 data cache tagor data error, bits 16:10 contains data RAM ECC information, and bits 6:0 contain tag RAM error information.

For For more information, refer to [Section 2.2.6.5, "DTagHi \(CP0 Register 29, Select 2\): L1 Data Cache and DSPRAM ECC"](#).

### 5.3.8.5 L1 Data Cache DataLo Register (CP0 register 28, Select 3)

In the interAptiv core, software can read or write cache data using a **cache** index load tag/index store data instruction. Which word of the cache line is transferred depends on the low address fed to the **cache** instruction.

Note that the interAptiv core does not implement the *DDataHi* register.

For more information, refer to [Section 2.2.6.6, "Level 1 Data Cache Data Low — DDataLo \(CP0 Register 28, Select 3\)"](#).

## 5.4 L1 Instruction and Data Cache Software Testing

Typically, the cache RAM arrays will be tested using BIST. It is, however, possible for software running on the processor to test all of the arrays. Of course, testing of the I-cache arrays should be done from an uncacheable space with interrupts disabled in order to maintain the cache contents. There are multiple methods for testing these arrays in software, some of which are described in the following subsections.

### 5.4.1 L1 Instruction Cache Tag Array

The L1 instruction cache tag array can be tested via the **Index Load Tag** and **Index Store Tag** varieties of the **CACHE** instruction. An **Index Store Tag** writes the contents of the *ITagLo* and *ITagHi* registers into the selected tag entry. An **Index Load Tag** reads the selected tag entry into the *ITagLo* and *ITagHi* registers.

If parity is implemented, the parity bits can be tested as normal bits by setting the *PO* (parity override) bit in the *ErrCtl* register. This will override the parity calculation and use the parity bits in *ITagLo* and *ITagHi* as the parity values.

### 5.4.2 Instruction Cache Data Array

This array can be tested using the Index Store Data and Index Load Tag varieties of the **CACHE** instruction. The Index Store Data variety is enabled by setting the *WST* bit in the *ErrCtl* register.

The Index Store Data instruction can optionally update the corresponding precode field in the tag array. The precode bits in the array are updated if the *PCD* bit in the *ErrCtl* register is zero when executing the Index Store Data instruction. The precode value is generated by the hardware automatically if the *PCO* bit in the *ErrCtl* register is zero. Otherwise, the corresponding precode value (*PREC\_01*/*PREC\_23*/*PREC\_45*/*PREC\_67*) from the *ITagHi* register is used in updating the tag array.

The parity bits in the array can be tested by setting the *PO* bit in the *ErrCtl* register. This will use the *PI* field in *ErrCtl* instead of calculating the parity on a write.

The rest of the data bits are read/written to/from the *IDataLo* and *IDataHi* registers.

### 5.4.3 Instruction Cache Way Select Array

The testing of this array is done with via Index Load Tag and Index Store Tag **CACHE** instructions. By setting the *WST* bit in the *ErrCtl* register, these operations will read and write the *WS* array instead of the tag array.

### 5.4.4 L1 Data Cache Tag Array

The L1 data cache tag array can be tested via the **Index Load Tag** and **Index Store Tag** varieties of the **CACHE** instruction. An **Index Store Tag** writes the contents of the *DTagLo* register into the selected tag entry. An **Index Load Tag** will read the selected tag entry into the *DTagLo* register.

If parity is implemented, the parity bits can be tested as normal bits by setting the *PO* (parity override) bit in the *ErrCtl* register. This will override the parity calculation and use the parity bits in *DTagLo* as the parity values.

If ECC is implemented, the ECC bits can be tested as normal bits by setting the *PO* (parity override) bit in the *ErrCtl* register. This overrides the ECC calculation and uses the ECC bits in the *DTagHi* register as the ECC values.

### 5.4.5 Duplicate Data Cache Tag Array

This array can be tested via the Index Load Tag and Index Store Tag varieties of the **CACHE** instruction. In order to access the duplicate tags, the *WST* and *SPR* bits of *ErrCtl* should both be set. Index Store Tag will write the contents of the *TagLo* register into the selected tag entry. Index Load Tag will read the selected tag entry into the *TagLo*. In normal mode, with *WST* and *SPR* cleared, *IndexStoreTags* will write into both the primary and duplicate tags, while *IndexLoadTags* will read the primary tag.

If parity is implemented, the parity bit can be tested as a normal bit by setting the *PO* bit in the *ErrCtl* register. This will override the parity calculation and write *P* bit in *TagLo* as the parity value.

### 5.4.6 Data Cache Data Array

This array can be tested using the Index Store Tag CACHE, SW, and LW instructions. First, use Index Store Tag to set the initial state of the tags to valid with a known physical address (PA). Write the array using SW instructions to the PAs that are resident in the cache. The value can then be read using LW instructions and compared to the expected data.

The parity bits can be implicitly tested using this mechanism. The parity bits can be explicitly tested by setting the *PO* bit in *ErrCtl* and using Index Store Data and Index Load Tag CACHE operations. The parity bits (one bit per byte) are read/written to/from the *PD* field in *ErrCtl*. Unlike the I-cache, the *DataHi* register is not used, and only 32b of data is read/written per operation.

### 5.4.7 Data Cache Way Select Array

The dirty and LRU bits can be tested using the same mechanism as the I-cache WS array.

## 5.5 L2 Cache

The L2 cache processes transactions that are not serviced by the L1 cache. L2 is generally larger than the L1 cache, but slower, due to the use of higher-density memories.

The L2 communicates with external memory via an Open Core Protocol (OCP) interface. Because the L2 cache is integrated into the Coherence Manager (CM2) in the interAptiv core, no OCP interface between the two is required, reducing both latency and complexity.

The L2 also communicates with the CPU(s) through the performance counter interface, error reporting interface, and other side band signals. In addition to these interfaces, the L2 has the clock, reset, and bypass signals as well as some static input signals which can be used to configure it for different operating modes.

### 5.5.1 L2 Cache General Features

- 7-stage pipeline. (Optional 8th stage<sup>1</sup> for pipelined memory arrays.)
- 32-bit address paths and 256-bit internal data paths
- Associativity: 8-way
- Cache size: 0K, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, 1 MB, 2 MB, 4 MB, 8 MB
- Line Size: 32 or 64 bytes (4 or 8 doublewords)
- Locking Support: Yes
- Replacement Algorithm: Pseudo LRU for 8-way
- Write policy: Write Back and Write through
- Write miss allocation policy: No-Write-Allocate and Write-Allocate.
- Error Checking and Correction (ECC): Optional 2-bit error detection and 1-bit error correction covering the tag and data arrays. 1-bit error detection covering the WS array.
- Maximum read misses outstanding: 8, 12 or 15. Build-time configuration option.
- Out-Of-Order processing (OOO): Yes
- Coherency: Non-coherent
- 64-bit, 128-bit, or 256-bit OCP SData/MData width on memory-side OCP interface.
- OCP Burst Size on the memory interface: 64-byte line size: 8 beats of 64-bit data or 2 beats of 256-bit data
- Bypass Mode Support: In bypass mode, all processor requests are routed to the system. This mode is used only for debug purposes and should not be used during normal operation.
- Multi-cycle Data Rams: 0, 1, 2, or 3 stalls can set Data RAM access times to 1, 2, 3, or 4 clocks.
- Multi-cycle Tag Rams: 0, 1, 2, or 3 stalls can set Tag RAM access times to 1, 2, 3, or 4 clocks.
- Multi-cycle Way-Select Rams: 0, 1, 2, or 3 stalls can set the Way-Select RAM access times to 1, 2, 3, or 4 clocks.

---

1. Build time option. The customer must choose this option if they are using pipelined RAM's in the wrappers instead of standard RAM cells (that are not pipelined in this way).

- Endianness: Independent of endianness

**Table 5.7 L2 Cache Attributes**

Attribute	With Parity
Size	0 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, 1 MB, 2 MB, 4 MB, or 8 MB
Line Size	32-byte or 64-byte
Number of Cache Sets	512, 1024, 2048, 4096, 8192, 16384 or 32768
Associativity	8 way

In the table above, the associativity of the L2 cache is fixed at 8 ways. As a result, changes to the number of sets per way and the line size determine the overall size of the L2 cache. The interAptiv core only supports the cache sizes shown in [Table 5.7](#) above. As a result, some of the options for line size and sets per way are invalid as they would result in cache sizes being either smaller or larger than those listed above. [Table 5.8](#) shows the list of possible configurations and which ones are valid or invalid. The invalid configurations are shaded in the table.

**Table 5.8 Valid and Invalid L2 Cache Configurations**

Line Size	Sets per Way	Number of Ways	Total L2 Cache Size	Valid L2 Cache Configuration	Notes
32 bytes	512	8	128 KBytes	Yes	
32 bytes	1024	8	256 KBytes	Yes	
32 bytes	2048	8	512 KBytes	Yes	
32 bytes	4096	8	1 MByte	Yes	
32 bytes	8192	8	2 MByte	Yes	
32 bytes	16384	8	4 MByte	Yes	
32 bytes	32768	8	8 MByte	Yes	
64 bytes	64	8	32 KBytes	Yes	No ECC
64 bytes	128	8	64 KBytes	Yes	No ECC
64 bytes	256	8	128 KBytes	Yes	No ECC
64 bytes	512	8	256 KBytes	Yes	
64 bytes	1024	8	512 KBytes	Yes	
64 bytes	2048	8	1 MByte	Yes	
64 bytes	4096	8	2 MByte	Yes	
64 bytes	8192	8	4 MByte	Yes	
64 bytes	16384	8	8 MByte	Yes	
64 bytes	32768	8	16 MByte	No	32768 sets/way valid only with 32 byte line size

### 5.5.2 0K L2 Cache Option

The interAptiv core contains a 0K cache option that is selected during IP configuration. If the ‘Enable L2’ option is selected, then the cache size field can be used to select the cache size between 128K and 8MB. If this option is not selected, the L2 cache is disabled and the cache size field is not available. This is the 0K option.


### 5.5.3 32K and 64K L2 Cache Options

The interAptiv core contains 32K and 64K cache options that are selected during IP configuration. These options are supported only when the L2 cache is configured with a 64-byte line size and without ECC.

### 5.5.4 OCP Interface

In the interAptiv core, the L2 cache is integrated into the CM2. This integration improves performance by eliminating the OCP interface that originally connected the L2 cache to the CM, or the L2 cache to the CPU depending on configuration. The OCP interface between the CM2 and the memory is programmable for widths of either 64-bit, 128-bit, or 256-bit and has a fixed 64-byte line size. This is shown in [Figure 5.4](#).

**Figure 5.4 .OCP Interface Between CM2 and Memory**



### 5.5.5 L2 Replacement Policy

The interAptiv core uses a pseudo-LRU replacement algorithm. The system memory configuration does not affect the replacement policy.

### 5.5.6 L2 Allocation Policy

The L2 cache controller can change its allocation policy based on the state of the *PB\_MReqInfo[4]* pin. This feature is only supported with 32-byte cache line size. With 64-byte cache line, *PB\_MReqInfo[4]* is ignored and the L2 cache always defaults to the behavior as if *PB\_MReqInfo[4]* is zero - no allocate on write miss and allocate on read miss.

#### 5.5.6.1 Write

The L2 will process the cached full-line writes in different ways depending on the allocation policy attribute on *PB\_MReqInfo*. [Table 5.9](#) shows the operations of the L2 according to the different allocation types.

When a write hits in the L2 cache, the data in the cache is always updated with the new data from the OCP master.

**Table 5.9 L2 writes - Full Line AND Write-Back Cacheable (cca 3)**

Allocation policy ( <i>PB_MReqInfo[4]</i> )	L2 hit/miss	What L2 does <sup>1</sup>
No allocate (0)	hit	overwrite, mark dirty
	miss	write out to the main memory, no-allocate
Allocate (1)	hit	overwrite, mark dirty
	miss	write-allocate, mark dirty

1. If not mentioned, L2 does not generate a main memory write.

When a write misses in the L2 cache, the data is written into either the L2 cache or the main memory depending on *PB\_MReqInfo[4]*.

Note that the L2 never allocate on miss if the write is write-through type or is writing a partial data. [Table 5.10](#) shows the L2 behavior.

**Table 5.10 L2 writes - partial data OR write-through cacheable (cca 0)**

Allocation policy ( <i>PB_MReqInfo[4]</i> )	L2 hit/miss	What L2 does
Don't care	hit	update the cache. write out to memory if it is write-through type
	miss	no-allocate, just write out to the main memory

The interAptiv MR2 core adds supports for allocation of partial writes in the L2 cache. Cacheable or coherent writes of any size driven into the IOCU are allocated in the L2 if the *IC\_MReqInfo[2]* pin is asserted with the request. The L2 supports partial write allocation when configured with either 32B or 64B cachelines.

The L2 cache also supports merging of write data when there is already outstanding reads to memory. For example, if an I/O device issues a 64B coherent write request with *IC\_MReqInfo[2]* asserted, the IOCU splits that request into two 32B writes. The L2 receives the first 32B write and, assuming the L2 is configured with 64B cachelines, issues a 32B read from memory (for the 2nd of the line). When the IOCU issues the 2nd 32B write (for the 2nd of the line), the L2 merges that data with the 1st write before filling the L2 cache. In this case, the memory data is dropped.

Note that the a write will only be allocated if the request is cacheable (writeback) or coherent after the CCA override is applied (if applicable). A write request that has a CCA override of WT will not be allocated in the L2.

### 5.5.6.2 Read

*MReqInfo[4]* field also controls whether or not the L2 allocates read data after a miss in the cache array.



As shown in [Table 5.11](#), on a read hit, there are no differences; the line will be returned to the core. However, when the read misses, depending on the value of *PB\_MReqInfo[4]*, the line that has been brought in from the main memory will end up residing in the L2 cache in the end of the operation.

**Table 5.11 L2 reads - cacheable (cca 0/3)**

L2 hit/miss	Allocation policy ( <i>MReqInfo[4]</i> )	What L2 does
hit	Don't care	Return the data. Keep the line in the L2.
miss	Allocate (1)	Get the data from memory. Return the data to the core. Allocate into the L2.
	No allocate (0)	Get the data from memory. Return the data to the core. Don't allocate into the L2.

### 5.5.7 Write-Through vs. Write-Back

Write-through and write-back operations are both supported. The L2 decodes *MReqInfo[2:0]* fields and determines which way to handle the write data.

When a write hits in the L2 cache, the data is written into the L2 cache, and also sent to the main memory when it was write-through type (*MReqInfo[2:0]* = 0).

When a write misses, the no-write-allocation policy is employed in most cases. That is, the write data is forwarded to the main memory without updating the L2 cache contents. However, for the write-back type write with full line data, usually resulting from the L1 D-cache eviction, the L2 supports write-allocate on miss as well as the normal no-allocate policy. This is controlled by the value on *MReqInfo[4]* that is set by the OCP requester. Please refer to the [Section 5.5.6 “L2 Allocation Policy”](#) for more details.

### 5.5.8 Cacheable vs. Uncacheable vs. Uncached Accelerated

The L2 cache supports cacheable and uncacheable accesses. This information also is conveyed on the *MReqInfo[2:0]* field. Cacheable operations access the cache memories, whereas an uncached access bypasses the L2 cache arrays and is sent directly to the main memory.

Uncached accelerated accesses are treated the same way as non-accelerated uncached accesses. This CCA enables uncached transactions to better utilize bus bandwidth via burst transactions. L2 supports single-beat as well as 4-beat burst uncacheable transactions for both read and write operations.

### 5.5.9 Cache Aliases

The L2 cache is physically addressed and physically tagged. It is not subject to virtual aliasing.

### 5.5.10 Performance Counters

The L2 tracks and reports to core the number of the following events.

- the number of cached accesses
- the number of misses
- the number of write backs

- the amount of cycles the L2 is held due to misses
- the number of single bit errors that were corrected
- L2 pipeline utilization — Counts the number of starts into the TA stage of the L2 pipeline
- L2 hit qualifier — Counts different types of L2 cache hits and misses, crossed with the instruction being requested

## 5.5.11 Sleep Modes

The L2 cache contains two basic sleep modes:

- Instruction controlled sleep mode using the WAIT instruction
- Internal dynamic sleep mode

### 5.5.11.1 Sleep Mode Using the WAIT Instruction

In addition to slowing down or stopping the primary *cm\_clk* input, software may initiate low-power Sleep Mode via the execution of the WAIT instruction in the processor.

When the processor enters into Sleep Mode, it will assert *SI\_Sleep*. The *SI\_Sleep* drives the *SI\_L2\_Sleep* input to the L2. The L2 then enters a low-power state and asserts the *L2\_Sleep* output once all outstanding bus activity has completed. Most clocks in the L2 will be stopped, but a handful of flops will remain active to sense the wake up call from the processor, which is the deassertion of *SI\_L2\_Sleep*.

Power is reduced since the global clock goes to the vast majority of flops within the L2, which are held idle during this period. There is no bus activity while the L2 is in sleep mode, so the system bus logic which interfaces to the L2 could be placed into a low power state as well.

When the L2 samples *SI\_L2\_Sleep* asserted and there is no activity in the L2, the L2 will assert *L2\_Sleep* two *cm\_clks* later. Any activity in the L2 will delay the start of *L2\_Sleep* assertion.

When *SI\_L2\_Sleep* is deasserted, the L2 will deassert *L2\_Sleep* and assert *PB\_SCmdAccept* two clocks later. If there is a valid *PB\_MCcmd* waiting at the L2 pins at the *cm\_clk*, then the following *cm\_clk* will have a coincident internal *l2\_clk* edge (clocks are now enabled) and the command that was accepted is launched into the pipeline as indicated by *inst\_ta*. The following clock after that will have an *l2\_tram\_clk* that initiates the tag ram access for that command. Thus, there is a four *cm\_clk* latency from *SI\_L2\_Sleep* deassertion to the start of a tag ram access.

### 5.5.11.2 Internal Dynamic Sleep Mode

When there is no activity at the input pins of the L2 cache and all pending transactions from the CPU are completed, the L2 cache will eventually empty. When this occurs, the L2 cache will turn off the *l2\_clk* signal after some small delay. Only data of value in the CMOS SRAM's retains state.

Beside the WAIT instruction induced sleep mode, the L2 is also equipped with the dynamic global clock gating. When there are no pending transactions in the L2 cache, the L2 shuts down the majority of internal clocks to save power. While the most part of the L2 cache can be turned off, the minimum required logic on the core-side OCP interface remain active. Thus, the L2 cache can accept a new OCP request from core at any time, and this will wake up the whole L2 cache controller.

## 5.5.12 Bypass Mode

**Note:** Bypass mode is strictly a debug feature and is not intended to be a normal mode of operation. It was not intended for active switching during normal operation.

Bypass mode is a test/bringup feature that causes the L2 cache to forward all requests received from either the core or the Coherency Manager to the OCP system interface to main memory. Entering or exiting from Bypass Mode other than at reset requires flushing of the L2 cache while running from uncached memory to restore the L2 cache state to a stable state. In bypass mode, all requests are forwarded to the system as received including L2 CACHE instructions and SYNCs.

## 5.5.13 Reduced L2 Hit Latency

The CM2 integrates the CM and L2 cache into a single, more tightly-coupled component, providing reduced L2 hit latency. [Table 5.12](#) provides the latencies for a read request from a interAptiv core to an idle CM2.

- The system is idle prior to this request
- The L2 cache is configured with no L2 Tag RAM or Data RAM stalls
- The L2 is configured without ECC
- L2-to-memory clock ratio is 1:1
- The L2 is configured with non-pipelined Data RAMs

The column labeled “Original CM/L2 (1:1)” shows the latency for the legacy CM and L2, assuming the CM-L2 clock ratio is 1:1. The “Original CM/L2 (2:1)” column is the same, but assumes the CM-L2 clock ratio is 2:1.

**Table 5.12 Comparison of CM2 and Legacy CM + L2 Read Latencies (in core clock cycles)**

Request CCA	Cache Hit/Miss	CM2	Original CM/L2 (1:1)	Original CM/L2 (2:1)
Coherent (CWB, CWBE)	L1 Miss/L2 Hit	10	15	25
	L1 Hit	15	15	15
	L1 Miss/L2 Miss	14	16	27
Cached/Non-coherent (WB)	L2 Hit	11	16	26
	L2 Miss	15	17	28
Uncached (UC)	---	12	14	22
GCR Read	---	8	8	8
Coherent Upgrade	Intervention Response of SHARED	11	11	11

## 5.5.14 L2-only Sync

The CM2 adds the ability to issue a barrier-sync to the L2 without executing a SYNC instruction, thus reducing the latency incurred for the sync. The L2-only sync provides a mechanism to guarantee that a uncached request does not pass previous cached requests in the L2 pipeline. For example, the L2-only SYNC can be used between a L2 HitWB cacheop and a subsequent uncached write to ensure that the uncached write does not pass the writeback from the L2.

The following sequence could be used to flush a cacheline from the L1 and L2 and then provide a sentinel to a consuming device as follows:

```
L1HitWB (flush L1 data to L2. will be globalized to all cores if coherent)
L2HitWB (flush L2 data to memory. CM2 ensures this does not pass the L1 HitWB)
L2-only SYNC (ensures subsequent uncached write does not pass L2HitWB)
uncached Store (sentinel to consuming device)
consuming device receives sentinel and reads memory
```

The L2-only sync is achieved by executing an Uncached store to an address that maps to the address region specified by the CM2's GCR\_L2\_ONLY\_SYNC\_BASE register. When the L2-only SYNC write is ready to be issued to the L2 pipeline the following actions occur:

- 1) Stop issuing new L2 requests until the L2 pipeline is empty and eviction queue is empty
- 2) The L2-only sync request is dropped and subsequent L2 requests continue.

Notice that the L2-only sync does not ensure any ordering in the coherent portion of the CM2.

The CM\_L2\_ONLY\_SYNC\_EN in bit 0 of the GCR\_L2\_ONLY\_SYNC\_BASE register must be set to a 1 for this feature to be enabled. The address match is performed on a 4KB boundary. An uncached write request address [31:12] that matches the address [31:12] in the GCR\_L2\_ONLY\_SYNC\_BASE will cause the CM2 to treat the uncached write request as an L2 only Sync.

The GCR\_L2\_ONLY\_SYNC\_BASE register is programmed through the Global Control Block Register Map located at offset 0x0070.

## 5.5.15 L2 Cache Initialization

The L2 cache controller contains minimal hardware initialization logic. It normally relies on software to fully initialize the L2 arrays. The registers used to support cache initialization are described in [Section 5.5.16, "L2 Cache CPO Interface"](#). For additional information, refer to the *CPO Registers* chapter of this manual.

The L1 data cache must be initialized during power-up or reset in order to place the lines of the cache in a known state. This is accomplished via the cache initialization routine, which is normally part of the boot code. For experienced user's, a sample boot code is shown in the following subsection.

### 5.5.15.1 init\_l2u Cache Initialization Routine

The following assembly provides an example initialization routine for the L2 cache.

```
LEAF(init_l2u)
    # Use CCA Override to allow cached execution of L2 init.
    # Check for CCA_Override_Enable by writing a one.
    lw r4_temp_data, 0x0008(r22_gcr_addr) # Read GCR_BASE register
    li r7_temp_mark, 0x50 # CM_DEFAULT_TARGET Memory
    # CCA Override Uncached enabled
    ins r4_temp_data, r7_temp_mark, 0, 8
    sw r4_temp_data, 0x0008(r22_gcr_addr)
    lw r4_temp_data, 0x0008(r22_gcr_addr) # GCR_BASE
    ext r4_temp_data, r4_temp_data, 4, 1 # Extract CCA_Override_Enable
    bnez r4_temp_data, done_l2 # Skip if CCA Override is implemented.
    nop
```

```

        b init_l2u
        nop
    END(init_l2u)

```

### 5.5.15.2 init\_l2c Cache Initialization Routine

The code in this function will be called from start.S after the L1 caches have been initialized. It will check to see if the core implements CCA Override. If it does, it will call the code to initialize the L2 cache.

```

LEAF(init_l2c)

    # Skip cached execution if CCA Override is not implemented.
    # If CCA override is not implemented the L2 cache would have already
    # been initialized when init_l2u was called.

    lw r4_temp_data, 0x0008(r22_gcr_addr) # Read GCR_BASE
    bnez r16_core_num, done_l2 # Only done from core 0.
    ext r4_temp_data, r4_temp_data, 4, 1 # CCA_Override_Enable
    beqz r4_temp_data, done_l2
    nop

END(init_l2c)

```

### 5.5.15.3 init\_L2u Initialization Routine Details

This section provides a detailed description of each line of code in the init\_l2u initialization routine.

The L2 cache is a system resource used by all cores in the system. Initialization of the L2 cache is done only by Core 0, because it only needs to be done once. The initialization of the L2 cache can be time consuming depending on its size. For example, a 256 KByte cache initializes quicker than an 8 MB cache.

The L2 cache initialization code executes faster if it is being run out of the instruction cache, so ideally the L2 initialization should be done after the L1 instruction cache in core 0 has been initialized. The instruction cache is a per-core resource and not initialized in the system initialization section of the code. Therefore, to be efficient and run the L2 cache initialization code out of the I-cache, the boot code tries to put off L2 cache initialization until the core 0 resources have been initialized. This can only be done if the L2 cache can be disabled before other cores are released to run this boot code. Otherwise there is a danger that other cores will use the L2 cache before it has been initialized by core 0.

The CCA override feature controls the cache attributes for the L2 cache. It allows for the disabling of the L2 cache by enabling the CCA override and setting the CCA to uncached. The CCA override works along with the L2 cache implementation.

The init\_l2u function tries to enable the CCA override and set the L2 cache to uncached in the GCR\_BASE register, thus disabling it. On systems that do not support CCA override, writes to the CCA override field have no effect, and reading back the GCR\_BASE register will not show the CCA override being set.

The code reads the GCR Base register.

```

    lw r4_temp_data, 0x0008(r22_gcr_addr) # GCR_BASE

```

The next 3 lines of code are used to enable CCA Override and set the L2 cache CCA to uncached.

```

    li r7_temp_mark, 0x50 # CM_DEFAULT_TARGET Memory
    # CCA Override Uncached enabled

```

```

ins r4_temp_data, r7_temp_mark, 0, 8
sw r4_temp_data, 0x0008(r22_gcr_addr)

```

Now the code reads back the GCR\_BASE register. If the CCA override bit is set, it means the code above worked, and the L2 cache is set to uncached. In this case, the code skips the initialization for now. The routine will be recalled later once the code is executing out of the L1 instruction cache. If not, the code branches to the `init_l2` function, which initializes the L2 cache.

```

lw r4_temp_data, 0x0008(r22_gcr_addr) # GCR_BASE
ext r4_temp_data, r4_temp_data, 4, 1 # CCA_Override_Enable
bnez r4_temp_data, done_l23 # Skip if CCA Override is implemented.
nop
b init_l2
nop

```

```

END(init_l2u)

```

#### 5.5.15.4 `init_L2c` Initialization Routine Details

This section provides a detailed description of each line of code in the `init_l2c` initialization routine. The code in this function is called from the `start.S` function after the L1 caches have been initialized. It checks to see if the core implements CCA Override. If it does, it calls the code to initialize the L2 cache.

In [Section 5.5.15.3](#) the code also checks to see if CCA override was implemented. If it was not, then it initialized the L2 cache while the code was executing in uncached mode, so there is no need to do it again here.

```

LEAF(init_l2c)

# Skip cached execution if CCA Override is not implemented.
# If CCA override is not implemented the L2 cache
# would have already been initialized when init_l2u was called.

lw r4_temp_data, 0x0008(r22_gcr_addr)      # GCR_BASE
bnez r16_core_num, done_l2                # Only done from core 0
ext r4_temp_data, r4_temp_data, 4, 1      # CCA_Override_Enable
beqz r4_temp_data, done_l23 nop

END(init_l2c)

```

## 5.5.16 L2 Cache CP0 Interface

The interAptiv core uses different CP0 registers for L2 cache operations.

**Table 5.13 L2 Cache CP0 Register Interface**

CP0 Registers	CP0 number
<i>Config2</i>	16.2
<i>ErrCtl</i>	26.0
<i>CacheErr</i>	27.0
<i>L23TagLo</i>	28.4
<i>L23DataLo</i>	28.5
<i>L23DataHi</i>	29.5

This section describes the base processor core CP0 registers that support the L2 cache. A complete description and bit assignments for each register listed is described in [Chapter 2, “CP0 Registers”](#).

### 5.5.16.1 Config2 Register (CP0 register 16, Select 2)

Asserting *Config2.L2B* (bit 12) enables the bypass-mode of the L2 cache. This bit is reflected on the *L2\_Bypass* output from the core. When L2 goes into bypass-mode, L2 responds by asserting *L2\_Bypassed* output, and the value of *L2\_Bypassed* is returned when *Config2.L2B* is read by software. Thus, reading this *Config2.L2B* bit does not read back what was written: it reflects the value of a signal sent back from the L2. The feedback signal, *L2\_Bypassed*, will reflect the previously written value with some implementation and clock ratio dependent delay.

Changing the value of *Config2.L2B* field in the middle of the normal operation may cause an unwanted loss of an OCP transaction in the L2 cache. For the safe transition into the L2 bypass-mode, an externalized SYNC before the MTC0 *Config2.L2B* is necessary to make sure all the pending transactions in L2 are completed. And, these instructions should run from the uncached space. It might be also a good idea to check if L2 is really in bypass-mode by reading the *Config2.L2B* field before moving onto the next instructions.

The *Config2.SS* field (bits 11:8) indicates the number of sets per way in the data cache. The interAptiv L2 cache supports from 512 up to 32768 sets per way, which is used to configure cache sizes from 256 KBytes to 8 MBytes.

The *Config2.SL* field (bits 7:4) indicates the line size for the L2 cache. The interAptiv L2 cache can be configured for a 32-byte or 64 byte line size.

The *Config2.SA* field (bits 3:0) indicates the set associativity for the L2 cache. The interAptiv L2 cache is fixed at 8-way set associative as indicated by a default value of 4 for this field.

For more information, refer to [Section 2.2.1.3, "Device Configuration 2 — Config2 \(CP0 Register 16, Select 2\)"](#).

### 5.5.16.2 Error Control Register (CP0 register 26, Select 0)

*ErrorControl.L2P* (bit 23) is used to enable L2 ECC checking and correction. This bit is read-only if the L2 has not been built with ECC/Parity support. Specific parity support is enabled using both *L2P* and *ErrorControl.PE* (bit 31) as described in [Table 5.14](#). *L2P* is also reflected on the *L2\_ECCEnable* output from the core.

These encodings were chosen such that legacy code which is unaware of L2P, will by default enable L2 ECC logic when it enables L1 parity. For more information, refer to [Section 2.2.6.10, "ErrCtl \(CP0 Register 26, Select 0\)"](#)

**Table 5.14 L2\_ECC\_Enable**

PE	L2P	L2_ECCEnable
1	0	1
1	1	0
0	0	0
0	1	1

### 5.5.16.3 Cache Error Register (CP0 register 27, Select 0)

When the L2 detects an uncorrectable error, CacheError.EC is set, identifying the exception as an L2 error. The Cache Error register stores information such as the cache way where the error was detected, the cache index of the double word in which the error was detected, the cache level at which the error was detected, if the tag RAM was involved, etc.

For more information, refer to [Section 2.2.6.11, "Cache Error — CacheErr \(CP0 Register 27, Select 0\)"](#).

### 5.5.16.4 L23TagLo Register (CP0 register 28, Select 4)

The L23TagLo register contains the contents of the L2 tag array at the location accessed by the L2 Index Load Tag cache-op. It is also used as the source register for the L2 Index Store Tag cache-op.

For more information, refer to [Section 2.2.6.7, "Level 2/3 Cache Tag Low — L23TagLo \(CP0 Register 28, Select 4\)"](#).

### 5.5.16.5 L23DataHi Register(CP0 register 29, Select 5) / L23DataLo Register(CP0 register 28, Select 5)

For the L2 Index Load Tag cache-op, L23DataHi and L23DataLo hold the contents of the doubleword from the L2 data array at the indexed location. (L23DataHi holds the most-significant word and L23DataLo holds the least-significant word). For the L2 Index Load WS cache-op, L23DataHi and L23DataLo each hold the ECC parity of the doubleword from the L2 data array at the indexed location.

These registers are also used for the source data for the Index Store Data cache-op. Finally, L23DataLo is used as the data source for the ECC to be written by the Index Store ECC cache-ops. For more details on the data returned by the L2 on a Index Load Tag/Data cache-op, please refer to [Section 5.6 “The CACHE Instruction”](#).

For more information on the L23DataLo register, refer to [Section 2.2.6.8, "Level 2/3 Cache Data Low — L23DataLo \(CP0 Register 28, Select 5\)"](#). For more information on the L23DataHi register, refer to [Section 2.2.6.9, "Level 2/3 Cache Data High — L23DataHi \(CP0 Register 29, Select 5\)"](#).

## 5.5.17 L2 Cache Operations

Cache-ops are used for control operations such as initialization, invalidation, eviction, etc. A brief description of the cache-ops implemented by the L2 are given below:

**Index Writeback Invalidate:** If the state of the cache line at the specified index is valid and dirty, the line is written back to the memory address specified by the cache tag. After that operation is completed, the state of the cache line is set to invalid. If the line is valid but not dirty, the state of the line is set to invalid.

**Index Load Tag:** The tag, valid, lock, dirty, parity and LRU bits for the cache line at the specified index are read. The doubleword indexed in the data RAM is also read.



**Index Load WS:** The LRU, dirty, and dirty parity bits for the cache line at the specified index are read. ECC for the doubleword indexed in the data RAM is also read.

**Hit Invalidate:** If the cache contains the specified address, the state of that cache line is set to invalid.

**Hit Writeback Inv:** If the cache contains the specified address and it is valid and dirty, the contents of that line are written back to main memory. After that operation is completed, the state of the cache line is set to invalid. If the line is valid but not dirty, the state of the line is set to invalid.

**Hit Writeback:** If the cache contains the specified address and it is valid and dirty, the contents of that line are written back to main memory. After the operation is completed, the state of the line is left valid, but the dirty state is cleared.

**Index Store Tag:** Write the tag for the cache line at the specified index.

**Index Store WS:** Write the WS array for the cache line at the specified index.

**Fetch And Lock:** If the cache contains the specified address, lock the line. If the cache does not contain the specified address, refill the line from main memory and then lock the line.

**Index Store Data:** Write the data and ECC for the cache line at the specified index. Proper ECC is generated for the written data and written into the ECC field.

**Index Store ECC:** Write the ECC for the cache line at the specified index.

Most CP0 instructions are used rarely, in code which is not timing-critical. But an OS which has to manage caches around I/O operations or otherwise may have to sit in a tight loop issuing hundreds of **cache** operations at a time, so performance can be important.

### 5.5.17.1 Bus Transaction Equivalence

When the base processor executes an L2 CACHE instruction, the operands and as well as data to be written to CP0 registers is transferred to and from L2. Index Load Tag and Index Load WS generate burst read transactions. All other L2 cache-ops generate single write transactions.

For 64 byte line configurations, bit 5 (the LSB of the Index field) is the selector to which 32 byte half of the 64 byte line is targeted (essentially it becomes an additional DW bit). For *tag* and *ws* type cache-ops, this bit is disregarded and cache-ops with either value of bit 5 impact the exact same tag or ws entry. For data type cache-ops, bit 5 selects which half of the 64 byte cache line is being accessed.

**Figure 5.5 Index Encoding for PB\_MAddr (1MB, 8-way)**

31	23 22	20 19	5 4 3 2	0
Unused	Way	Index	DW	Unused

### 5.5.17.2 Details of Cache-ops

Table 5.15 indicates the operation and behavior of the L2 cache for each cache-op.

**Table 5.15 Cache-ops**

Cache-op	Effective Address Operand Type	Operation
Index WB inv/ Indx Inv (OPCODE: 0)	INDEX	<ul style="list-style-type: none"> <li>• If the state of the cache line at the specified index is valid and dirty, the line is written back to the memory address specified by the cache tag. After that operation is completed, the state of the cache line is set to invalid.</li> <li>• If the line is valid but not dirty, the state of the line is set to invalid</li> <li>• The LRU bits are updated to Least-recently-used.</li> <li>• The dirty bits are updated to clean for that way.</li> </ul>
Index Load Tag (OPCODE: 1) ErrCtl.WST = 0	INDEX	<ul style="list-style-type: none"> <li>• The tag, valid, lock, and parity fields from the tag array for the cache line at the specified index are written into L23TagLo. Furthermore, the dirty bit from the WS array corresponding to the specified index is also written into L23TagLo. (First beat of return data)</li> <li>• For the first beat of return data, the two halves of the 64-bit data bus are identical.</li> <li>• The indexed doubleword is written into {L23DataHi, L23DataLo}. (2nd beat of return data)</li> <li>• ErrCtl.PO is treated as a don't care</li> <li>• The LRU bits are unchanged</li> </ul>
Index Load WS (OPCODE: 1) ErrCtl.WST = 1	INDEX	<ul style="list-style-type: none"> <li>• The dirty, dirty parity, and LRU fields from the WS array for the cache line at the specified index are written into L23TagLo. (First beat of return data)</li> <li>• For the first beat of return data, the two halves of the 64-bit data bus are identical.</li> <li>• The WS data at the indexed location is written into L23TagLo. (First beat of return data)</li> <li>• The indexed doubleword's ECC is written into {L23DataHi, L23DataLo}. (2nd beat of return data)</li> <li>• ErrCtl.PO is treated as a don't care</li> <li>• The LRU bits are unchanged</li> <li>• Data RAM:</li> <li>• The DW ECC to be read in the line is determined by <i>PB_MAddr[4:3]</i></li> </ul>
Index Store Tag (OPCODE: 2) ErrCtl.WST = 0	INDEX	<ul style="list-style-type: none"> <li>• The tag, valid, and lock fields in the Tag array at the indexed location are written from L23TagLo.</li> <li>• If ErrCtl.PO==1, the parity and total parity fields in the Tag array at the indexed location are written from L23TagLo.</li> <li>• If ErrCtl.PO==0, the parity and total parity fields in the Tag array at the indexed location are written with hardware generated values.</li> <li>• If valid==1, the LRU bits in the WS array are updated to make the indexed way most-recently-used. If valid==0, the LRU bits are updated with least-recently-used.</li> <li>• If valid==1, the dirty bit in the WS array at the indexed location is written from L23TagLo.</li> <li>• If valid==0, the dirty bit in the WS array at the indexed location is cleared.</li> <li>• The dirty parity bit in the WS array at the indexed location is written with the correct hardware generated values.</li> </ul>
Index Store WS (OPCODE: 2) ErrCtl.WST = 1	INDEX	<ul style="list-style-type: none"> <li>• The dirty and LRU fields for all 8 ways of the WS array at the indexed location are written from L23TagLo</li> <li>• If ErrCtl.PO==1, the dirty parity fields for all 8 ways of the WS array at the indexed location are written from L23TagLo</li> <li>• If ErrCtl.PO==0, the dirty parity fields for all 8 ways of the WS array at the indexed location are written with hardware generated values</li> </ul>

**Table 5.15 Cache-ops (continued)**

Cache-op	Effective Address Operand Type	Operation
Index Store Data (OPCODE: 3) ErrCtl.WST = 0	INDEX	<ul style="list-style-type: none"> <li>The doubleword in the data array at the indexed location and doubleword offset is written from {L23DataHi, L23DataLo} regardless of the PB_MDataByteEn value.</li> <li>The Parity/ECC field in the data array at the indexed location and doubleword offset is written with a hardware generated value.</li> <li>The LRU bits in the WS array are updated to make the indexed way most-recently-used.</li> </ul>
Index Store ECC (OPCODE: 3) ErrCtl.WST = 1	INDEX	<ul style="list-style-type: none"> <li>The Parity/ECC field in the data array at the indexed location and doubleword offset is written from L23DataLo[7:0].</li> <li>The LRU bits in the WS array are updated to make the indexed way most-recently-used.</li> </ul>
HIT Inv (OPCODE: 4)	ADDRESS	<ul style="list-style-type: none"> <li>If the address is not contained in L2, nothing happens.</li> <li>If the address hits in L2, it is invalidated and the dirty bit is cleared.</li> <li>If any arrays are written, the appropriate parity fields are updated by hardware.</li> </ul>
HIT WB Inv (OPCODE: 5)	ADDRESS	<ul style="list-style-type: none"> <li>If the address is not contained in L2, nothing happens.</li> <li>If the address hits in L2, and it is dirty, the line is written back to main memory. It is then invalidated and the dirty bit is cleared.</li> <li>If the address hits in L2, and it is clean, it is invalidated.</li> <li>If any arrays are written, the appropriate parity fields are updated by hardware.</li> </ul>
HIT WB (OPCODE: 6)	ADDRESS	<ul style="list-style-type: none"> <li>If the address is not contained in L2, nothing happens.</li> <li>If the address hits in L2, and it is dirty, the line is written back to main memory and the dirty bit is cleared.</li> <li>If the address hits in L2, and it is clean, nothing happens.</li> <li>If any arrays are written, the appropriate parity fields are updated by hardware.</li> </ul>
Fetch and Lock (OPCODE: 7)	ADDRESS	<ul style="list-style-type: none"> <li>If the address is not contained in L2, the line is refilled. The refilled line is then locked in the cache. The LRU bits in the WS array are updated to make the fetched way most-recently-used. The Dirty bit and the dirty parity bit are set to clean.</li> <li>On a hit the line is locked and the operation retires. The LRU bits or the dirty bits are not affected.</li> </ul>

### 5.5.17.3 Sync in L2

A Sync operation can be used to guarantee ordering of transactions. The L2 ensures that all transactions preceding a Sync request will be ordered in front of transactions received after the Sync request. Within the L2 only requests are ordered, not responses, i.e., there is no guarantee of the ordering between a read response vs. the Sync.

One example of the use of a Sync involves cache operations. Normally, the L2 does not guarantee the ordering between a cache operation, such as a Hit-Writeback-Invalidate, vs. an subsequent uncached request. If the software wants to ensure that any writes on the system interface due to the Hit-Writeback-Invalidate will be ordered in front of a subsequent uncached write, then a Sync must be issued between the cache operation and uncached write. Note that in order for a core to externalize a Sync request, *Config7.ES* bit must be set before the sync instruction.

The L2 issues a response to a Sync after all 3 of the following have completed:

- All previous requests have cleared the L2 pipeline
- The L2 has issued all requests to the system interface that are required by previous transactions, such as uncached requests, cache operations, cache misses, evictions, or previous Syncs.

- If the downstream system can take a sync OCP transaction ( $L2\_SyncTxEn=1$ ), it will externalize the sync transaction to the system once the above criteria has been satisfied. When the Sync response is received from the system interface, the L2 will return a Sync response to the processor interface.

#### 5.5.17.4 L2 Cache Fetch and Lock

In the L2 cache, each line in a way can be locked independently. If a line is locked it will not be evicted. Software is not allowed to lock all available ways at the same cache index, since L2 would be unable to refill any other addresses at that index.

If the requested address is not contained in the L2 cache, the line is refilled and then locked in the cache. The LRU bits in the WS array are updated to make the fetched way most-recently-used. The dirty bit and the dirty parity bit are set to clean.

On a hit the L2 cache line is locked and the operation retires. The LRU bits or the dirty bits are not affected.

### 5.5.18 L2 Cache Error Management

This section describes ECC, parity, and bus error support for the L2 cache.

#### 5.5.18.1 ECC/Parity Support

If ECC support is selected at build time, and this support is enabled via software by setting the *ErrCtl.L2EccEn* bit in the Error Control register (CP0 register 26, Select 0), then the tag and the data arrays are protected with single-error correction logic as well as double-error detection logic.

If Parity support is selected at build time, and this support is enabled via software by setting the *ErrCtl.pE* bit in the Error Control register (CP0 register 26, Select 0), then the tag and the data arrays are protected with single-error correction logic.

The Way Select RAM is protected with single-error detection logic. Correctable errors are not reported to the processor, but uncorrectable errors are reported to the processor. If ECC/Parity support is either not selected at build time or disabled, then no errors are detected (or corrected) on any of the cache arrays.

The ECC logic uses Hamming's error correcting code. In the data array, each 64-bit doubleword is independently ECC protected. This requires 8 parity bits per doubleword. The tag array requires 6 parity bits.

To perform a single detection and correction the parity bits are placed at  $2^n$  locations among the data bits. The bits at different locations are then grouped together. The grouping is done by analyzing the binary weights of the particular location.

For example, to protect 8 data bits, 4 parity bits are needed which will be placed as below:

**Table 5.16 Parity Bit Distribution**

Bit Location	12	11	10	9	8	7	6	5	4	3	2	1
Parity and data bits	d7	d6	d5	d4	p3	d3	d2	d1	p2	d0	p1	p0

Note that Bit location 0 does not exist.

The binary weight of bit location 3 is  $2^0$  and  $2^1$ , which is derived from its binary value 0011b. Therefore, bit location 3 falls in group g0 and g1. Similarly, Bit location 11 falls into groups g0, g1 and g3.

Parity bit p0 will belong to g0 and its value will be generated such that g0 will have an even parity. Similarly all other parity bits are generated such that their respective group ends up in even parity.

This sharing of binary weights across groups enables the L2 to determine precisely which data or parity bit was in error. That is achieved by recreating the parity bits from the data read from the memory and XORing it with the parity bits read from the memory. The XORed value, or the syndrome, points to the bit in error. Once this error is detected the L2 corrects it. A value of zero on the syndrome indicates that there was no error in the parity and data bits.

To achieve double bit error detection an even parity is generated across the parity and data bits, which is termed as the total parity bit. The total parity bit will be flipped in case of a single bit error, whereas for a double bit error it will remain the same. The syndrome along with the total parity bit is then used to detect a double bit error.

The WSRAM's dirty bits are protected, whereas the LRU bits are not. For each dirty bit there is one more bit added called the dirty parity bit. The value of the dirty parity bit enforces even parity protection.

### 5.5.18.2 Tag and Data Formats

#### **Tag Array Format**

The width of the tag in an 8 way 128 MB cache is 18 bits per way. The data array format is as shown in [Table 5.17](#).

**Table 5.17 Tag Array Format for a 8 Way 128 MB Cache**

Bit position	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
Content	TP	L	V	d17	d16	d15	d14	d13	d12	d11	p4	d10	d9	d8	d7	d6	d5	d4	p3	d3	d2	d1	p2	d0	p1	p0

Where, d0-d17 : Tag  
V : Valid bit  
L : Lock bit  
p0-p4 : parity bits  
TP : Total parity bit

For larger caches, the width of the tag reduces. In that case, the upper data bits are ignored from the calculation as appropriate.

#### **Data Array Format**

The data array format is as shown in [Table 5.18](#).

**Table 5.18 Data Array Format**

Bit position	72	71:65	64	63:33	32	31:17	16	15:9	8	7:5	4	3	2	1
Content	TP	[63:57]	p6	[56:26]	p5	[25:11]	p4	[10:4]	p3	[3:1]	p2	[0]	p1	p0

### 5.5.18.3 Cache Parity Error Handling

The three types of memory arrays in the L2 have an option for parity. If selected, this option provides single bit correction and double bit detection of the tag rams and data rams.

- The Tag RAM coverage is for each way.
- The Data RAM coverage is for each way and each double-word in each way.
- The Way Select RAM has parity for each dirty bit. A correctable bit failure is corrected and no notification of this event is present at the L2 pins.

The five types of ECC/parity errors are handled internally as follows.

### ***Single Bit (correctable) Tag RAM Error in a Way***

The corrected tag value is written back into the tag ram by replaying the request.

### ***Single Bit (correctable) Data RAM Error in a Dword of a Way***

The corrected data value is written back into the data ram by replaying the request. This may occur due to a read that hits or a partial write where a dword in the way has a single bit failure.

### ***Double Bit (uncorrectable) Tag RAM Error in a Way***

An uncorrectable tag ram failure kills the request in the L2 pipeline. A write request is dropped and a read request is treated as a hit to an arbitrary way.

### ***Double Bit (uncorrectable) Data RAM Error in a Dword of a Way***

For a read hit, the uncorrected data is returned.

### ***Parity Error (uncorrectable) on a Dirty Bit in the Wsram***

When a dirty parity error is detected, the L2 treats the state as dirty by default. This means that a victim line being evicted due to either allocation or invalidation by a request might not have really have needed to be written.

## **5.5.18.4 Multiple Uncorrectable Errors**

This error is reported when more than one uncorrectable error is being reported on the same L2 clock cycle. Since double-bit Tag RAM errors, double-bit Data RAM error, and parity bit errors in the Way Select RAM are each reported in different L2 pipeline stages, this assertion indicates that different requests have encountered uncorrectable requests. In other words, if a single request suffers all three uncorrectable errors, the error will be reported three times.

## **5.5.18.5 Bus Error Handling**

Bus errors are never originated by the L2. However, bus errors may be received from the system on an OCP read from the L2 to the system. The error is indicated when the read-data is returned back to the L2. The L2 propagates the bus error when returning data to the processor or CM2.

If a bus error is received on a 64-byte burst read to the system, the L2 signals the bus error for the processor read that originated the request. If the L2 receives a subsequent read to the same 64-byte cache line before all the data has been received from memory for the previous request, the new request also receives a bus error response.

In general, a bus error reported in a system response due to a processor/CM request is considered to be reporting the entire cache line as having a bus error. However, if the original request is satisfied before the L2 detects the system bus error, then the response to the processor/CM will not have a bus error.

There is no capability for signalling bus errors on writes.

## 5.6 The CACHE Instruction

The L1 instruction, L1 data, and L2 caches in the interAptiv core each support the CACHE instruction, which allows users to manipulate the contents of the Data and Tag arrays, including the locking of individual cache lines. The behavior of the CACHE instruction is identical for both the L1 instruction and data caches.

### 5.6.1 Decoding the Type of Cache Operation

The type of cache operation performed is encoded using a combination of the 5-bit *op* field of the CACHE instruction, and selected bits from the *ErrCtl* register (CP0 Register 26, Select 0). In addition to performing operations on the caches themselves, there are other CACHE operations that are performed on internal memories such as the way selection RAM, the scratch pad RAM, and the Dirty Bit RAM. The *ErrCtl* bits determine the type internal memory where the CACHE operation will be performed.

The selected bits of the *ErrCtl* register used to determine the type of CACHE operation are as follows:

- Bit 29, *WST*: If this bit is set, execution of a **cache IndexLoadTag** or **cache IndexStoreTag** instruction reads or writes the cache's internal *way-selection RAM* instead of the cache tags.
- Bit 28, *SPR*: If this bit is set, index-type cache instructions work on the data scratch pad (*DSPRAM*) and instruction scratch pad (*ISPRAM*), if implemented. Read the *ConfigDSP* and *ConfigISP* bits to determine if the associated scratch pad RAM is present.
- Bit 21, *DYT*: Setting this bit allows **cache** load/store data operations to work on the "dirty array" associated with the L1 data cache.

### 5.6.2 CACHE Instruction Opcodes

Refer to the implementation-specific CACHE instruction at the back of this manual for a list of CACHE instruction opcodes.

### 5.6.3 Way Selection RAM Encoding

The CACHE Index Load Tag and Index Store Tag instructions can be used to read and write the Way Select (WS) RAM by setting the *WST* bit in the *ErrCtl* register. Similarly, the *SPR* bit in the *ErrCtl* register causes the **Index Load Tag** and **Index Store Tag** instructions to read the pseudo-tags associated with the scratch-pad RAM array. Note that when the *WST* and *SPR* bits are zero, the CACHE index instructions access the cache Tag array.

Not all values of the WS field are valid for defining the order in which the ways are selected. This is only an issue, however, if the WS RAM is written after the initialization (invalidation) of the Tag array. Valid WS field encodings for way selection order is shown in [Table 5.19](#).

**Table 5.19 Way Selection Encoding, 4 Ways**

Selection Order <sup>1</sup>	WS[5:0]	Selection Order	WS[5:0]
0123	000000	2013	100010
0132	000001	2031	110010
0213	000010	2103	100110
0231	010010	2130	101110
0312	010001	2301	111010
0321	010011	2310	111110

**Table 5.19 Way Selection Encoding, 4 Ways (continued)**

<b>Selection Order<sup>1</sup></b>	<b>WS[5:0]</b>	<b>Selection Order</b>	<b>WS[5:0]</b>
1023	000100	3012	011001
1032	000101	3021	011011
1203	100100	3102	011101
1230	101100	3120	111101
1302	001101	3201	111011
1320	101101	3210	111111

1. The order is indicated by listing the least-recently used way to the left and the most-recently used way to the right, etc.



## 5.7 Write Back Buffer

The Bus Interface Unit (BIU) includes a Write Back Buffer (WBB) that holds data from the L1 cache that is going to memory. This includes evictions from the data cache, uncached stores, and uncached accelerated stores. The WBB consists of eight 32-byte entries. The WBB also holds L2 **CACHE** instructions that are to be sent out on the bus. The WBB gathers uncached accelerated (UCA) stores to allow full line burst writes.

WBB entries are ‘flushed’ under a variety of conditions. When a buffer is flushed, the write command is queued in the BIU and the WBB entry will not accept any more activity until the data has been written to the bus and the buffer is freed up. Some flush conditions are shown here:

- Uncached (non-accelerated) stores flush immediately
- L2 CACHE instruction commands are also flushed immediately
- Entries for L1 data cache evictions are flushed when all 4 double-words (32B) of data have been gathered

When coherence is enabled, the CPU is the ‘owner’ of a cache line until the self-intervention for the writeback request has been seen. The WBB entry cannot be deallocated until that point so that the CPU can respond with the data if another CPU requests it. The WBB is also used for staging data responses to interventions. To avoid deadlock, one WBB entry must be reserved for this purpose.



## Exceptions and Interrupts

The interAptiv core receives exceptions from a number of sources, including arithmetic overflows, misses in the translation lookaside buffer (TLB), I/O interrupts, and system calls. When the CPU detects an exception, the normal sequence of instruction execution is suspended and the processor enters kernel mode, disables interrupts, loads the *Exception Program Counter (EPC)* register with the location where execution can restart after the exception has been serviced, and forces execution of a software exception handler located at a specific address.

The software exception handler saves the context of the processor, including the contents of the program counter, the current operating mode, and the status of the interrupts (enabled or disabled). This context is saved so it can be restored when the exception has been serviced.

Exceptions may be precise or imprecise. Precise exceptions are those for which the *EPC* can be used to identify the instruction that caused the exception. For precise exceptions, the restart location in the *EPC* register is the address of the instruction that caused the exception or, if the instruction was executing in a branch delay slot (as indicated by the *BD* bit in the *Cause* register), the address of the branch instruction immediately preceding the delay slot. Imprecise exceptions, on the other hand, are those for which no return address can be identified. Bus error exceptions and CP2 exceptions are examples of imprecise exceptions.

This chapter contains the following sections:

- [Section 6.1 “Exception Conditions”](#)
- [Section 6.2 “Exception Priority”](#)
- [Section 6.3 “Exception Vector Locations”](#)
- [Section 6.4 “General Exception Processing”](#)
- [Section 6.5 “Debug Exception Processing”](#)
- [Section 6.6 “Exception Descriptions”](#)
- [Section 6.7 “Exception Handling and Servicing Flowcharts”](#)
- [Section 6.8 “Interrupts”](#)

### 6.1 Exception Conditions

When an exception condition occurs, the instruction causing the exception and all those that follow it in the pipeline are cancelled. Accordingly, any stall conditions and any later exception conditions that may have referenced this instruction are inhibited.

When the exception condition is detected on an instruction fetch, the CPU aborts that instruction and all instructions that follow. When the instruction graduates, the exception flag causes it to write various CP0 registers with the exception state, change the current program counter (PC) to the appropriate exception vector address, and clear the exception bits of earlier pipeline stages.

For most types of exceptions, this implementation allows all preceding instructions to complete execution and prevents all subsequent instructions from completing. Thus, the value in the *EPC* (or *ErrorEPC* for errors or *DEPC* for debug exceptions) is sufficient to restart execution. It also ensures that exceptions are taken in program order. An instruction taking an exception may itself be aborted by an instruction further down the pipeline that takes an exception in a later cycle.

Imprecise exceptions are taken after the instruction that caused them has completed and potentially after following instructions have completed.

## 6.2 Exception Priority

Table 6.1 contains a list and a brief description of all exception conditions. The exceptions are listed in the order of their relative priority, from highest priority (Reset) to lowest priority (Load/store bus error). When several exceptions occur simultaneously, the exception with the highest priority is taken.

**Table 6.1 Priority of Exceptions**

Exception	Description
Reset	Assertion of <i>SI_Reset</i> signal.
DSS	EJTAG Debug Single Step.
DINT	EJTAG Debug Interrupt. Caused by the assertion of the external <i>EJ_DINT</i> input, or by setting the <i>EjtagBrk</i> bit in the <i>ECR</i> register.
DDBLImpr/DDBSImpr	Debug Data Break Load/Store. Imprecise.
NMI	Asserting edge of <i>SI_NMI</i> signal.
Interrupt	Assertion of unmasked hardware or software interrupt signal.
Deferred Watch	Deferred Watch (unmasked by $K DM \rightarrow !(K DM)$ transition).
DIB	EJTAG debug hardware instruction break matched.
WATCH	A reference to an address in one of the watch registers (fetch).
AdEL	Fetch address alignment error. Fetch reference to protected address.
TLBL	Fetch TLB miss . Fetch TLB hit to page with $V=0$ . This exception is at the same priority level as MPUL below and is only taken if an MMU is implemented.
MPUL	Fetch MPU miss. This exception is at the same priority level as TLBL above and is only taken in an MPU is implemented.
TLBXI	TLB Execute Inhibit. Occurs when there is an execute access from a page table whose XI bit is set. This exception is at the same priority level as MPUL below and is only taken if an MMU is implemented.
I-cache Error	Parity error on I-cache access.
DBE	From Instruction Fetch Unit (IFU) instruction cache ops.
D\$/L2\$ Error	Both of these errors are signaled as data cache errors.
DBE	Load or store bus error.
IBE	Instruction fetch bus error.
DBp	EJTAG Breakpoint (execution of SDBBP instruction).

**Table 6.1 Priority of Exceptions (continued)**

<b>Exception</b>	<b>Description</b>
Sys (Execution exception)	Execution of SYSCALL instruction. Note that all of the execution exceptions have the same priority.
Bp (Execution exception)	Execution of BREAK instruction. Note that all of the execution exceptions have the same priority.
CpU (Execution exception)	Execution of a coprocessor instruction for a coprocessor that is not enabled. Note that all of the execution exceptions have the same priority.
CEU (Execution exception)	Execution of a CorExtend instruction modifying local state when CorExtend is not enabled. Note that all of the execution exceptions have the same priority.
DSPDis (Execution exception)	DSP ASE state disabled.
RI (Execution exception)	Execution of a Reserved Instruction. Note that all of the execution exceptions have the same priority.
FPE (Execution exception)	Floating Point exception. Note that all of the execution exceptions have the same priority.
C2E (Execution exception)	Coprocessor 2 unusable exception. Note that all of the execution exceptions have the same priority.
ISI (Execution exception)	Implementation specific Coprocessor 2 exception. Note that all of the execution exceptions have the same priority.
Ov (Execution exception)	Execution of an arithmetic instruction that overflowed. Note that all of the execution exceptions have the same priority.
Tr (Execution exception)	Execution of a trap (when trap condition is true). Note that all of the execution exceptions have the same priority.
MT_ov (Execution exception)	Thread overflow condition, where a TC allocation request cannot be satisfied. Note that all of the execution exceptions have the same priority.
MT_under (Execution exception)	Thread underflow condition, where the termination and deallocation of a thread leaves no TCs activated on a VPE. Note that all of the execution exceptions have the same priority.
MT_invalid (Execution exception)	Invalid qualifier condition, where a YIELD instruction specifies an invalid condition for resuming execution. Note that all of the execution exceptions have the same priority.
MT_yield_sched (Execution exception)	YIELD scheduler exception condition, where a valid YIELD instruction could have caused a rescheduling of a TC, and the YIELD intercept bit is set. Note that all of the execution exceptions have the same priority.
DDBL / DDBS	EJTAG Data Address Break (address only).
WATCH	A reference to an address in one of the watch registers (data).
AdEL	Load address alignment error. Load reference to protected address.
AdES	Store address alignment error. Store to protected address.
TLBL	Load TLB miss. Load TLB hit to page with V=0. This exception is at the same priority level as the MPUL load miss below and is only taken if an MMU is implemented.
MPUL	Load MPU miss. This exception is at the same priority level as the TLBL load miss above and is only taken if an MPU is implemented.

**Table 6.1 Priority of Exceptions (continued)**

Exception	Description
TLBS	Store TLB miss. Store TLB hit to page with V=0. This exception is at the same priority level as the MPUS load miss below and is only taken if an MMU is implemented.
MPUS	Store MPU miss. This exception is at the same priority level as the TLBS load miss above and is only taken if an MPU is implemented.
TLBRI	TLB Read Inhibit. Occurs when there is an attempt to access a page table whose RI bit is set.
TLB Mod	Store to TLB page with D = 0.
MT_GSS (Thread exception)	Gating storage scheduler exception, where a gating storage load or store would have been blocked and caused a rescheduling or a TC, and the GS intercept bit is set. Note that both thread exception have the same priority.
MT_GS (Thread exception)	Gating storage exception condition, where implementation-dependent logic associated with gating or inter-thread communication (ITC) storage requires software intervention. Note that both thread exception have the same priority.

## 6.3 Exception Vector Locations

The location of the exception vector in the interAptiv core depends on the operating mode. If the core is in the legacy setting, the exception vector location is the same as in previous MIPS processors. However, if the core is configured for Enhanced Virtual Address (EVA), the exception vector can effectively be placed anywhere within kernel address space.

The *SI\_EVAReset* pin determines the addressing scheme and whether the device boots up in the legacy setting or the EVA setting. The legacy setting is defined as having the traditional MIPS virtual memory map used in previous generation processors. The EVA setting places the device in the enhanced virtual address configuration, where the initial size and function of each segment in the virtual memory map is determined from the segmentation control registers (*SegCtl0* - *SegCtl2*).

If the *SI\_EVAReset* pin is deasserted at reset, the interAptiv core comes up in the legacy configuration and hardware takes the following actions:

- The *CONFIG5.K* bit becomes read-write and is programmed by hardware to a value of 0 to indicate the legacy configuration. In this case, the cache coherency attributes for the kseg0 segment are derived from the *Config.K0* field as described in the previous subsection. In addition to selecting the location of the cache coherency attributes, the *CONFIG5.K* bit also causes hardware to generate two boot exception overlay segments, one for kseg0 and one for kseg1, as described in [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#).
- Hardware programs the CP0 memory segmentation registers (*SegCtl0* - *SegCtl2*) for the legacy setting. Note that these registers are new in the interAptiv core and are not used by legacy software. However, they are used by hardware during normal operation, so their default values should not be changed.

If the *SI\_EVAReset* pin is asserted at reset, the interAptiv core comes up in the EVA configuration (default is *xkseg0* space = 3 GB) and hardware takes the following actions:

- The *CONFIG5.K* bit becomes read-only and is forced to a value of 1 to indicate the EVA configuration. In this case, the *CONFIG.K0* field is ignored and is no longer used to determine the kseg0 cache coherency attributes (CCA). Rather, the values in bits 2:0 (segments 0, 2, and 4) and bits 18:16 (segments 1, 3, and 5) of the *SegCtl0* - *SegCtl2* registers are used to define the CCA for each memory segment. In this case, hardware generates only one BEV overlay segment as described in [Section 3.6, "Boot Exception Vector Relocation in Kernel Mode"](#).

- Hardware sets the CP0 memory segmentation registers (*SegCtl0* - *SegCtl2*) for the EVA configuration.

When the *SI\_UseExceptionBase* pin is 0 and the *Config5.K* bit is cleared, the device is in legacy mode. In this mode the exception vector location defaults of 0xBFC0\_0000 and the *SI\_ExceptionBase[31:12]* pins are ignored.

When the *SI\_UseExceptionBase* pin is 1 and the *Config5.K* bit is cleared, the device is still in legacy mode, but the *SI\_ExceptionBase[29:12]* pins are used to indicate the location of the exception vector. Bits 31:30 are forced to a value of 2'b10, placing the exception vector somewhere in kseg0/kseg1 space.

If the *Config5.K* bit is set, the device is in EVA mode. In this case the *SI\_UseExceptionBase* pin is ignored and the *SI\_ExceptionBase[31:12]* pins are used to derive the location of the exception vector.

The function of the *Config5.K* bit and the *SI\_UseExceptionBase* pin is shown in [Table 6.2](#). For more information on EVA mode, refer to the MMU chapter.

**Table 6.2 *SI\_UseExceptionBase* Pin and CONFIG5.K Encoding**

CONFIG5.K Bit	<i>SI_UseExceptionBase</i> Pin	Condition	Action
0	0	Legacy Mode <i>SI_ExceptionBase[31:12]</i> pins are not used.	Use default BEV location of 0xBFC0_0000.
0	1	Legacy Mode Use only <i>SI_ExceptionBase[29:12]</i> for the BEV base location. Bits 31:30 are forced to a value of 2'b10 to put the BEV vector into KSEG0/KSEG1 virtual address space.	The BEV location is determined as follows: <i>SI_ExceptionBase[31:12]</i> = 2'b10, <i>SI_ExceptionBase[29:12]</i> pins, 12'b0 Bits 31:30 are forced to a value of 2'b10 to put the BEV vector into KSEG0/KSEG1 virtual address space.
1	Don't care	EVA Mode Use <i>SI_ExceptionBase[31:12]</i> pins.	The <i>SI_ExceptionBase[31:12]</i> pins are used directly to derive the BEV location. The <i>SI_UseExceptionBase</i> pin is ignored.

Another degree of flexibility in the selection of the vector base address, for use when *StatusBEV* equals 1, is provided via a set of input pins, *SI\_UseExceptionBase*, *SI\_ExceptionBase[31:12]*, and *SI\_ExceptionBaseMask[27:20]*.

In the legacy setting, when the *SI\_UseExceptionBase* pin is 0, the Reset, Soft Reset, NMI, and EJTAG Debug exceptions are vectored to a specific location, as shown in [Table 6.3](#). Addresses for all other exceptions are a combination of a vector offset and a vector base address. In the interAptiv core, software is allowed to specify the vector base address via the *EBase* register for exceptions that occur when *StatusBEV* equals 0. [Table 6.3](#) shows the vector base address when the core is in legacy setting and the *SI\_UseExceptionBase* pin is 0.

[Table 6.4](#) shows the vector base addresses when the core is in legacy setting and the *SI\_UseExceptionBase* equals 1. As can be seen in [Table 6.4](#), when *SI\_UseExceptionBase* equals 1, the exception vectors for cases where *StatusBEV* = 0 are not affected.

**Table 6.3 Exception Vector Base Addresses — Legacy Mode,  $SI\_UseExceptionBase = 0$**

Exception	Status <sub>BEV</sub>	
	0	1
Reset, NMI	0xBFC0.0000	
EJTAG Debug (with $ProbEn = 0$ , in the EJTAG_Control_register and $DCR.RDVec=0$ )	0xBFC0.0480	
EJTAG Debug (with $ProbEn = 0$ , in the EJTAG_Control_register and $DCR.RDVec=1$ )	<i>DebugVectorAddr</i> [31:7]    2b0000000	
EJTAG Debug (with $ProbEn = 1$ in the EJTAG_Control_register)	0xFF20.0200	
Cache Error	<i>EBase</i> <sub>31 30</sub>    1    <i>EBase</i> <sub>28 12</sub>    0x000 Note that <i>EBase</i> <sub>31 30</sub> have the fixed value of 2'b10	0xBFC0.0300
Other	<i>EBase</i> <sub>31 12</sub>    0x000 Note that <i>EBase</i> <sub>31 30</sub> have the fixed value of 2'b10	0xBFC0.0200
‘  ’ denotes bit string concatenation		

In legacy mode, when the  $SI\_UseExceptionBase$  pin is 0, the Reset, Soft Reset, NMI, and EJTAG Debug exceptions are vectored to a specific location, as shown in Table 6.4.

**Table 6.4 Exception Vector Base Addresses — Legacy Mode,  $SI\_UseExceptionBase = 1$**

Exception	Status <sub>BEV</sub>	
	0	1
Reset, NMI	0b10    <i>SI_ExceptionBase</i> [29:12]    0x000	
EJTAG Debug (with $ProbEn = 0$ in the EJTAG_Control_register and $DCR.RDVec=0$ )	0b10    <i>SI_ExceptionBase</i> [29:12]    0x480	
EJTAG Debug (with $ProbEn = 0$ in the EJTAG_Control_register and $DCR.RDVec=1$ )	<i>DebugVectorAddr</i> [31:7]    2b0000000	
EJTAG Debug (with $ProbEn = 1$ in the EJTAG_Control_register)	0xFF20.0200	
Cache Error	<i>EBase</i> <sub>31 30</sub>    1    <i>EBase</i> <sub>28 12</sub>    0x000 Note that <i>EBase</i> <sub>31 30</sub> have the fixed value 2'b10. Exception vector resides in kseg1.	0b101    <i>SI_ExceptionBase</i> [28:12]    0x300 Exception vector resides in kseg1.



**Table 6.4 Exception Vector Base Addresses — Legacy Mode, *SI\_UseExceptionBase* = 1 (continued)**

Exception	Status <sub>BEV</sub>	
	0	1
Other	$EBase_{31:12} \parallel 0x000$ Note that $EBase_{31:30}$ have the fixed value 2'b10 Exception vector resides in kseg0/kseg1.	$0b10 \parallel SI\_ExceptionBase[29:12] \parallel 0x200$ Exception vector resides in kseg0/kseg1.
‘  ’ denotes bit string concatenation		

Table 6.5 shows the offsets from the vector base address as a function of the exception. Note that the IV bit in the Cause register causes interrupts to use a dedicated exception vector offset, rather than the general exception vector. Table 6.25 (on page 344) shows the offset from the base address in the case where Status<sub>BEV</sub> = 0 and Cause<sub>IV</sub> = 1. Table 6.7 combines these three tables into one that contains all possible vector addresses as a function of the state that can affect the vector selection. To avoid complexity in the table, it is assumed that IntCtl<sub>VS</sub> = 0.

**Table 6.5 Exception Vector Offsets**

Exception	Vector Offset
TLB Refill, EXL = 0	0x000
General Exception	0x180
Interrupt, Cause <sub>IV</sub> = 1	0x200 (In Release 3 implementations, this is the base of the vectored interrupt table when Status <sub>BEV</sub> = 0)
Reset, NMI	None (uses reset base address)

In EVA mode, when the *SI\_UseExceptionBase* pin is ignored and the Reset, Soft Reset, NMI, and EJTAG Debug exceptions are vectored to a location determined by the programming of the three Segment Control registers (*SegCtl0* - *SegCtl2*), as shown in Table 6.6.

**Table 6.6 Exception Vector Base Addresses — EVA Mode**

Exception	Status <sub>BEV</sub>	
	0	1
Reset, NMI	$SI\_ExceptionBase[31:12] \parallel 0x000$	
EJTAG Debug (with ProbEn = 0 in the EJTAG_Control_register and DCR.RDVec=0)	$SI\_ExceptionBase[31:12] \parallel 0x480$	
EJTAG Debug (with ProbEn = 0 in the EJTAG_Control_register and DCR.RDVec=1)	$DebugVectorAddr[31:7] \parallel 2b0000000$	
EJTAG Debug (with ProbEn = 1 in the EJTAG_Control_register)	0xFF20.0200	
Cache Error	$EBase_{31:12} \parallel 0x000$	$SI\_ExceptionBase[31:12] \parallel 0x300$ (Forced uncached)
Other	$EBase_{31:12} \parallel 0x000$	$SI\_ExceptionBase[31:12] \parallel 0x200$
‘  ’ denotes bit string concatenation		

Table 6.7 Exception Vectors

Exception	Config5 <sub>K</sub>	SI_UseExceptionBase	Status <sub>BEV</sub>	Status <sub>EXL</sub>	Cause <sub>V</sub>	EJTAG ProbEn	Vector (IntCtl <sub>VS</sub> = 0)
Reset, NMI	0	0	x	x	x	x	0xBFC0.0000
Reset, NMI	0	1	x	x	x	x	2'b10    <i>SI_ExceptionBase</i> [29:12]    0x000
Reset, NMI	1	x	x	x	x	x	<i>SI_ExceptionBase</i> [31:12]    0x000
EJTAG Debug	0	0	x	x	x	0	0xBFC0.0480 (if <i>DCR.RDVec</i> =0) <i>DebugVectorAddr</i> [31:7]    2b0000000 (if <i>DCR.RDVec</i> =1)
EJTAG Debug	0	1	x	x	x	0	2'b10    <i>SI_ExceptionBase</i> [29:12]    0x480 (if <i>DCR.RDVec</i> =0) <i>DebugVectorAddr</i> [31:7]    2b0000000 (if <i>DCR.RDVec</i> =1)
EJTAG Debug	1	x	x	x	x	0	<i>SI_ExceptionBase</i> [31:12]    0x480 (if <i>DCR.RDVec</i> =0) <i>DebugVectorAddr</i> [31:7]    2b0000000 (if <i>DCR.RDVec</i> =1)
EJTAG Debug	x	x	x	x	x	1	0xFF20.0200
TLB Refill	x	x	0	0	x	x	<i>EBase</i> [31:12]    0x000
TLB Refill	x	x	0	1	x	x	<i>EBase</i> [31:12]    0x180
TLB Refill	0	0	1	0	x	x	0xBFC0.0200
TLB Refill	0	1	1	0	x	x	2'b10    <i>SI_ExceptionBase</i> [29:12]    0x200
TLB Refill	1	x	1	0	x	x	<i>SI_ExceptionBase</i> [31:12]    0x200
TLB Refill	0	0	1	1	x	x	0xBFC0.0380
TLB Refill	0	1	1	1	x	x	2'b10    <i>SI_ExceptionBase</i> [29:12]    0x380
TLB Refill	1	x	1	1	x	x	<i>SI_ExceptionBase</i> [31:12]    0x380
Cache Error	0	x	0	x	x	x	<i>EBase</i> [31:30]    0b1    <i>EBase</i> [28:12]    0x100
Cache Error	1	x	0	x	x	x	0xBFC0.0100 (Config5 <sub>CV</sub> = 0)
Cache Error	1	x	0	x	x	x	<i>EBase</i> [31:12]    0x100 (Config5 <sub>CV</sub> = 1)
Cache Error	0	0	1	x	x	x	0xBFC0.0300
Cache Error	0	1	1	x	x	x	2'b101    <i>SI_ExceptionBase</i> [28:12]    0x300
Cache Error	1	x	1	x	x	x	<i>SI_ExceptionBase</i> [31:12]    0x300
Interrupt	x	x	0	0	0	x	<i>EBase</i> [31:12]    0x180
Interrupt	x	x	0	0	1	x	<i>EBase</i> [31:12]    0x200
Interrupt	0	0	1	0	0	x	0xBFC0.0380
Interrupt	0	1	1	0	0	x	2'b10    <i>SI_ExceptionBase</i> [29:12]    0x380
Interrupt	1	x	1	0	0	x	<i>SI_ExceptionBase</i> [31:12]    0x380
Interrupt	0	0	1	0	1	x	0xBFC0.0400

**Table 6.7 Exception Vectors (continued)**

Exception	Config5k	SI_UseExceptionBase	StatusBEV	StatusEXL	CauseIV	EJTAG Proben	Vector (IntCtl <sub>VS</sub> = 0)
Interrupt	0	1	1	0	1	x	$2'b10 \parallel SI\_ExceptionBase[29:12] \parallel 0x400$
Interrupt	1	x	1	0	1	x	$SI\_ExceptionBase[31:12] \parallel 0x400$
All others	x	x	0	x	x	x	$EBase[31:12] \parallel 0x180$
All others	0	0	1	x	x	x	$0xBFC0.0380$
All others	0	1	1	x	x	x	$2'b10 \parallel SI\_ExceptionBase[29:12] \parallel 0x380$
All others	1	x	1	x	x	x	$SI\_ExceptionBase[31:12] \parallel 0x380$
'x' denotes don't care, '  ' denotes bit string concatenation							

## 6.4 General Exception Processing

With the exception of Reset, NMI, cache error, and EJTAG Debug exceptions, which have their own special processing as described below, exceptions have the same basic processing flow:

- If the *EXL* bit in the *Status* register is zero, the *EPC* register is loaded with the PC at which execution will be restarted, and the *BD* bit is set appropriately in the *Cause* register. The value loaded into the *EPC* register is dependent on whether the processor implements the MIPS16 Module, and whether the instruction is in the delay slot of a branch or jump which has delay slots. Table 6.8 shows the value stored in each of the CP0 PC registers, including *EPC*.

If the *EXL* bit in the *Status* register is set, the *EPC* register is not loaded and the *BD* bit is not changed in the *Cause* register.

**Table 6.8 Value Stored in EPC, ErrorEPC, or DEPC on Exception**

MIPS16 Implemented?	In Branch/Jump Delay Slot?	Value stored in EPC/ErrorEPC/DEPC
No	No	Address of the instruction
No	Yes	Address of the branch or jump instruction (PC-4)
Yes	No	Upper 31 bits of the address of the instruction, combined with the ISA Mode bit
Yes	Yes	Upper 31 bits of the branch or jump instruction (PC-2 in the MIPS16 ISA Mode and PC-4 in the 32-bit ISA Mode), combined with the ISA Mode bit

- The *CE*, and *ExcCode* fields of the *Cause* registers are loaded with the values appropriate to the exception. The *CE* field is loaded, but not defined, for any exception type other than a coprocessor unusable exception.
- The *EXL* bit is set in the *Status* register.
- The processor begins executing at the exception vector.

The value loaded into *EPC* represents the restart address for the exception and need not be modified by exception handler software in the normal case. Software need not look at the *BD* bit in the *Cause* register unless it wishes to identify the address of the instruction that actually caused the exception.

Note that individual exception types may load additional information into other registers. This is noted in the description of each exception type below.

### Operation:

```

/* If Status_EXL is 1, all exceptions go through the general exception vector */
/* and neither the EPC nor Cause_BD are modified */
if Status_EXL = 1 then
    vectorOffset ← 0x180
else
    /* For implementations that include the MIPS16e Module, calculate potential */
    /* PC adjustment for exceptions in the delay slot */
    if Config1_CA = 0 then
        restartPC ← PC
        branchAdjust ← 4          /* Possible adjustment for delay slot */
    else
        restartPC ← PC31..1 || ISAMode

```

```

    if (ISAMode = 0) or ExtendedMIPS16Instruction
        branchAdjust ← 4 /* Possible adjustment for 32-bit MIPS delay slot */
    else
        branchAdjust ← 2 /* Possible adjustment for MIPS16 delay slot */
    endif
endif
if InstructionInBranchDelaySlot then
    EPC ← restartPC - branchAdjust /* PC of branch/jump */
    CauseBD ← 1
else
    EPC ← restartPC /* PC of instruction */
    CauseBD ← 0
endif

/* Compute vector offsets as a function of the type of exception */
if ExceptionType = TLBRefill then
    vectorOffset ← 0x000
elseif (ExceptionType = Interrupt) then
    if (CauseIV = 0) then
        vectorOffset ← 0x180
    else
        if (StatusBEV = 1) or (IntCtlVS = 0) then
            vectorOffset ← 0x200
        else
            if Config3VEIC = 1 then
                VecNum ← CauseR IPL
            else
                VecNum ← VIntPriorityEncoder()
            endif
            vectorOffset ← 0x200 + (VecNum × (IntCtlVS || 0b00000))
        endif /* if (StatusBEV = 1) or (IntCtlVS = 0) then */
    endif /* if (CauseIV = 0) then */
endif /* elseif (ExceptionType = Interrupt) then */
endif /* if StatusEXL = 1 then */

CauseCE ← FaultingCoproprocessorNumber
CauseExcCode ← ExceptionType
StatusEXL ← 1

if Config1CA = 1 then
    ISAMode ← 0
endif

/* Calculate the vector base address */
if StatusBEV = 1 then
    vectorBase ← 0xBFC0.0200
else
    if ArchitectureRevision ≥ 2 then
        /* The fixed value of EBase31..30 forces the base to be in kseg0 or kseg1 */
        vectorBase ← EBase31..12 || 0x000
    else
        vectorBase ← 0x8000.0000
    endif
endif

/* Exception PC is the sum of vectorBase and vectorOffset */
PC ← vectorBase31..30 || (vectorBase29..0 + vectorOffset29..0)
/* No carry between bits 29 and 30 */

```

## 6.5 Debug Exception Processing

All debug exceptions have the same basic processing flow:

- The *DEPC* register is loaded with the program counter (PC) value at which execution will be restarted and the *DBD* bit is set appropriately in the *Debug* register. The value loaded into the *DEPC* register is the current PC if the instruction is not in the delay slot of a branch, or the PC-4 of the branch if the instruction is in the delay slot of a branch.
- The *DSS*, *DBp*, *DDBL*, *DDBS*, *DIB*, and *DINT* bits in the *Debug* register are updated appropriately, depending on the debug exception type.
- *Halt* and *Doze* bits in the *Debug* register are updated appropriately.
- The *DM* bit in the *Debug* register is set to 1.
- The processor is started at the debug exception vector.

The value loaded into *DEPC* represents the restart address for the debug exception and need not be modified by the debug exception handler software in the usual case. Debug software need not look at the *DBD* bit in the *Debug* register unless it wishes to identify the address of the instruction that actually caused the debug exception.

A unique debug exception is indicated through the *DSS*, *DBp*, *DDBL*, *DDBS*, *DIB* and *DINT* bits (D\* bits [5:0]) in the *Debug* register.

No other CP0 registers or fields are changed due to the debug exception, and thus no additional state is saved.

### Operation:

```
if InstructionInBranchDelaySlot then
    DEPC ← PC-4
    DebugDBD ← 1
else
    DEPC ← PC
    DebugDBD ← 0
endif
DebugD* bits at [5:0] ← DebugExceptionType
DebugHalt ← HaltStatusAtDebugException
DebugDoze ← DozeStatusAtDebugException
DebugDM ← 1
if EJTAGControlRegisterProbTrap = 1 then
    PC ← 0xFF20_0200
else
    if DebugControlRegisterRDVec = 1 then
        if CacheErr then
            PC ← 2#101 || DebugVectorAddr28..7 || 2#0000000
        else
            PC ← 2#10 || DebugVectorAddr29..7 || 2#0000000
        else
            if SI_UseExceptionBase
                if CacheErr then
                    PC ← 2#101 || SI_ExceptionBase[28:12] || 0x000
                else
                    PC ← 2#10 || SI_ExceptionBase[29:12] || 0x000
                else
                    PC ← 0xBFC0_0480
            endif
        endif
    endif
```


The location of the debug exception vector is determined by the *ProbTrap* bit in the *EJTAG Control* register (*ECR*) and the *RDVec* bit in the *Debug Control* register (*DCR*), as shown in [Table 6.9](#).

**Table 6.9 Debug Exception Vector Addresses**

ProbTrap bit in ECR Register	RDVec bit in DCR Register	Debug Exception Vector Address
0	0	0xBFC0 0480
0	1	DebugVectorAddr <sub>31:7</sub>    0000000
1	0	0xFF20 0200 in dmseg
1	1	

The value in the optional drseg register *DebugVectorAddr* (offset 0x00020) is used as the debug exception vector when the *ECR ProbTrap* bit is 0 and when enabled through the optional *RDVec* control bit in the *Debug Control Register (DCR)*. Bit 0 of *DebugVectorAddr* determines the ISA mode used to execute the handler. [Figure 6.1](#) shows the format of the *DebugVectorAddr* register; [Table 6.10](#) describes the *DebugVectorAddr* register fields.

**Figure 6.1 DebugVectorAddr Register Format**



**Table 6.10 DebugVectorAddr Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
1	31	Ignored on write; returns one on read.	R	1
DebugVectorOffset	29:7	Programmable Debug Exception Vector Offset	R/W	Preset to 0x7F8009
IM	0	ISA mode to be used for exception handler	R	0
0	30,6:1	Ignored on write; returns zero on read.	R	0

Bits 31..30 of the *DebugVectorAddr* register are fixed with the value 0b10, and the addition of the base address and the exception offset is done inhibiting a carry between bit 29 and bit 30 of the final exception address. The combination of these two restrictions forces the final exception address to be in the kseg0 or kseg1 unmapped virtual address segments. For cache error exceptions, bit 29 is forced to a 1 in the ultimate exception base address, so that this exception always runs in the kseg1 unmapped, uncached virtual address segment.

When MIPS16 is implemented, the power-up state of *IM* is zero. If the implementation does not include MIPS16, the *IM* field is read-only, should be written with zero and will return 0 on a read.

If the TAP is not implemented, the debug exception vector location is as if *ProbTrap*=0.

## 6.6 Exception Descriptions

The following subsections describe each of the exceptions listed in the same sequence as shown in [Table 6.1](#).

### 6.6.1 Reset Exception (Reset)

A reset exception occurs when the *SI\_Reset* signal is asserted to the processor. This exception is not maskable. When a Reset exception occurs, the processor performs a full reset initialization, including aborting state machines, establishing critical state, and generally placing the processor in a state in which it can execute instructions from uncached, unmapped address space. On a Reset exception, the state of the processor is not defined, with the following exceptions:

- The *Random* register is initialized to the number of TLB entries - 1.
- The *Wired* register is initialized to zero.
- The *Config* register is initialized with its boot state.
- The *RP*, *BEV*, *TS*, *SR*, *NMI*, and *ERL* fields of the *Status* register are initialized to a specified state.
- The *I*, *R*, and *W* fields of the *WatchLo* register are initialized to 0.
- The *ErrorEPC* register is loaded with PC-4 if the state of the processor indicates that it was executing an instruction in the delay slot of a branch. Otherwise, the *ErrorEPC* register is loaded with PC. Note that this value may or may not be predictable.
- PC is loaded with 0xBFC0\_0000.

**Cause Register ExcCode Value:**

None

**Additional State Saved:**

None

**Entry Vector Used:**

Reset (exact vector address depends on mode of operation - Legacy/EVA)

**Operation:**

```
Random ← TLBEntries - 1
Wired ← 0
Config ← ConfigurationState
StatusRP ← 0
StatusBEV ← 1
StatusTS ← 0
StatusSR ← 0
StatusNMI ← 0
StatusERL ← 1
WatchLoI ← 0
WatchLoR ← 0
WatchLoW ← 0
if InstructionInBranchDelaySlot then
    ErrorEPC ← PC - 4
else
    ErrorEPC ← PC
endif
```



PC ← 0xBFC0\_0000

## 6.6.2 Debug Single Step Exception (DSS)

A debug single step exception occurs after the CPU has executed one/two instructions in non-debug mode, when returning to non-debug mode after debug mode. One instruction is allowed to execute when returning to a non-jump/branch instruction, otherwise two instructions are allowed to execute since the jump/branch and the instruction in the delay slot are executed as one step. Debug single step exceptions are enabled by the *SSt* bit in the *Debug* register, and are always disabled for the first one/two instructions after a DERET.

The *DEPC* register points to the instruction on which the debug single step exception occurred, which is also the next instruction to single step or execute when returning from debug mode. So the *DEPC* register will not point to the instruction which has just been single stepped, but rather the following instruction. The *DBD* bit in the *Debug* register is never set for a debug single step exception, since the jump/branch and the instruction in the delay slot is executed in one step.

Exceptions occurring on the instruction(s) executed with debug single step exception enabled are taken even though debug single step was enabled. For a normal exception (other than reset), a debug single step exception is then taken on the first instruction in the normal exception handler. Debug exceptions are unaffected by single step mode, e.g. returning to a SDBBP instruction with debug single step exceptions enabled causes a debug software breakpoint exception, and *DEPC* will point to the SDBBP instruction. However, returning to an instruction (not jump/branch) just before the SDBBP instruction, causes a debug single step exception with *DEPC* pointing to the SDBBP instruction.

To ensure proper functionality of single step, the debug single step exception has priority over all other exceptions, except reset and soft reset.

### Debug Register Debug Status Bit Set

*DSS*

### Additional State Saved

None

### Entry Vector Used

Debug exception vector

## 6.6.3 Debug Interrupt Exception (DINT)

A debug interrupt exception is either caused by the *EjtagBrk* bit in the *EJTAG Control* register (controlled through the TAP), or caused by the debug interrupt request signal to the CPU.

The debug interrupt exception is an asynchronous debug exception which is taken as soon as possible, but with no specific relation to the executed instructions. The *DEPC* register is set to the instruction where execution should continue after the debug handler is through. The *DBD* bit is set based on whether the interrupted instruction was executing in the delay slot of a branch.

### Debug Register Debug Status Bit Set

*DINT*

### Additional State Saved

None

**Entry Vector Used**

Debug exception vector

**6.6.4 Non-Maskable Interrupt (NMI) Exception**

A non maskable interrupt exception occurs when the *SI\_NMI* signal is asserted to the processor. *SI\_NMI* is an edge sensitive signal - only one NMI exception will be taken each time it is asserted. An NMI exception occurs only at instruction boundaries, so it does not cause any reset or other hardware initialization. The state of the cache, memory, and other processor states are consistent and all registers are preserved, with the following exceptions:

- The *BEV*, *TS*, *SR*, *NMI*, and *ERL* fields of the *Status* register are initialized to a specified state.
- The *ErrorEPC* register is loaded with PC-4 if the state of the processor indicates that it was executing an instruction in the delay slot of a branch. Otherwise, the *ErrorEPC* register is loaded with PC.
- PC is loaded with 0xBFC0\_0000.

**Cause Register ExcCode Value:**

None

**Additional State Saved:**

None

**Entry Vector Used:**

Reset (exact vector address depends on mode of operation - Legacy/EVA)

**Operation:**

```

StatusBEV ← 1
StatusTS ← 0
StatusSR ← 0
StatusNMI ← 1
StatusERL ← 1
if InstructionInBranchDelaySlot then
    ErrorEPC ← PC - 4
else
    ErrorEPC ← PC
endif
PC ← 0xBFC0_0000

```

**6.6.5 Interrupt Exception (Int)**

The interrupt exception occurs when one or more of the six hardware, two software, or timer interrupt requests is enabled by the *Status* register and the interrupt input is asserted. See 6.8 “Interrupts” on page 338 for more details about the processing of interrupts.

**Register ExcCode Value:**

Int

**Additional State Saved:**

**Table 6.11 Register States an Interrupt Exception**

Register State	Value
<i>CauseIP</i>	Indicates the interrupts that are pending.

**Entry Vector Used:**

See 6.8.2 “Generation of Exception Vector Offsets for Vectored Interrupts” on page 344 for the entry vector used, depending on the interrupt mode the processor is operating in.

## 6.6.6 Debug Instruction Break Exception (DIB)

A debug instruction break exception occurs when an instruction hardware breakpoint matches an executed instruction. The *DEPC* register and *DBD* bit in the *Debug* register indicate the instruction that caused the instruction hardware breakpoint to match. This exception can only occur if instruction hardware breakpoints are implemented.

**Debug Register Debug Status Bit Set:**

*DIB*

**Additional State Saved:**

None

**Entry Vector Used:**

Debug exception vector

## 6.6.7 Watch Exception — Instruction Fetch or Data Access (WATCH)

The Watch facility provides a software debugging vehicle by initiating a watch exception when an instruction or data reference matches the address information stored in the *WatchHi* and *WatchLo* registers. A Watch exception is taken immediately if the *EXL* and *ERL* bits of the *Status* register are both zero and the *DM* bit of the *Debug* register is also zero. If any of those bits is a one at the time that a watch exception would normally be taken, then the *WP* bit in the *Cause* register is set, and the exception is deferred until all three bits are zero. Software may use the *WP* bit in the *Cause* register to determine if the *EPC* register points at the instruction that caused the watch exception, or if the exception actually occurred while in kernel mode.

The Watch exception can occur on either an instruction fetch or a data access. Watch exceptions that occur on an instruction fetch have a higher priority than watch exceptions that occur on a data access.

**Register ExcCode Value:**

WATCH

**Additional State Saved:**

**Table 6.12 Register States on Watch Exception**

Register State	Value
<i>Cause<sub>WP</sub></i>	Indicates that the watch exception was deferred until after <i>Status<sub>EXL</sub></i> , <i>Status<sub>ERL</sub></i> , and <i>Debug<sub>DM</sub></i> were zero. This bit directly causes a watch exception, so software must clear this bit as part of the exception handler to prevent a watch exception loop at the end of the current handler execution.
<i>WatchHi<sub>I,R,W</sub></i>	Set for the watch channel that matched, and indicates which type of match there was.

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.8 Address Error Exception — Instruction Fetch/Data Access (AdEL/AdES)

An address error exception occurs on an instruction or data access when an attempt is made to execute one of the following:

- Fetch an instruction, load a word, or store a word that is not aligned on a word boundary
- Load or store a halfword that is not aligned on a halfword boundary
- Reference the kernel address space from user mode
- Reference to a non-user address space when using the new EVA instructions

Note that in the case of an instruction fetch that is not aligned on a word boundary, PC is updated before the condition is detected. Therefore, both *EPC* and *BadVAddr* point to the unaligned instruction address. In the case of a data access the exception is taken if either an unaligned address or an address that was inaccessible in the current processor mode was referenced by a load or store instruction.

**Cause Register ExcCode Value:**

ADEL: Reference was a load or an instruction fetch

ADES: Reference was a store

**Additional State Saved:**

**Table 6.13 CP0 Register States on Address Exception Error**

Register State	Value
<i>BadVAddr</i>	Failing address
<i>Context<sub>VPN2</sub></i>	UNPREDICTABLE
<i>EntryHi<sub>VPN2</sub></i>	UNPREDICTABLE
<i>EntryLo0</i>	UNPREDICTABLE
<i>EntryLo1</i>	UNPREDICTABLE

**Entry Vector Used:**

General exception vector (offset 0x180)

## 6.6.9 TLB Refill Exception — Instruction Fetch or Data Access (TLBL/TLBS)

During an instruction fetch or data access, a TLB refill exception occurs when no TLB entry matches a reference to a mapped address space and the *EXL* bit is 0 in the *Status* register. Note that this is distinct from the case in which an entry matches but has the valid bit off. In that case, a TLB Invalid exception occurs.

### Cause Register ExcCode Value:

TLBL: Reference was a load or an instruction fetch

TLBS: Reference was a store

### Additional State Saved:

**Table 6.14 CP0 Register States on TLB Refill Exception**

Register State	Value
<i>BadVAddr</i>	Failing address.
<i>Context</i>	The <i>BadVPN2</i> field contains VA <sub>31:13</sub> of the failing address.
<i>EntryHi</i>	The <i>VPN2</i> field contains VA <sub>31:13</sub> of the failing address; the <i>ASID</i> field contains the ASID of the reference that missed.
<i>EntryLo0</i>	UNPREDICTABLE
<i>EntryLo1</i>	UNPREDICTABLE

### Entry Vector Used:

TLB refill vector (offset 0x000) if *Status*<sub>EXL</sub> = 0 at the time of exception;

General exception vector (offset 0x180) if *Status*<sub>EXL</sub> = 1 at the time of exception

## 6.6.10 TLB Invalid Exception — Instruction Fetch or Data Access (TLBINV)

During an instruction fetch or data access, a TLB invalid exception occurs in one of the following cases:

- No TLB entry matches a reference to a mapped address space; and the *EXL* bit is 1 in the *Status* register.
- A TLB entry matches a reference to a mapped address space, but the matched entry has the valid bit off.

### Cause Register ExcCode Value:

TLBL: Reference was a load or an instruction fetch

TLBS: Reference was a store

**Additional State Saved:**

**Table 6.15 CP0 Register States on TLB Invalid Exception**

Register State	Value
<i>BadVAddr</i>	Failing address
<i>Context</i>	The BadVPN2 field contains VA <sub>31:13</sub> of the failing address.
<i>EntryHi</i>	The VPN2 field contains VA <sub>31:13</sub> of the failing address; the ASID field contains the ASID of the reference that missed.
<i>EntryLo0</i>	UNPREDICTABLE
<i>EntryLo1</i>	UNPREDICTABLE

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.11 TLB Execute-Inhibit Exception (TLBXI)

A *TLB execute-inhibit* exception occurs when there is a execute access from a TLB entry whose XI bit is set. The *TLB execute-inhibit* exception type can only occur if execute-inhibit exceptions are enabled by setting bit 30 (XIE) in the *PageGrain* register.

In addition, the type of exception taken depends on the state of the *PageGrain*<sub>IEC</sub> bit. If the XI bit of the entry is set, and the *PageGrain*<sub>IEC</sub> bit is set, a TLBXI exception is taken. If the *PageGrain*<sub>IEC</sub> bit is cleared, a *TLBL* exception is taken.

**Cause Register ExcCode Value:**

if *PageGrain*<sub>IEC</sub> == 0 TLBL

if *PageGrain*<sub>IEC</sub> == 1 TLBXI

### Additional State Saved:

**Table 6.16 CP0 Register States on TLB Execute-Inhibit Exception**

Register State	Value
<i>BadVAddr</i>	Failing address.
<i>Context</i>	If the <i>Config3.CTXTC</i> bit is set, then the bits of the <i>Context</i> register corresponding to the set bits of the <i>VirtualIndex</i> field of the <i>ContextConfig</i> register are loaded with the high-order bits of the virtual address that missed.  If the <i>Config3.CTXTC</i> bit is clear, then the <i>BadVPN2</i> field contains $VA_{31:13}$ of the failing address.
<i>EntryHi</i>	The <i>VPN2</i> field contains $VA_{31:13}$ of the failing address; the <i>ASID</i> field contains the ASID of the reference that missed.
<i>EntryLo0</i>	UNPREDICTABLE
<i>EntryLo1</i>	UNPREDICTABLE

### Entry Vector Used:

General exception vector (offset 0x180)

## 6.6.12 TLB Read-Inhibit Exception (TLBRI)

A *TLB read-inhibit* exception occurs when there is an attempt to read a TLB entry whose RI bit is set. The *TLB read-inhibit* exception type can only occur if read-inhibit exceptions are enabled by setting bit 31 (RIE) in the *PageGrain* register.

In addition, the type of exception taken depends on the state of the *PageGrain<sub>IEC</sub>* bit. If the RI bit of the entry is set, and the *PageGrain<sub>IEC</sub>* bit is set, a TLBRI exception is taken. If the *PageGrain<sub>IEC</sub>* bit is cleared, a *TLBL* exception is taken.

### Cause Register ExcCode Value:

if *PageGrain<sub>IEC</sub>* == 0 TLBL

if *PageGrain<sub>IEC</sub>* == 1 TLBRI

## Additional State Saved:

**Table 6.17 CP0 Register States on TLB Read-Inhibit Exception**

Register State	Value
<i>BadVAddr</i>	Failing address.
<i>Context</i>	If the <i>Config3.CTXTC</i> bit is set, then the bits of the <i>Context</i> register corresponding to the set bits of the <i>VirtualIndex</i> field of the <i>ContextConfig</i> register are loaded with the high-order bits of the virtual address that missed.  If the <i>Config3.CTXTC</i> bit is clear, then the <i>BadVPN2</i> field contains VA <sub>31:13</sub> of the failing address.
<i>EntryHi</i>	The <i>VPN2</i> field contains VA <sub>31:13</sub> of the failing address; the <i>ASID</i> field contains the ASID of the reference that missed.
<i>EntryLo0</i>	UNPREDICTABLE
<i>EntryLo1</i>	UNPREDICTABLE

## Entry Vector Used:

General exception vector (offset 0x180)

### 6.6.13 Cache Error Exception (ICache Error/DCache Error)

A cache error exception occurs when an instruction or data reference detects a cache tag or data error. This exception is not maskable. Because the error was in a cache, the exception vector is to an unmapped, uncached address. This exception can be imprecise and the *ErrorEPC* may not point to the instruction that saw the error. Additionally, because the caches on the cores within the interAptiv core are coherent, cache errors detected on other cores could indicate data corruption for a process on this CPU. An error on another CPU will still cause a Cache Error exception, with the *CacheErr<sub>EE</sub>* indicating that the error occurred on another processor.

Instruction cache parity errors are precise and are taken only if the core is going to execute an instruction that saw the parity error on a read of the instruction cache. If multiple instruction cache errors are detected prior to the exception being taken, the *CacheErr* register contents capture the information for the most recent error, which may not correlate to the instruction indicated by *ErrorEPC*. Error handling code should use the *CacheErr* contents to process the exception. Upon returning to the instruction indicated by *ErrorEPC*, that error would be seen again.

For data cache parity or uncorrectable ECC errors, this exception can be imprecise and the *ErrorEPC* may not point to the instruction that saw the error.

L2 cache errors are considered to be imprecise. An L2 cache error on a data load operation can potentially corrupt the target GPR.

## Cause Register ExcCode Value

N/A



## Additional State Saved

**Table 6.18 CP0 Register States on Cache Error Exception**

Register State	Value
<i>CacheErr</i>	Error state
<i>ErrorEPC</i>	Restart PC

### Entry Vector Used

Cache error vector (offset 0x100)

## 6.6.14 Bus Error Exception — Instruction Fetch or Data Access (IBE)

A bus error exception occurs when an instruction or data access makes a bus request (due to a cache miss or an uncacheable reference) and that request terminates in an error. The bus error exception can occur on either an instruction fetch or a data read. Bus error exceptions cannot be generated on data writes. Bus error exceptions that occur on an instruction fetch have a higher priority than bus error exceptions that occur on a data access.

Instruction errors are precise, while data bus errors can be imprecise. These errors are taken when the ERR code is returned on the *OC\_SResp* input.

### Cause Register ExcCode Value:

IBE: Error on an instruction reference

DBE: Error on a data reference

### Additional State Saved:

None

### Entry Vector Used:

General exception vector (offset 0x180)

## 6.6.15 Debug Software Breakpoint Exception (DBp)

A debug software breakpoint exception occurs when an SDBBP instruction is executed. The *DEPC* register and DBD bit in the *Debug* register will indicate the SDBBP instruction that caused the debug exception.

### Debug Register Debug Status Bit Set:

*DBp*

### Additional State Saved:

None

### Entry Vector Used:

Debug exception vector

## 6.6.16 Execution Exception — System Call (Sys)

The system call exception is one of the execution exceptions. All of these exceptions have the same priority. A system call exception occurs when a SYSCALL instruction is executed.

**Cause Register ExcCode Value:**

Sys

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

**6.6.17 Execution Exception — Breakpoint (Bp)**

The breakpoint exception is one of the execution exceptions. All of these exceptions have the same priority. A breakpoint exception occurs when a BREAK instruction is executed.

**Cause Register ExcCode Value:**

Bp

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

**6.6.18 Execution Exception — Coprocessor Unusable (CpU)**

The coprocessor unusable exception is one of the execution exceptions. All of these exceptions have the same priority. A coprocessor unusable exception occurs when an attempt is made to execute a coprocessor instruction for one of the following:

- a corresponding coprocessor unit that has not been marked usable by setting its CU bit in the *Status* register
- CP0 instructions, when the unit has not been marked usable, and the processor is executing in user mode

**Cause Register ExcCode Value:**

CpU

**Additional State Saved:****Table 6.19 Register States on Coprocessor Unusable Exception**

Register State	Value
<i>Cause<sub>CE</sub></i>	Unit number of the coprocessor being referenced

**Entry Vector Used:**

General exception vector (offset 0x180)

**6.6.19 Execution Exception — CorExtend Block Unusable (CEU)**

The CorExtend block unusable exception is one of the execution exceptions. All of these exceptions have the same priority. A CEU exception occurs when an attempt is made to execute a CorExtend instruction when the CEE bit in

the *Status* register is not set. It is dependent on the implementation of the CorExtend block, but this exception should be taken on any CorExtend instruction that modifies local state within the CorExtend block and can optionally be taken on other CorExtend instructions.

**Cause Register ExcCode Value:**

CEU

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.20 Execution Exception — DSP ASE State Disabled (DSPDis)

The DSP ASE State Disabled exception is an execution exception. It occurs when an attempt is made to execute a DSP ASE instruction when the MX bit in the Status register is not set. This allows an OS to do “lazy” context switching.

**Cause Register ExcCode Value:**

DSPDis

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.21 Execution Exception — Reserved Instruction (RI)

The reserved instruction exception is one of the execution exceptions. All of these exceptions have the same priority. A reserved instruction exception occurs when a reserved or undefined major opcode or function field is executed. This includes Coprocessor 2 instructions which are decoded reserved in the Coprocessor 2.

**Cause Register ExcCode Value:**

RI

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.22 Execution Exception — Floating Point Exception (FPE)

A floating point exception is initiated by the floating point coprocessor.

**Cause Register ExcCode Value:**

FPE

**Additional State Saved:**

**Table 6.20 Register States on Floating Point Exception**

Register State	Value
FCSR	Indicates the cause of the floating point exception

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.23 Execution Exception — Coprocessor2 (C2E)

A C2E exception is signalled from the optional coprocessor2 block on a coprocessor instruction.

**Cause Register ExcCode Value:**

C2E

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.24 Execution Exception — Coprocessor2 (IS1)

An IS1 exception is signalled from the optional coprocessor2 block on a coprocessor instruction.

**Cause Register ExcCode Value:**

IS1

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.25 Execution Exception — Integer Overflow (Ov)

The integer overflow exception is one of the execution exceptions. All of these exceptions have the same priority. An integer overflow exception occurs when selected integer instructions result in a 2's complement overflow.

**Cause Register ExcCode Value:**

Ov

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

**6.6.26 Execution Exception — Trap (Tr)**

The trap exception is one of the execution exceptions. All of these exceptions have the same priority. A trap exception occurs when a trap instruction results in a TRUE value.

**Cause Register ExcCode Value:**

Tr

**Additional State Saved:**

None

**Entry Vector Used:**

General exception vector (offset 0x180)

**6.6.27 Execution Exceptions — MT\_ov, MT\_under, MT\_invalid, MT\_yield\_sched**

- **MT\_ov** - A Thread Overflow condition on a FORK, where a TC allocation request cannot be satisfied.
- **MT\_under** - A Thread Underflow condition on a YIELD, where the termination and deallocation of a thread leaves no dynamically allocatable TCs activated on a VPE.
- **MT\_invalid** - An Invalid qualifier condition, where a YIELD instruction specifies an invalid condition for resuming execution.
- **MT\_yield\_sched** - A YIELD scheduler exception, where a valid YIELD instruction could have caused a rescheduling of a TC, and the YIELD Intercept bit is set. This happens when a YIELD is executed with a yield qualifier of -1, 0, or any positive value when *VPEControl*[YSI]=1 and *TCStatusDT*=1. Lower priority than *MT\_under* or *MT\_Invalid* - if one of those conditions is met by YIELD, that exception will be taken instead.

**Cause Register ExcCode Value:**

Thread.

**Additional State Saved:**

There is a sub-cause filed in *VPEControl*[EXCPT], which indicates the type of Thread exception. [Table 6.21](#) shows the different Thread Exception Codes.

**Table 6.21 Thread Exception Codes in VPEControl[EXCPT]**

Value	Exception Name	Exception Type
0	MT_ov	Execution
1	MT_under	Execution
2	MT_invalid	Execution
3	MT_gs	Thread (Section 5.7.30)
4	MT_yield_sched	Execution

**Table 6.21 Thread Exception Codes in VPEControl[EXCPT]**

Value	Exception Name	Exception Type
5	MT_gss	Thread (Section 5.7.30)

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.28 Debug Data Break Exception (DDBL/DDBS)

A debug data break exception occurs when a data hardware breakpoint matches the load/store transaction of an executed load/store instruction. The *DEPC* register and *DBD* bit in the *Debug* register will indicate the load/store instruction that caused the data hardware breakpoint to match. The load/store instruction that caused the debug exception has not completed e.g. not updated the register file, and the instruction can be re-executed after returning from the debug handler.

**Debug Register Debug Status Bit Set:**

*DDBL* for a load instruction or *DDBS* for a store instruction

**Additional State Saved:**

None

**Entry Vector Used:**

Debug exception vector

### 6.6.29 TLB Modified Exception (TLB Mod)

During a data access, a TLB modified exception occurs on a store reference to a mapped address if the following condition is true:

- The matching TLB entry is valid, but not dirty.

**Cause Register ExcCode Value:**

Mod

**Additional State Saved:**

**Table 6.22 Register States on TLB Modified Exception**

Register State	Value
<i>BadVAddr</i>	Failing address
<i>Context</i>	The BadVPN2 field contains VA <sub>31:13</sub> of the failing address.
<i>EntryHi</i>	The VPN2 field contains VA <sub>31:13</sub> of the failing address; the ASID field contains the ASID of the reference that missed.
<i>EntryLo0</i>	UNPREDICTABLE
<i>EntryLo1</i>	UNPREDICTABLE

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.30 Thread Exceptions — MT\_gs, MT\_gss

- **MT\_gs** - A Gating Storage exception condition, where implementation dependent logic associated with gating or inter-thread communication (ITC) storage requires software intervention.
- **MT\_gss** - A Gating Storage Scheduler exception, where a Gating Storage load or store would have blocked and caused a rescheduling of a TC, and the GS Intercept bit is set.

**Cause Register ExcCode Value:**

Thread.

**Additional State Saved:**

There is a sub-cause filed in *VPEControl*[EXCPT], which indicates the type of Thread exception. [Table 6.21](#) shows the different Thread Exception Codes.

**Entry Vector Used:**

General exception vector (offset 0x180)

### 6.6.31 Memory Protection Unit (MPU) Exception — MPUL/MPUS

The protection exception occurs when an access to memory that has been protected by the Memory Protection Unit has been attempted.

**Register ExcCode Value**

Prot (Cause Code 29)

**Additional State Saved**

MPU Config Register - Triggered fields

BadVAddr - Causing address

**Entry Vector Used**

General exception vector (offset 0x180)

## 6.7 Exception Handling and Servicing Flowcharts

The remainder of this chapter contains flowcharts for the following exceptions and guidelines for their handlers:

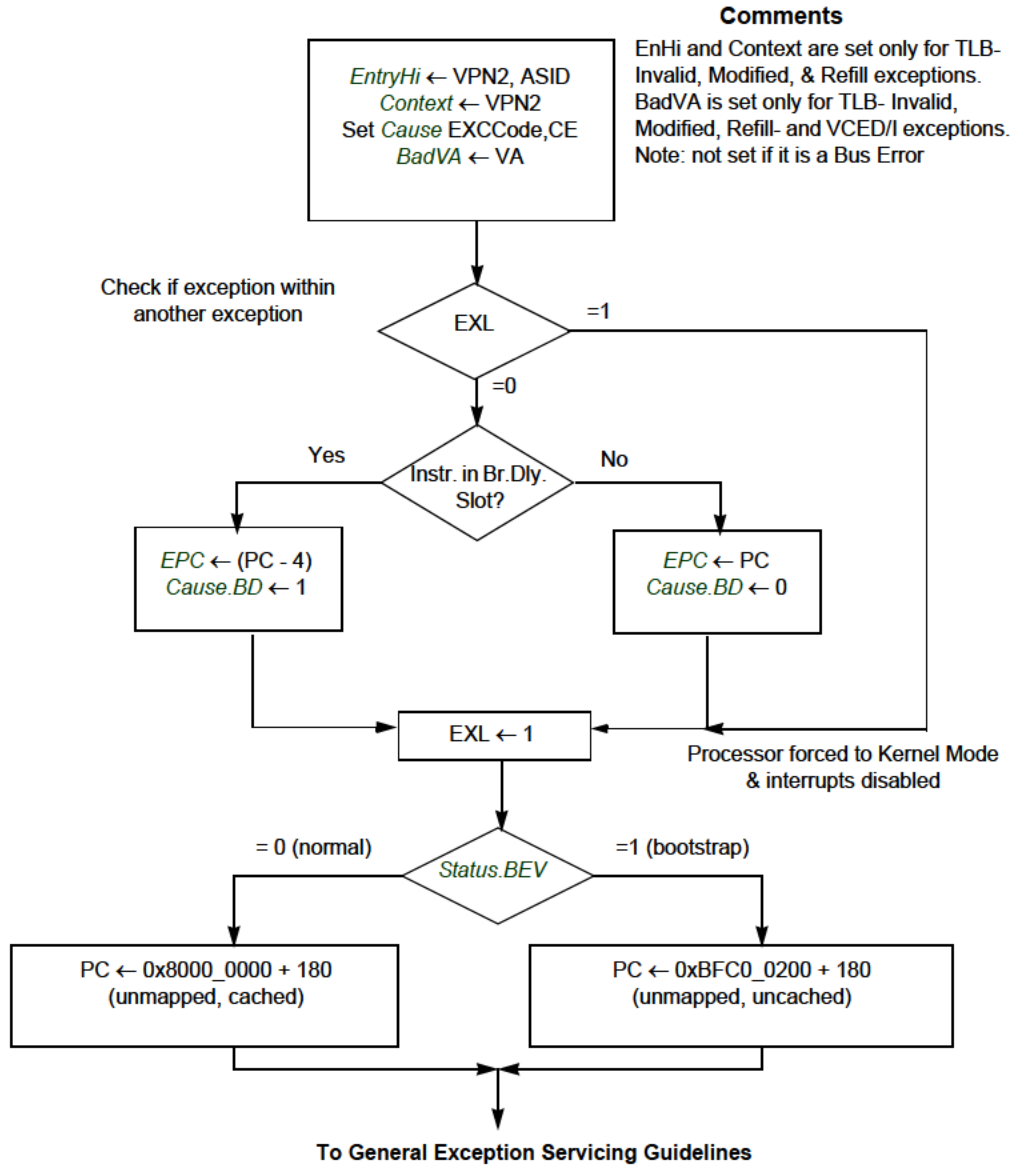
- General exceptions
- TLB miss exceptions
- Reset and NMI exceptions
- Debug exceptions

Generally speaking, exceptions are handled by hardware and then serviced by software. Note that unexpected debug exceptions to the debug exception vector at 0xBFC0\_0200 may be viewed as a reserved instruction since uncontrolled execution of an SDBBP instruction caused the exception. The DERET instruction must be used at return from the debug exception handler, in order to leave debug mode and return to non-debug mode. The DERET instruction returns to the address in the *DEPC* register.



### Figure 6.2 General Exception Handler (HW)

Exceptions other than Reset, NMI, or first-level TLB miss. Note: Interrupts can be masked by IE or IMs, and Watch is masked if EXL = 1.



**Figure 6.3 General Exception Servicing Guidelines (SW)**

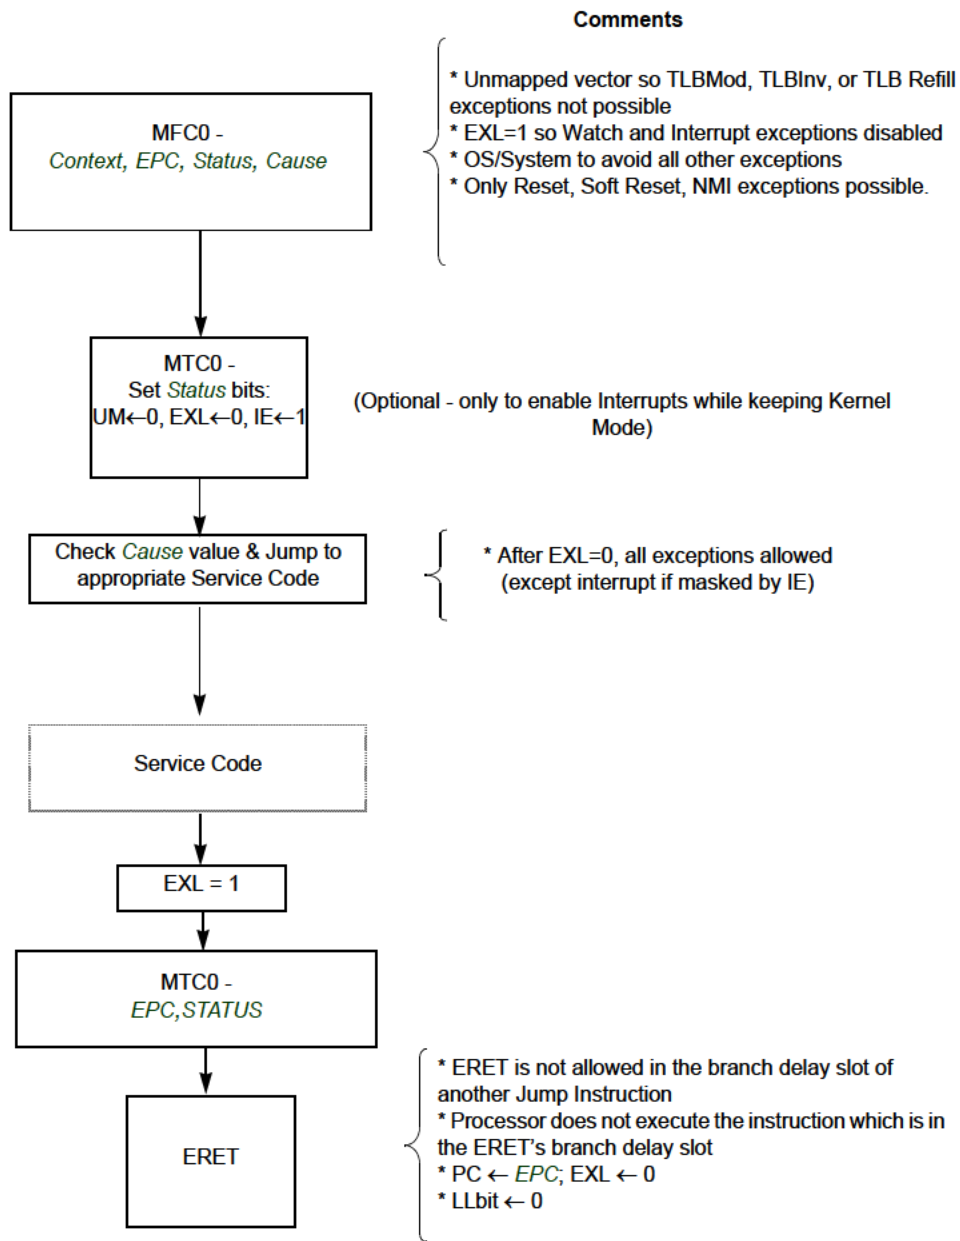


Figure 6.4 TLB Miss Exception Handler (HW)

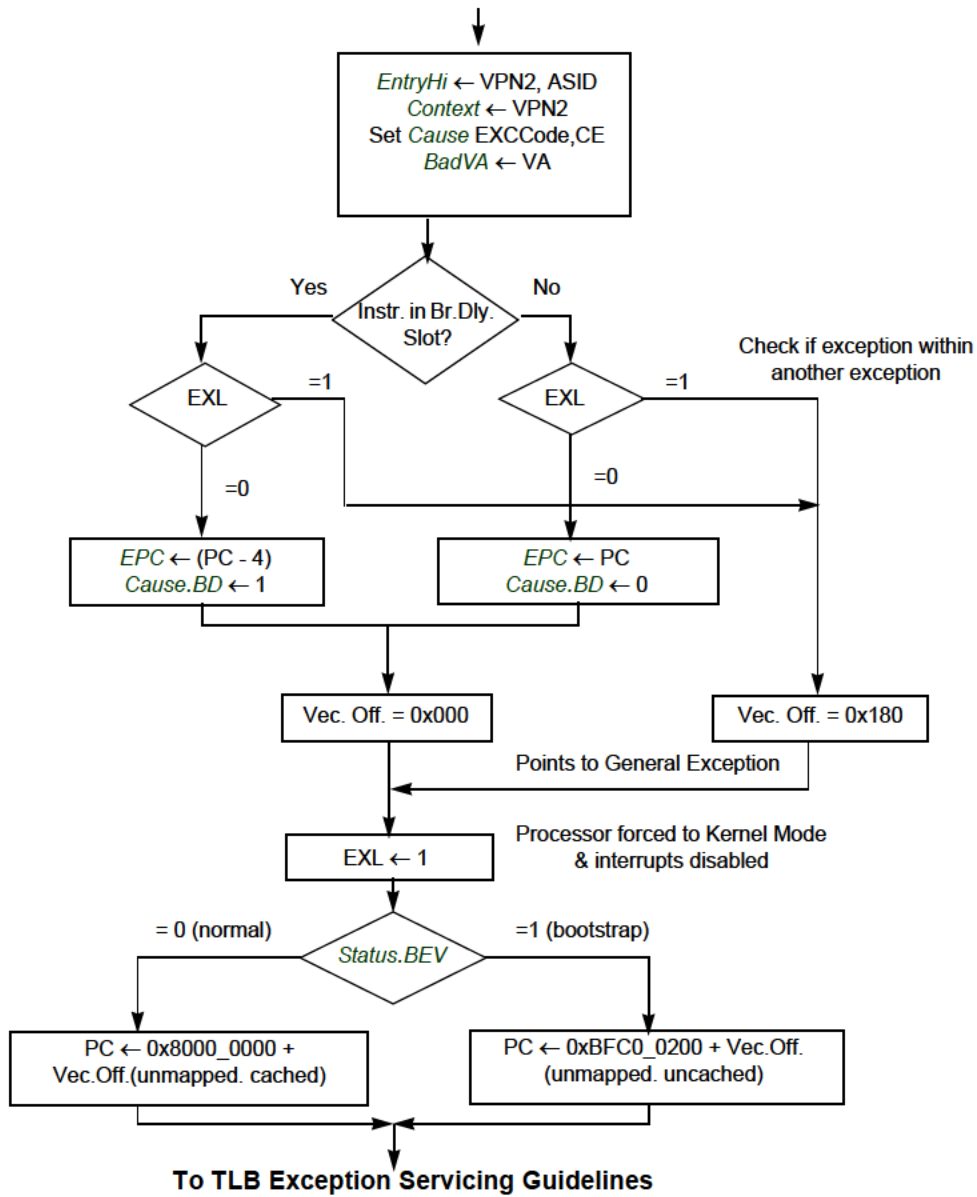


Figure 6.5 TLB Exception Servicing Guidelines (SW)

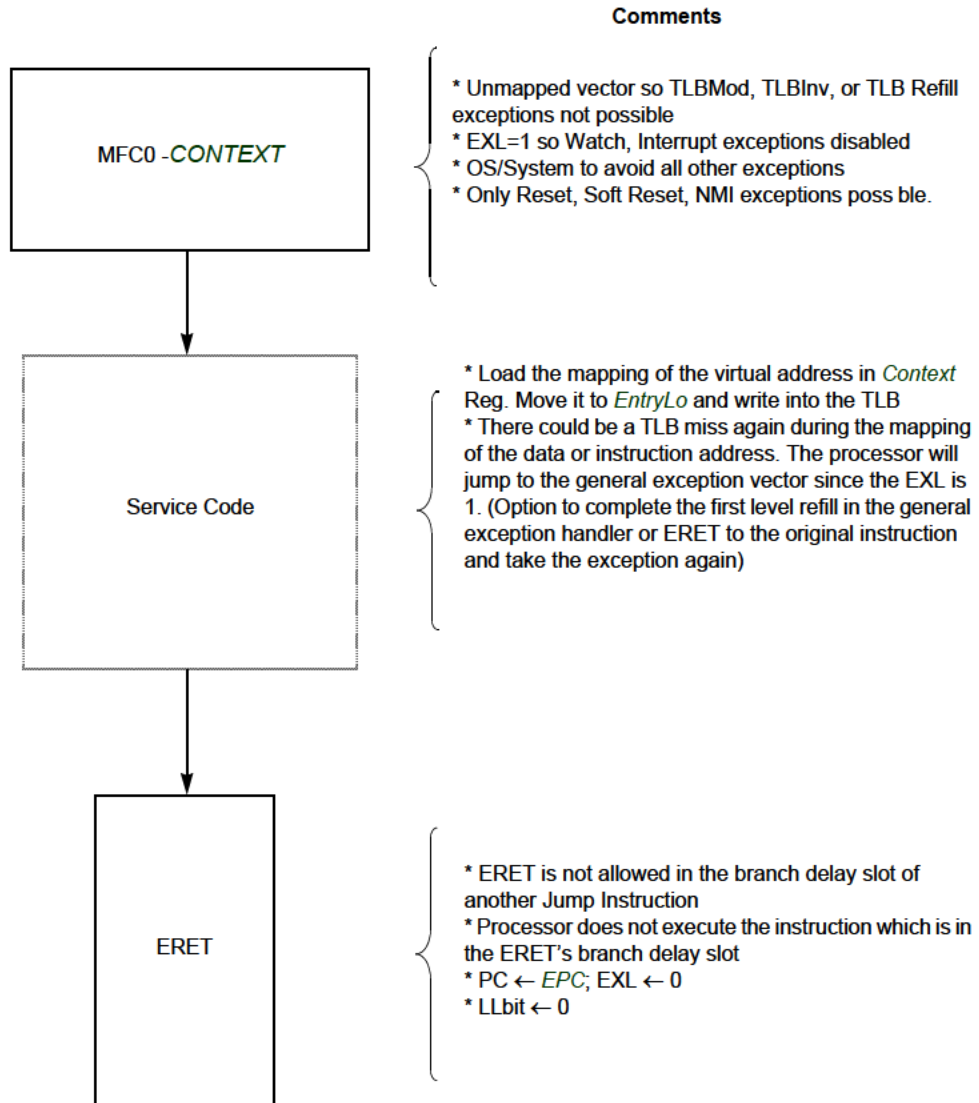
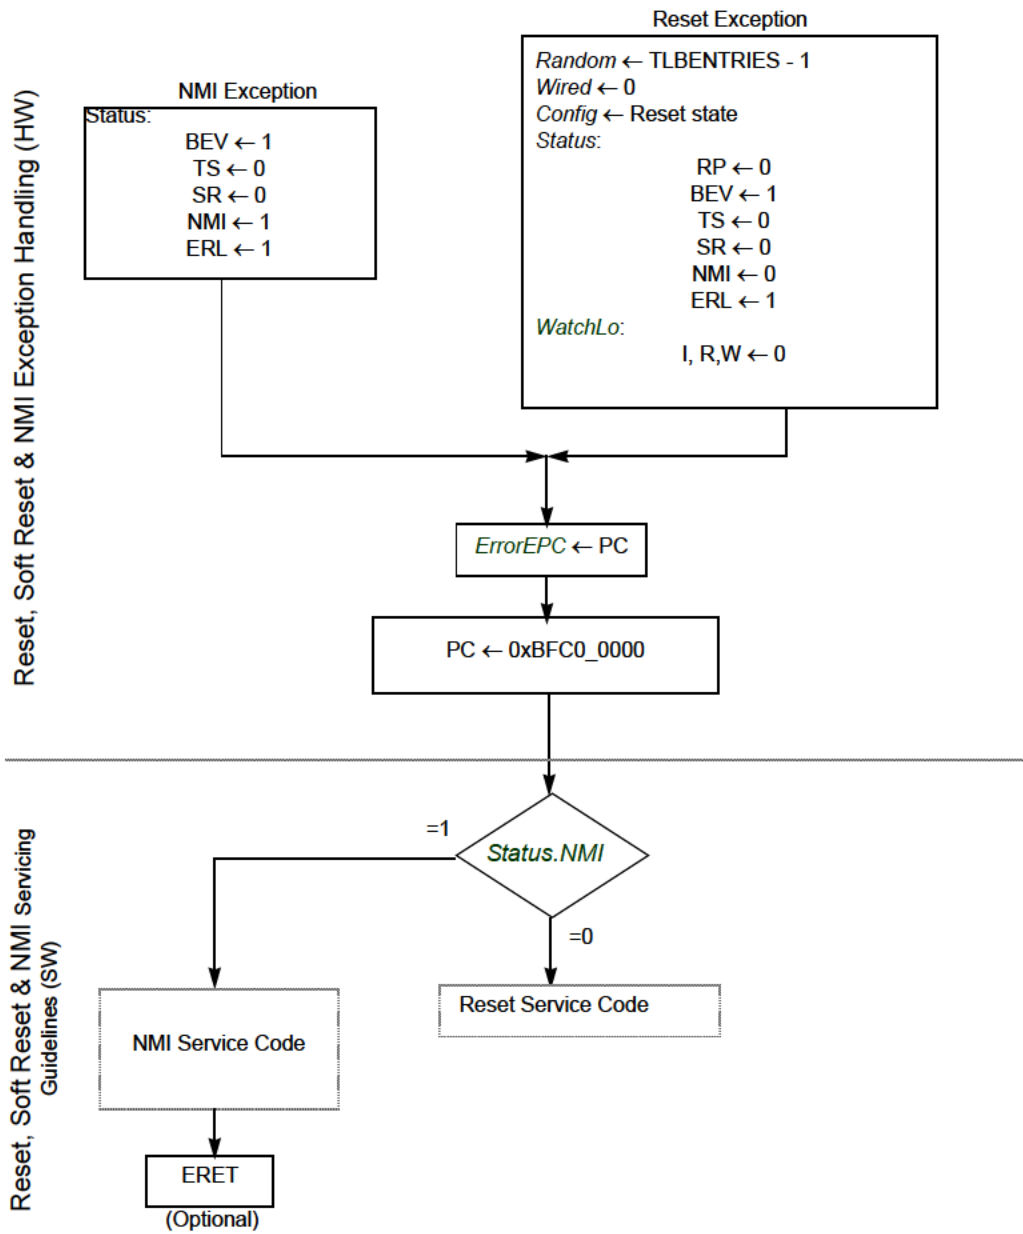


Figure 6.6 Reset and NMI Exception Handling and Servicing Guidelines



## 6.8 Interrupts

Release 3 of the MIPS32 architecture, implemented by the interAptiv core, includes support for vectored interrupts and the implementation of a new interrupt mode that permits the use of an external interrupt controller.

Additionally, internal performance counters have been added to the interAptiv core. These counters can be configured to count various events within the CPU. When the MSB of the counter is set, it can trigger a performance counter interrupt. This interrupt, like the timer interrupt, is an output from the core that can be brought back into the cores interrupt pins in a system-dependent manner.

The Fast Debug Channel feature in EJTAG provides a low overhead means for sending data between CPU software and the EJTAG probe. It includes a pair of FIFOs for transmit and receive data. Software can define FIFO thresholds for generating an interrupt. The fast debug channel interrupt is also routed similarly to the timer and performance counter interrupts. The interrupt status is made available on an output pin and can be brought back into the cores interrupt pins.

### 6.8.1 Interrupt Modes

The interAptiv core includes support for three interrupt modes, as defined by Release 3 of the Architecture:

- Interrupt Compatibility mode, in which the behavior of the interAptiv core is identical to the behavior of an implementation of Release 1 of the Architecture.
- Vectored Interrupt (VI) mode adds the ability to prioritize and vector interrupts to a handler dedicated to that interrupt. The presence of this mode is denoted by the *VInt* bit in the *Config3* register. Although this mode is architecturally optional, it is always present on the interAptiv core, so the *VInt* bit will always read as a 1.
- External Interrupt Controller (EIC) mode, which redefines the way interrupts are handled to provide full support for an external interrupt controller that handles prioritization and vectoring of interrupts. As with VI mode, this mode is architecturally optional. The presence of this mode is denoted by the *VEIC* bit in the *Config3* register. On the interAptiv core, the *VEIC* bit is set externally by the static input, *SI\_EICPresent*, to allow system logic to indicate the presence of an external interrupt controller.

Following reset, the interAptiv core defaults to Compatibility mode, which is fully compatible with all implementations of Release 1 of the Architecture.

Table 6.23 shows the current interrupt mode of the processor as a function of the Coprocessor 0 register fields that can affect the mode.

**Table 6.23 Interrupt Modes**

<i>StatusBEV</i>	<i>CauseIV</i>	<i>IntCtlVS</i>	<i>Config3VINT</i>	<i>Config3VEIC</i>	<b>Interrupt Mode</b>
1	x	x	x	x	Compatibility
x	0	x	x	x	Compatibility
x	x	=0	x	x	Compatibility
0	1	≠0	1	0	Vectored Interrupt
0	1	≠0	x	1	External Interrupt Controller
0	1	≠0	0	0	Cannot occur because <i>IntCtl<sub>VS</sub></i> cannot be non-zero if neither Vectored Interrupt nor External Interrupt Controller mode is implemented.
“x” denotes don’t care					

### 6.8.1.1 Interrupt Compatibility Mode

This is the default interrupt mode for the processor and is entered when a Reset exception occurs. In this mode, interrupts are non-vectored and dispatched through exception vector offset 0x180 (if  $Cause_{IV} = 0$ ) or vector offset 0x200 (if  $Cause_{IV} = 1$ ). This mode is in effect when any of the following conditions are true:

- $Cause_{IV} = 0$
- $Status_{BEV} = 1$
- $IntCtl_{VS} = 0$ , which is the case if vectored interrupts are not implemented or have been disabled.

Here is a typical software handler for compatibility mode:

```
/*
 * Assumptions:
 * - CauseIV = 1 (if it were zero, the interrupt exception would have to
 *   be isolated from the general exception vector before arriving
 *   here)
 * - GPRs k0 and k1 are available
 * - The software priority is IP7..IP0 (HW5..HW0, SW1..SW0)
 *
 * Location: Offset 0x200 from exception base
 */

IVexception:
    mfc0    k0, CO_Cause      /* Read Cause register for IP bits */
    mfc0    k1, CO_Status    /* and Status register for IM bits */
    andi   k0, k0, M_CauseIM /* Keep only IP bits from Cause */
    and    k0, k0, k1        /* and mask with IM bits */
    beq    k0, zero, Dismiss /* no bits set - spurious interrupt */
    clz    k0, k0            /* Find first bit set, IP7..IP0; k0 = 16..23 */
    xori   k0, k0, 0x17      /* 16..23 => 7..0 */
    sll    k0, k0, VS        /* Shift to emulate software IntCtlVS */
    la     k1, VectorBase    /* Get base of 8 interrupt vectors */
    addu   k0, k0, k1        /* Compute target from base and offset */
    jr     k0                /* Jump to specific exception routine */
    nop

/*
 * Each interrupt processing routine processes a specific interrupt, analogous
 * to those reached in VI or EIC interrupt mode. Since each processing routine
 * is dedicated to a particular interrupt line, it has the context to know
 * which line was asserted. Each processing routine may need to look further
 * to determine the actual source of the interrupt if multiple interrupt requests
 * are ORed together on a single IP line. Once that task is performed, the
 * interrupt may be processed in one of two ways:
 *
 * - Completely at interrupt level (e.g., a simple UART interrupt). The
 *   SimpleInterrupt routine below is an example of this type.
 * - By saving sufficient state and re-enabling other interrupts. In this
 *   case the software model determines which interrupts are disabled during
 *   the processing of this interrupt. Typically, this is either the single
 *   StatusIM bit that corresponds to the interrupt being processed, or some
 *   collection of other StatusIM bits so that "lower" priority interrupts are
 *   also disabled. The NestedInterrupt routine below is an example of this type.
 */
```

```

SimpleInterrupt:
/*
 * Process the device interrupt here and clear the interrupt request
 * at the device. In order to do this, some registers may need to be
 * saved and restored. The coprocessor 0 state is such that an ERET
 * will simply return to the interrupted code.
 */
    eret                                /* Return to interrupted code */

NestedException:
/*
 * Nested exceptions typically require saving the EPC and Status registers,
 * saving any GPRs that may be modified by the nested exception routine, disabling
 * the appropriate IM bits in Status to prevent an interrupt loop, putting
 * the processor in kernel mode, and re-enabling interrupts. The sample code
 * below cannot cover all nuances of this processing and is intended only
 * to demonstrate the concepts.
 */

    /* Save GPRs here, and setup software context */
    mfc0    k0, CO_EPC                /* Get restart address */
    sw      k0, EPCSave                /* Save in memory */
    mfc0    k0, CO_Status              /* Get Status value */
    sw      k0, StatusSave             /* Save in memory */
    li      k1, ~IMbitsToClear        /* Get IM bits to clear for this interrupt */
                                           /* this must include at least the IM bit */
                                           /* for the current interrupt, and may include */
                                           /* others */
    and     k0, k0, k1                /* Clear bits in copy of Status */
    ins     k0, zero, S_StatusEXL, (W_StatusKSU+W_StatusERL+W_StatusEXL)
                                           /* Clear KSU, ERL, EXL bits in k0 */
    mtc0    k0, CO_Status              /* Modify mask, switch to kernel mode, */
                                           /* re-enable interrupts */

/*
 * Process interrupt here, including clearing device interrupt.
 * In some environments this may be done with the core running in
 * kernel or user mode. Such an environment is well beyond the scope of
 * this example.
 */

/*
 * To complete interrupt processing, the saved values must be restored
 * and the original interrupted code restarted.
 */

    di                                /* Disable interrupts - may not be required */
    lw      k0, StatusSave              /* Get saved Status (including EXL set) */
    lw      k1, EPCSave                /* and EPC */
    mtc0    k0, CO_Status              /* Restore the original value */
    mtc0    k1, CO_EPC                /* and EPC */
    /* Restore GPRs and software state */
    eret                                /* Dismiss the interrupt */

```



### 6.8.1.2 Vectored Interrupt Mode

In Vectored Interrupt (VI) mode, a priority encoder prioritizes pending interrupts and generates a vector which can be used to direct each interrupt to a dedicated handler routine. VI mode is in effect when all the following conditions are true:

- $Config3_{VInt} = 1$
- $Config3_{VEIC} = 0$
- $IntCtl_{VS} \neq 0$
- $Cause_{IV} = 1$
- $Status_{BEV} = 0$

In VI interrupt mode, the six hardware interrupts are interpreted as individual hardware interrupt requests. The timer, performance counter, and fast debug channel interrupts are combined in a system-dependent way (external to the CPU) with the hardware interrupts (the interrupt with which they are combined is indicated by the  $IntCtl_{IPTI|IPC|IPFDCI}$  fields) to provide the appropriate relative priority of the those interrupts with that of the hardware interrupts. The processor interrupt logic ANDs each of the  $Cause_{IP}$  bits with the corresponding  $Status_{IM}$  bits. If any of these values is 1, and if interrupts are enabled ( $Status_{IE} = 1$ ,  $Status_{EXL} = 0$ , and  $Status_{ERL} = 0$ ), an interrupt is signaled and a priority encoder scans the values in the order shown in [Table 6.24](#).

**Table 6.24 Relative Interrupt Priority for Vectored Interrupt Mode**

Relative Priority	Interrupt Type	Interrupt Source	Interrupt Request Calculated From	Vector Number Generated by Priority Encoder
Highest Priority	Hardware	HW5	IP7 and IM7	7
		HW4	IP6 and IM6	6
		HW3	IP5 and IM5	5
		HW2	IP4 and IM4	4
		HW1	IP3 and IM3	3
		HW0	IP2 and IM2	2
Lowest Priority	Software	SW1	IP1 and IM1	1
		SW0	IP0 and IM0	0

A typical software handler for Vectored Interrupt mode bypasses the entire sequence of code following the `IVexception` label shown for the compatibility mode handler above. Instead, the hardware performs the prioritization, dispatching directly to the interrupt processing routine.

A nested interrupt is similar to that shown for compatibility mode. Such a routine might look as follows:

```
NestedException:
/*
 * Nested exceptions typically require saving the EPC and Status registers,
 * disabling the appropriate IM bits in Status to prevent an interrupt loop,
 * putting the processor in kernel mode, and re-enabling interrupts. The sample
 * code below cannot cover all nuances of this processing and is intended only
 * to demonstrate the concepts.
 */
```

```

mfc0 k0, C0_EPC          /* Get restart address */
sw   k0, EPCSave        /* Save in memory */
mfc0 k0, C0_Status      /* Get Status value */
sw   k0, StatusSave     /* Save in memory */
li   k1, ~IMbitsToClear /* Get IM bits to clear for this interrupt */
                                /* this must include at least the IM bit */
                                /* for the current interrupt, and may include */
                                /* others */
and  k0, k0, k1         /* Clear bits in copy of Status */
ins  k0, zero, S_StatusEXL, (W_StatusKSU+W_StatusERL+W_StatusEXL)
                                /* Clear KSU, ERL, EXL bits in k0 */
mtc0 k0, C0_Status      /* Modify mask, switch to kernel mode, */
                                /* re-enable interrupts */

/* Process interrupt here, including clearing device interrupt */

/*
 * To complete interrupt processing, the saved values must be restored
 * and the original interrupted code restarted.
 */

di   /* Disable interrupts - may not be required */
lw   k0, StatusSave     /* Get saved Status (including EXL set) */
lw   k1, EPCSave        /* and EPC */
mtc0 k0, C0_Status      /* Restore the original value */
mtc0 k1, C0_EPC         /* and EPC */
ehb  /* Clear hazard */
eret /* Dismiss the interrupt */

```

### 6.8.1.3 External Interrupt Controller Mode

External Interrupt Controller (EIC) mode redefines the way that the processor interrupt logic is configured to provide support for an external interrupt controller. The interrupt controller is responsible for prioritizing all interrupts, including hardware, software, timer, fast debug channel, and performance counter interrupts, and directly supplying to the processor the vector number of the highest priority interrupt.

EIC interrupt mode is in effect if all of the following conditions are true:

- $Config3_{VEIC} = 1$
- $IntCtl_{VS} \neq 0$
- $Cause_{IV} = 1$
- $Status_{BEV} = 0$

In EIC mode, the processor sends the state of the software interrupt requests ( $Cause_{IP1..IP0}$ ) and the timer, performance counter, and fast debug channel interrupt requests ( $Cause_{TI/PC/FDC}$ ) to the external interrupt controller, which prioritizes these interrupts in a system-dependent way with other hardware interrupts. The interrupt controller can be a hardwired logic block, or it can be configurable by control and status registers. This allows the interrupt controller to be more specific or more general as a function of the system environment and needs.

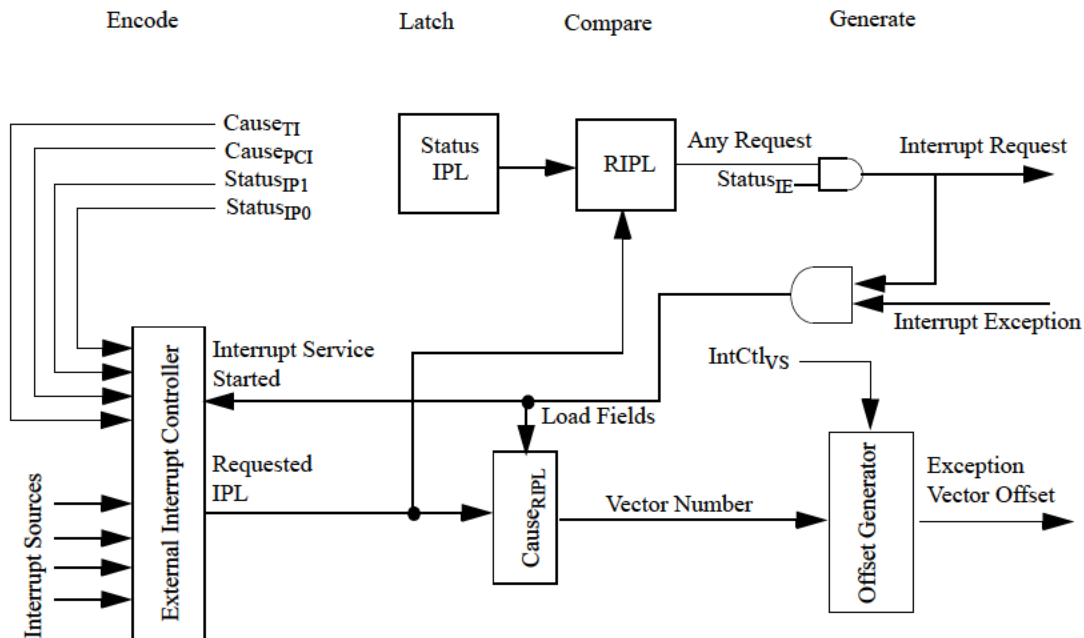
The external interrupt controller prioritizes its interrupt requests and produces the vector number of the highest priority interrupt to be serviced. The vector number, called the Requested Interrupt Priority Level (RIPL), is a 6-bit encoded value in the range 0..63, inclusive. The values 1..63 represent the lowest (1) to highest (63) RIPL for the

interrupt to be serviced. A value of 0 indicates that no interrupt requests are pending. The interrupt controller inputs this value on the 6 hardware interrupt lines, which are treated as an encoded value in EIC mode.

$Status_{IPL}$  (which overlays  $Status_{IM7..IM2}$ ) is interpreted as the Interrupt Priority Level (IPL) at which the processor is currently operating (a value of zero indicates that no interrupt is currently being serviced). When the interrupt controller requests service for an interrupt, the processor compares RIPL with  $Status_{IPL}$  to determine if the requested interrupt has a higher priority than the current IPL. If RIPL is strictly greater than  $Status_{IPL}$ , and interrupts are enabled ( $Status_{IE} = 1$ ,  $Status_{EXL} = 0$ , and  $Status_{ERL} = 0$ ), an interrupt request is signaled to the pipeline. When the processor starts the interrupt exception, it loads RIPL into  $Cause_{RIPL}$  (which overlays  $Cause_{IP7..IP2}$ ) and signals the external interrupt controller to notify it that the request is being serviced. The interrupt exception uses the value of  $Cause_{RIPL}$  as the vector number. Because  $Cause_{RIPL}$  is only loaded by the processor when an interrupt exception is signaled, it is available to software during interrupt processing.

The operation of EIC interrupt mode is shown in Figure 6.7.

**Figure 6.7 Interrupt Generation for External Interrupt Controller Interrupt Mode**



A typical software handler for EIC mode bypasses the entire sequence of code following the `IV exception` label shown for the compatibility-mode handler above. Instead, the hardware performs the prioritization, dispatching directly to the interrupt processing routine.

A nested interrupt is similar to that shown for compatibility mod. It also need only copy  $Cause_{RIPL}$  to  $Status_{IPL}$  to prevent lower priority interrupts from interrupting the handler. Here is an example of such a routine:

```
NestedException:
/*
 * Nested exceptions typically require saving the EPC and Status registers,
 * disabling the appropriate IM bits in Status to prevent an interrupt loop,
 * putting the processor in kernel mode, and re-enabling interrupts.
 * The sample code below can not cover all nuances of this processing and is
 * intended only to demonstrate the concepts.
 */
```

```

mfc0 k1, C0_Cause      /* Read Cause to get RIPL value */
mfc0 k0, C0_EPC        /* Get restart address */
srl  k1, k1, S_CauseRIPL /* Right justify RIPL field */
sw   k0, EPCSave       /* Save in memory */
mfc0 k0, C0_Status     /* Get Status value */
sw   k0, StatusSave    /* Save in memory */
ins  k0, k1, S_StatusIPL, 6 /* Set IPL to RIPL in copy of Status */
ins  k0, zero, S_StatusEXL, (W_StatusKSU+W_StatusERL+W_StatusEXL)
                                     /* Clear KSU, ERL, EXL bits in k0 */
mtc0 k0, C0_Status     /* Modify IPL, switch to kernel mode, */
                                     /* re-enable interrupts */

/* Process interrupt here, including clearing device interrupt */

/*
 * The interrupt completion code is identical to that shown for VI mode above.
 */

```

## 6.8.2 Generation of Exception Vector Offsets for Vectored Interrupts

For vectored interrupts (in either VI or EIC interrupt mode), a vector number is produced by the interrupt control logic. This number is combined with  $IntCtl_{VS}$  to create the interrupt offset, which is added to 0x200 to create the exception vector offset. For VI mode, the vector number is in the range 0..7, inclusive. For EIC interrupt mode, the vector number is in the range 1..63, inclusive (0 being the encoding for “no interrupt”). The  $IntCtl_{VS}$  field specifies the spacing between vector locations. If this value is zero (the default reset state), the vector spacing is zero and the processor reverts to Interrupt Compatibility mode. A non-zero value enables vectored interrupts. Table 6.25 shows the exception vector offset for a representative subset of the vector numbers and values of the  $IntCtl_{VS}$  field.

**Table 6.25 Exception Vector Offsets for Vectored Interrupts**

Vector Number	Value of $IntCtl_{VS}$ Field				
	5'b00001	5'b00010	5'b00100	5'b01000	5'b10000
0	0x0200	0x0200	0x0200	0x0200	0x0200
1	0x0220	0x0240	0x0280	0x0300	0x0400
2	0x0240	0x0280	0x0300	0x0400	0x0600
3	0x0260	0x02C0	0x0380	0x0500	0x0800
4	0x0280	0x0300	0x0400	0x0600	0x0A00
5	0x02A0	0x0340	0x0480	0x0700	0x0C00
6	0x02C0	0x0380	0x0500	0x0800	0x0E00
7	0x02E0	0x03C0	0x0580	0x0900	0x1000
		• • •			
61	0x09A0	0x1140	0x2080	0x3F00	0x7C00
62	0x09C0	0x1180	0x2100	0x4000	0x7E00
63	0x09E0	0x11C0	0x2180	0x4100	0x8000

The general equation for the exception vector offset for a vectored interrupt is:

```
vectorOffset ← 0x200 + (vectorNumber × (IntCtlVS || 0b00000))
```

### 6.8.3 Global Interrupt Controller

The Global Interrupt Controller (GIC) handles the routing and masking of local interrupts, such as the timer, performance counter, fast debug channel interrupts, inter-processor interrupts, and external interrupts. This block can be configured to support various numbers of external interrupts and to support any of the CPU interrupt modes.

An interactive GUI is available to simplify the setup of desired event-routing through the GIC. The tool outputs a C-language function covering all required programming registers of the GIC.



# Power Management and the Cluster Power Controller

This chapter describes the Cluster Power Controller (CPC) included in the interAptiv Multiprocessing System. The CPC organizes bootstrap, reset, tree root clock gating, and power gating of CPUs. The CPC also manages power cycling, reset, and clock gating of the Coherence Manager, dependent on the individual core status and shutdown policy.

The interAptiv Multiprocessing System includes an option for another version of the CPC that has reduced functionality. This version is called CPC-basic. At IP configuration time, the user selects between the standard CPC and the CPC-basic. For more information on the CPC-basic functionality, refer to [Section 7.8, "CPC-Basic"](#).

The chapter contains the following sections:

- [Section 7.1 “Introduction to the Cluster Power Controller”](#)
- [Section 7.2 “CPC Register Programming”](#)
- [Section 7.3 “Cluster Power Controller Address Map”](#)
- [Section 7.4 “Cluster Power Controller Commands”](#)
- [Section 7.5 “interAptiv Core Power Management Options”](#)
- [Section 7.6 “interAptiv Core Clock Gating”](#)
- [Section 7.7 “interAptiv Core Power Gating”](#)

## 7.1 Introduction to the Cluster Power Controller

The Cluster Power Controller (CPC) works in conjunction with the power management features of the individual interAptiv cores to provide a comprehensive power management scheme.

The main purpose of the Cluster Power Controller (CPC) is to manage static leakage and dynamic power consumption based on system-level power states assigned to the individual components of the interAptiv Multiprocessing System. As such, the CPC acts as a programmable platform peripheral, accessible through cluster CPU software and SOC-level hardware protocols.

The CPC is an integral part of the coherent cluster and is designed to bootstrap, reset, tree root clock-gate and power-gate cluster CPUs and the Coherence Manager. Implementors may or may not choose to support some or all of the physical features the CPC is architected to control. The following physical power-management features can be selected independently:

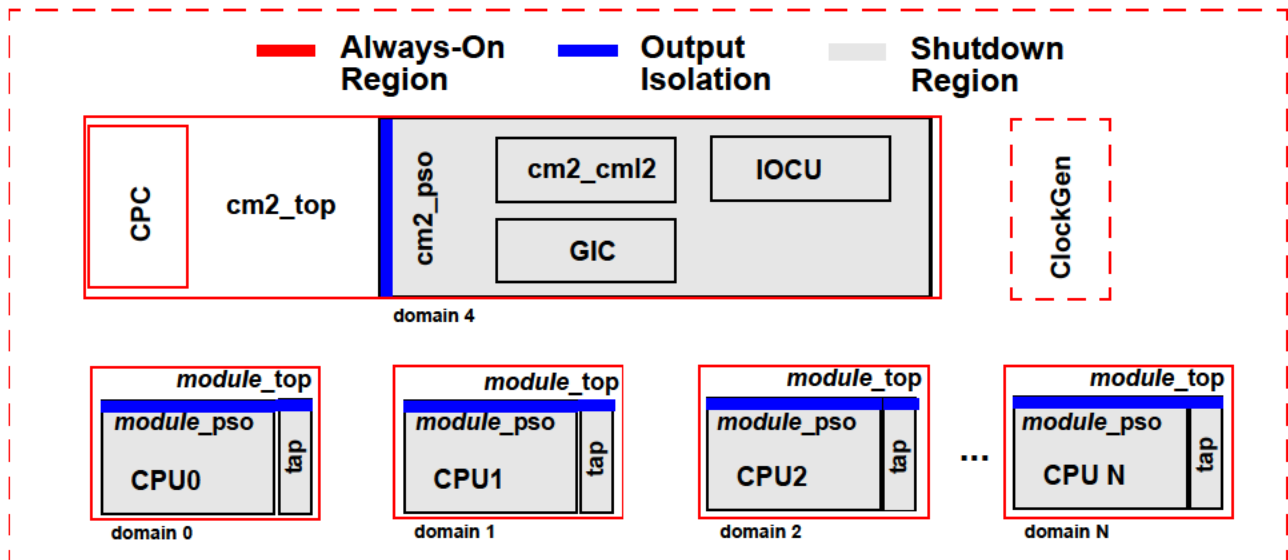
- **Power gating of selected CPUs and/or the CM.** Supported by industry-standard physical design flows, supply voltage of individual power domains can be switched on-chip. Currently, the Common Power Format (CPF) and Unified Power Format (UPF) are provided for a seamless front to back-end design flow. Besides CPF/UPF compliant EDA tools, standard cell libraries are required to provide power-gating header or footer cells, as well as isolate-high and isolate-low cells to separate unpowered domains from their active surroundings. The CPC provides a front-end RTL simulation environment and diagnostics to verify power-gating behavior.

- **Tree root clock gating.** Independent of CPU internal power-management features such as register-bank level clock gating and the sleep and doze modes, the CPC provides controls to gate clocks directly at or after the PLL in order to quiesce the entire clock tree of a CPU. CPC clock-gating signals are designed to bridge large clock insertion delays and are controlled through system-level power states.

In addition to power-management functions, the CPC also acts as reset and bootstrap controller of the Multiprocessing System (MPS) to initialize cores as they become operational, or re-initialize them upon system-level requests. The CPC also facilitates EJTAG debug probe access to cores by detecting the connection of a probe and enabling cores to respond to debug interrupt requests.

### 7.1.1 Power Domains of the interAptiv Multiprocessing System

Figure 7.1 interAptiv Multiprocessing System Power Domains



To individually power gate each core, independently controlled power domains are introduced to the interAptiv core. RTL simulation as well as physical implementation of the CPS support five distinct domains, cpu0-N and the Coherence Manager. These components are intended to be implemented with power rail switch cells to allow shutdown. Each controllable domain also is required to drive isolation values towards the system. This ensures proper logic values from shutdown domain boundaries into powered surroundings.

The CPS top level can be implemented to belong to a voltage scaled supply domain. This enables dynamic voltage and frequency scaling over the full CPS with shutdown features for individual subdomains.

With shutdown of all cores, the Coherence Manager becomes inactive unless IOCU traffic is requested. The CPC provides programmable power down for these components.

Level 2 cache is part of the CPS. However, power management of the L2 cache is not handled by the CPC. The CMP cluster implementation will ensure that power-down of cores and Coherence Manager does not affect L2 status.

### 7.1.2 Operating Level Transitions

To reach power-down and clock-off mode, software and hardware are required to go through a sequence of steps on each operating level to reach the next level.



### 7.1.2.1 Coherent to Non-Coherent Mode Transition

To leave the coherent domain and operate independently or prepare for shutdown, the following sequence should be followed:

1. Switch to non-coherent CCA.
2. Flush dirty data from data cache using IndexWritebackInvalidate CACHE instruction on all lines in the cache.
3. If the instruction cache contains lines that are expected to be maintained by software as coherent (via globalized CACHE instructions), and the CPU is not going to go through a reset sequence, the instruction cache should be flushed using IndexInvalidate CACHE instructions.
4. Write GCR\_CL\_COHERENCE (Core Local GCR address 0x0008). Write 0 to all bits except bit for "self", which should stay set to 1. This is required so that the core can issue a coherent SYNC (step 6) to make sure all previous interventions are complete.
5. Read GCR\_CL\_COHERENCE (ensures step 4 has completed).
6. Issue Coherent SYNC (intervention-only SYNC is fine).
7. Write 0 to GCR\_CL\_COHERENCE to completely remove core from coherence domain.
8. Read GCR\_CL\_COHERENCE to ensure step 7 is complete.

### 7.1.2.2 Non-Coherent to Coherent Mode Transition

An independently operating core becomes a member of a coherent cluster.

- Caches must be initialized first (since last reset)
- There should be no data in the caches that will later be accessed coherently. Non-coherent data is treated as exclusive/modified which can lead to violations of the coherence protocol if other caches have copies of the data.
- The GCR local coherence control register is programmed to add the core to the coherent domain.
- Switch to coherent Cache Coherence Attribute (CCA).
- Regular coherent programs can now start on this core.

### 7.1.2.3 Non-Coherent to Power Down Mode Transition

A core which is not member of a coherent domain is powered down. NOTE: When an EJTAG probe is detected, the CPC will prevent power down to preserve the connectivity of the TAP scan chain. A power-down command will instead cause the core to enter clock off mode.

- The GIC might be programmed to re-route interrupts away from this core.
- The CPC must be programmed to enter power-down mode.
- Core outputs are held inactive towards the CM. Completion of pending bus traffic is awaited and start of new traffic prevented using the *SI\_LPReq* protocol.
- The CPC initiates the clock and power shutdown micro-sequence.

### 7.1.2.4 Non-Coherent to Clock Off Mode Transition

A core is disconnected from bus and stops operation. Dynamic power consumption is removed.

- Programming a CPC ClkOff command will disable the clock tree root for this core.
- Core outputs are held inactive towards the CM. Completion of pending bus traffic is awaited and start of new traffic prevented using the *SI\_LPReq* protocol.
- The GIC might be programmed to re-route interrupts for this core to others.

#### 7.1.2.5 Clock Off to Power Down Mode Transition

Power supply is removed from a disconnected core. Dynamic and leakage power is removed.

- The CPC must be programmed to enter power-off mode.
- The CPC initiates the clock and power shutdown micro-sequence.

#### 7.1.2.6 Clock Off to Non-Coherent Mode Transition

A disconnected core is reconnected to the bus and starts operation.

- The CPC command register is programmed to bring the core back on-line. A CPC\_PwrUp command will let the core resume operation immediately, or, if a Reset command given, go through a reset sequence before becoming operational.
- If the core bus was isolated due to earlier power modes, this isolation is removed.
- The clock is applied and the core starts executing instructions.

#### 7.1.2.7 PowerDown to Non-Coherent Mode Transition

A core is powered up and becomes operational.

- The GCR local coherence control register must be set inactive for this core. Powering up into a coherent state with uninitialized caches may corrupt coherent data.
- Software on another core can send a PwrUp or Reset command for this core or an SOC hardware signal can request for the CPC to schedule a power-up sequence targeting non-coherent mode.
- The CPC will schedule a power-up sequence and the core becomes operational outside the coherent domain. After the core becomes operational, execution continues at the boot vector provided while power-up mode reset. NOTE: reset is not automatically applied unless the core really was in the power-down state prior to a PwrUp command or hardware PwrUp signal.
- The GIC might be reprogrammed to perform interrupt routing to this core.

## 7.2 CPC Register Programming

This section describes some of the programming functions that can be performed via the CPC registers.

### 7.2.1 Requestor Access to CPC Registers

The CPC allows up to eight requestor's in a system. A requestor can be either a core or an IOCU. The interAptiv core allows up to 6 requestors in a multiprocessing system; four cores and two IOCU's.

The requestor's may not have unrestricted access to the CPC registers. During boot time, software determines which requestor's are provided access to the CPC registers by programming the 8-bit *CPC\_ACCESS\_EN* field of the *Global CPC Access Privilege* register located at offset 0x000. Each bit in this field corresponds to a specific requestor.

The MIPS default for this field is 0xFF, meaning that all requestor's in the system have access to the CPC register set. To disable access to the registers for a particular requestor, software need only clear the corresponding bit of this field to zero and all write requests to the CPC registers by that requestor will be ignored.

## 7.2.2 Global Sequence Delay Count

The Sequence Delay register (*CPC\_SEQDEL\_REG*) located at offset 0x0008 in the CPC Global Control Block, contains a 10-bit field that describes the number of clock cycles each domain micro-sequencer will take to advance. It describes a set of worst-case timing of the physical implementation and is used to ensure electrical and bus protocol integrity. Typically, the *CPC\_SEQDEL\_REG* contents would be defined at IP configuration time. However, runtime write capability allows fine tuning to optimize sequencer timing. Domain sequencing begins once the RAILDELAY field has counted down to zero. Refer to [Section 7.2.3, "Rail Delay"](#) for more information.

The 10-bit MICROSTEP field is encoded as follows:

**Table 7.1 Encoding of MICROSTEP Field**

Encoding	Description
0x000	1-cycle delay
0x001	2-cycle delay
0x002	3-cycle delay
0x003	4-cycle delay
0x004	5-cycle delay
.....	.....
0x3FD	1022-cycle delay
0x3FE	1023-cycle delay
0x3FF	1024-cycle delay

Note that the physical implementation might not allow power sequence micro steps to advance with full cluster speed. At cluster cold start, the counter divides cluster frequency by a hardcoded IP configuration value to derive a micro step width.

## 7.2.3 Rail Delay

The Rail Delay register (*CPC\_RAIL\_REG*) located at offset 0x010 in the CPC Global Register Block contains a 10-bit counter field (*RAILDELAY*) used to schedule delayed start of power domain sequencing after the *RailEnable* signal has been activated by the CPC. This allows the CPC to compensate for slew rates at the gated rail, since hardware interlocks such as *SI\_VddOk* are either unavailable or don't reflect to complete power up time of a domain.

The 10-bit counter value delays the power-up sequence per domain after the *SI\_RailStable* and *VddOK* signals become active. The power-up micro-sequence starts after RAILDELAY has been loaded into the internal counter and a count-down to zero has concluded.

After completion of the domain power-up micro-sequence, the DomainReady signal is raised and can be used for domain daisy-chaining.

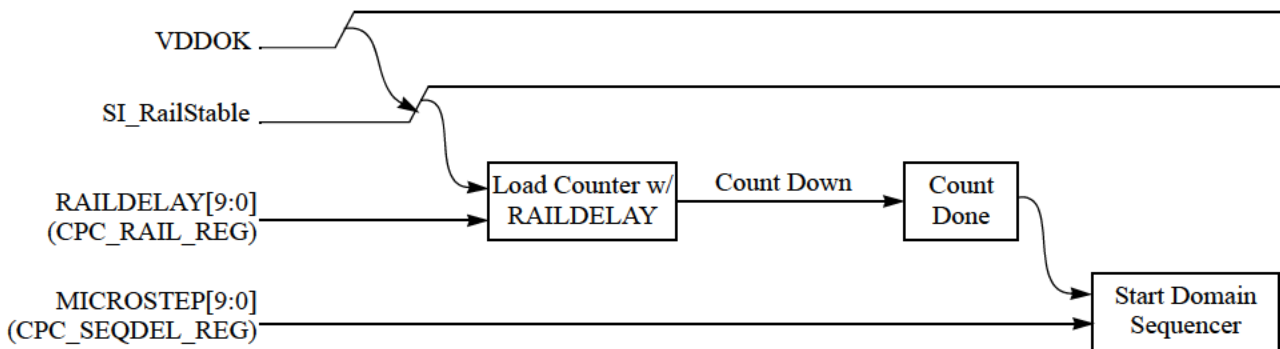
At IP configuration time, the contents of the *CPC\_RAIL\_REG* register are preset. However, for fine tuning, the register can be written at run time.

The 10-bit RAILDELAY field is encoded as follows:

**Table 7.2 Encoding of RAILDELAY Field**

Encoding	Description
0x000	1-cycle delay
0x001	2-cycle delay
0x002	3-cycle delay
0x003	4-cycle delay
0x004	5-cycle delay
.....	.....
0x3FD	1022-cycle delay
0x3FE	1023-cycle delay
0x3FF	1024-cycle delay


**Figure 7.2 Relationship Between RAILDELAY and MICROSTEP During Power-Up Sequence**



### 7.2.4 Reset Delay

Within the power-up micro-sequence, reset is applied. Typically, reset is active until the domain responds by asserting the internal *PB\_Reset\_N* signal. However, the *CPC\_RESETLEN\_REG* allows reset to be extended beyond the assertion of *PB\_Reset\_N*. The down-counter starts after the sequencer has detected the assertion of *PB\_Reset\_N*. Domains without a *PB\_Reset\_N* signal could tie this input low or connect it to an inverted reset signal.

**Figure 7.3 Extending the Reset Sequence Beyond the Assertion of the Reset Signal**



## 7.2.5 Executing a Power Sequence

The power sequence for the CPC block support the following commands:

- **ClockOff:** This command causes the domain to cycle into clock-off mode. It disables the clock to this power domain.
- **PwrDown:** This command uses the setup values in the CPC\_STAT\_CONF\_REG register.
- **PwrUp:** This command uses the setup values in the CPC\_STAT\_CONF\_REG register.
- **Reset:** When this command is issued, the domain is reset if it is in non-coherent mode.

A command can be executed in the local core by writing an encoded value to bits 3:0 of the Command register (CPC\_CL\_CMD\_REG) of the Core-Local block located at offset address 0x000. To write a command to another core, bits 3:0 of the Command register (CPC\_CO\_CMD\_REG) in the Core-Other block is used.

## 7.2.6 Accessing Another Core

To access another core, the number of the core to be accessed is programmed into bits 23:16 of the Core-Other Addressing register (CPC\_CL\_OTHER\_REG) located at offset 0x010 of the Core-Local block. This field selects the core number of the register set to be accessed in Core-Other address space. Refer to [Section 7.3.4.2, "Core-Other Addressing Register"](#) for more information.

## 7.3 Cluster Power Controller Address Map

The CPC uses memory locations within the global, core-local, and core-others address space. The CPC location within the CPU address map is determined by the GCR\_CPC\_BASE register. All address locations in this document are relative to this base address.

In [Table 7.3](#), all registers are accessed using 32-bit aligned uncached load/stores. In addition, the block offsets shown are relative to bits 31:15 of the GCR\_CPC\_Base register located in the CM2. Refer to Chapter 8, *CM2 Global Control Registers* for more information on this register.

**Table 7.3 CPC Address Map (Relative to GCR\_CPC\_BASE[31:15])**

Block Offset	Size (bytes)	Description
0x0000 - 0x1FFF	8 KB	<b>Global Control Block.</b> Contains registers pertaining to the global system functionality. This address section is visible to all CPUs.
0x2000 - 0x3FFF	8 KB	<b>Core-Local Control Block.</b> Aliased for each interAptiv core. Contains registers pertaining to the core issuing the request. Each core has its own copy of registers within this block.
0x4000 - 0x5FFF	8 KB	<b>Core-Other Control Block.</b> Aliased for each interAptiv core. This block of addresses gives each Core a window into another Core's Local Control Block. Before accessing this space, the <i>Core-Other Addressing Register</i> in the Local Control Block must be set to the CORENum of the target Core.

### 7.3.1 Block Offsets Relative to the Base Address

The block offsets for each of the three blocks listed in [Table 7.3](#) above are relative to a CPC base address and can be located anywhere in physical memory. The base address is a 17-bit value that is programmed into the `GCR_CPC_BASE` field of the *GCR CPC Base* register located at offset address 0x0088 in the Global Control Block of the CM2 registers. Note that this Global Control Block is different from the one listed in [Table 7.3](#) above. Refer to the *GCR\_CPC\_BASE Register* in Chapter 8, *CM2 Global Control Registers* for more information on this register.

To determine the physical address of each block listed in [Table 7.3](#), the base address written to the `GCR_CPC_BASE Register` this value would be added to the CPC block offset ranges to derive the absolute physical address as shown in [Table 7.4](#). Note that an example base address of 0x1BDE\_0 is used for these calculations.

**Table 7.4 Example Physical Address Calculation of the CPC Register Blocks**

Example Base Address		GCR Block Offset		Absolute Physical Address	Size (bytes)	Description
0x1BDE_0	+	0x0000 - 0x1FFF	=	0x1BDE_0000 - 0x1BDE_1FFF	8 KB	CPC Global Control Block.
0x1BDE_0	+	0x2000 - 0x3FFF	=	0x1BDE_2000 - 0x1BDE_3FFF	8 KB	CPC Core-Local Control Block.
0x1BDE_0	+	0x4000 - 0x5FFF	=	0x1BDE_4000 - 0x1BDE_5FFF	8 KB	CPC Core-Other Control Block.

### 7.3.2 Register Offsets Relative to the Block Offsets

In addition to the block offsets, the register offsets provided in each register description of this chapter are relative to the block offsets shown in [Table 7.4](#) above. To determine the physical address of each register, the base address programmed into the `GCR_CPC_BASE` register is added to the corresponding CPC block offset plus the actual register offset to derive the absolute physical address as shown in [Table 7.5](#). In this table an example base address of 0x1BDE\_0 is used.

**Table 7.5 Absolute Address of Individual CPC Global Control Block Registers**

MIPS Default Base		Global Register Block Offset		Global Register Offset		Absolute Physical Address	Global Control Register
0x1BDE_0	+	0x0000	+	0x0000	=	0x1BDE_0000	CPC Access Privilege.
0x1BDE_0	+	0x0000	+	0x0008	=	0x1BDE_0008	CPC Global Sequence Delay.
0x1BDE_0	+	0x0000	+	0x0010	=	0x1BDE_0010	CPC Rail Delay.
0x1BDE_0	+	0x0000	+	0x0018	=	0x1BDE_0018	CPC Reset Length.
0x1BDE_0	+	0x0000	+	0x0020	=	0x1BDE_0020	CPC Revision.

[Table 7.6](#) shows the absolute physical addresses for the CPC Core-Local block. In this table an example base address of 0x1BDE\_0 is used.

**Table 7.6 Absolute Address of Individual CPC Core-Local Block Registers**

MIPS Default Base		Core-Local Register Block Offset		Core-Local Register Offset		Absolute Physical Address	Core-Local Register
0x1BDE_0	+	0x2000	+	0x0000	=	0x1BDE_2000	CPC Core-Local Command.
0x1BDE_0	+	0x2000	+	0x0008	=	0x1BDE_2008	CPC Core-Local Status and Configuration.

**Table 7.6 Absolute Address of Individual CPC Core-Local Block Registers(continued)**

MIPS Default Base		Core-Local Register Block Offset		Core-Local Register Offset		Absolute Physical Address	Core-Local Register
0x1BDE_0	+	0x2000	+	0x0010	=	0x1BDE_2010	CPC Core-Other Addressing.

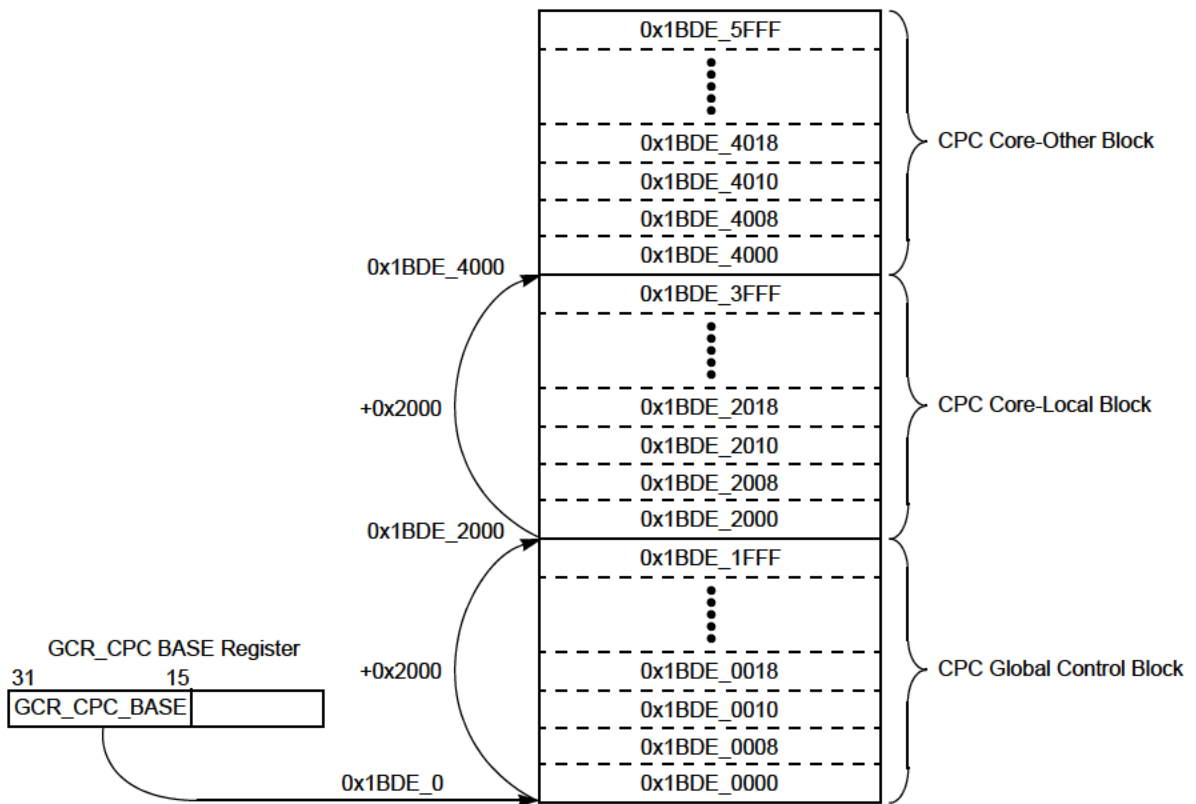
Table 7.6 shows the absolute physical addresses for the CPC Core-Other block. In this table an example base address of 0x1BDE\_0 is used.

**Table 7.7 Absolute Address of Individual CPC Core-Other Block Registers**

MIPS Default Base		Core-Other Register Block Offset		Core-Other Register Offset		Absolute Physical Address	Core-Other Register
0x1BDE_0	+	0x4000	+	0x0000	=	0x1BDE_4000	CPC Core-Other Command.
0x1BDE_0	+	0x4000	+	0x0008	=	0x1BDE_4008	CPC Core-Other Status and Configuration.
0x1BDE_0	+	0x4000	+	0x0010	=	0x1BDE_4010	CPC Core-Other Addressing.

This concept is described in Figure 7.4 below. In this figure an example base address of 0x1BDE\_0 is used.

**Figure 7.4 CPC Register Addressing Scheme Using an Example Base Address of 0x1BDE\_0**



### 7.3.3 Global Control Block Register Map

All registers in the Global Control Block are 32 bits wide and should only be accessed using aligned 32-bit uncached load/stores. Reads from unpopulated registers in the CPC address space return 0x0, and writes to those locations are silently dropped without generating any exceptions.

**Table 7.8 Global Control Block Register Map (Relative to Global Control Block offset)**

Register Offset in Block	Name	Type	Description
0x000	CPC Global CSR Access Privilege Register ( <i>CPC_ACCESS_REG</i> )	R/W	Controls which cores can modify the CPC Registers.
0x008	CPC Global Sequence Delay Counter ( <i>CPC_SEQDEL_REG</i> )	R/W	Time between microsteps of a CPC domain sequencer in CPC clock cycles.
0x010	CPC Global Rail Delay Counter Register ( <i>CPC_RAIL_REG</i> )	R/W	Rail power-up timer to delay CPS sequencer progress until the gated rail has stabilized.
0x018	CPC Global Reset Width Counter Register ( <i>CPC_RESETLEN_REG</i> )	R/W	Duration of any domain reset sequence.
0x020	CPC Global Revision Register ( <i>CPC_REVISION_REG</i> )	R	RTL Revision of CPC
0x028 0x0F8	CPC Global RESERVED registers.	-	For Future Extensions

#### 7.3.3.1 Global CSR Access Privilege Register

**Table 7.9 CPC Global CSR Access Privilege Register (CPC\_ACCESS\_REG Offset 0x000)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:8	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
CM_ACCESS_EN	7:0	Each bit in this field represents a power domain CPU.  If the bit is set, that requester is able to write to the CPC registers (this includes all registers within the Global, Core-Local and Core-Other blocks. If the bit is clear, any write request from that requestor to the CPC registers (Global, Core-Local, Core-Other) will be dropped.	R/W	0xff

The Access privilege register configures the CPU access rights towards CPC programming registers. Its function is defined equally to the GCR Access Privilege Register.



### 7.3.3.2 Global Sequence Delay Counter

The *CPC\_SEQDEL\_REG* describes globally the number of clock cycles each domain micro-sequencer will take to advance. It describes a set of worst-case timing of the physical implementation and is used to ensure electrical and bus protocol integrity. Mainly, buffer tree delays on *SI\_Isolate* and/or *SI\_RailEnable* can be used to set proper micro sequencer delay values.

Typically, the *CPC\_SEQDEL\_REG* contents would be defined at IP configuration time. However, runtime write capability allows fine tuning to optimize sequencer timing.

**Table 7.10 Global Sequence Delay Counter Register (CPC\_SEQDEL\_REG, Offset 0x008)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:10	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
MICROSTEP	9:0	This field reflects the delay in clock cycles, taken by each power domain micro-sequencer to advance between atomic micro steps. Cycles/Step = MICROSTEP[9:0] value + 1; 0 => 1cycle, 1 => 2cycles... Physical implementation might not allow power sequence micro steps to advance with full cluster speed. At cluster cold start, the counter divides cluster frequency by a hardcoded IP configuration value to derive a micro step width.	R/W	IP Configuration Value

### 7.3.3.3 Global Rail Delay Counter

The *CPC\_RAIL\_REG* represents a 10-bit counter register to schedule delayed start of domain operation after the *RailEnable* signal has been activated by the CPC. This allows to compensate for slew rates at the gated rail, since hardware interlocks such as *SI\_VddOk* are either unavailable or don't reflect to complete power up time of a domain.

At IP configuration time, the contents of *CPC\_RAIL\_REG* is preset. However, for fine tuning, the register can be written at run time.

**Table 7.11 Global Rail Delay Counter Register (CPC\_RAIL\_REG, Offset 0x010)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:10	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
RAILDELAY	9:0	10-bit counter value to delay power-up sequence per domain after <i>RailStable</i> and <i>VddOK</i> signals became active. The power-up micro-sequence starts after RAILDELAY has been loaded into the internal counter and a counted down to zero has concluded. After completion of the domain power-up micro-sequence, the <i>DomainReady</i> signal is raised and can be used for domain daisy-chaining.	R/W	IP Configuration Value

### 7.3.3.4 Global Reset Width Counter

Within the power-up micro-sequence, reset is applied. Typically, reset is active until the domain responds with *PB\_Reset\_N* feedback. However, the *CPC\_RESETLEN\_REG* allows reset to be extended beyond the *ResetN* feedback, or in case the reset feedback is unavailable. Counting down will start after the sequencer has received the *PB\_Reset\_N* feedback. Domains without *PB\_ResetN* feedback could tie this input low or connect it to an inverted reset signal.

**Table 7.12 Global Reset Width Counter Register (CPC\_RESETLEN\_REG, Offset 0x018)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:10	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
RESETLEN	9:0	10-bit counter value to extend reset duration beyond <i>PB_Reset_N</i> feedback. The domain behavior after reset is determined by the domain local setup register.	R/W	IP Configuration Value

### 7.3.3.5 Revision Register

**Table 7.13 Revision Register (CPC\_REVISION\_REG, Offset 0x020)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:16	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	
CPC_TYPE	15:14	This field reflects the type of CPC implemented and is encoded as follows: 00: Standard CPC 01: CPC-basic 10 - 11: Reserved	R	IP Configuration Value
MAJOR_REV	13:8	This field reflects the major revision of the CPC block. A major revision might reflect the changes from one product generation to another.	R	Preset
MINOR_REV	7:0	This field reflects the minor revision of the CPC block. A minor revision might reflect the changes from one release to another.	R	Preset

## 7.3.4 Local and Core-Other Control Blocks

All registers in the CPC Local Control Block are 32 bits wide and should only be accessed using aligned 32-bit uncached load/stores. Reads from unpopulated registers in the CPC address space return 0x0, and writes to those locations are silently dropped without generating any exceptions.

A set of these registers exists for each core in the interAptiv MPS. These registers can also be accessed from other cores by first writing the *CPC Core Other Addressing Register* (in the Core-Local Control Block) with the proper CoreNum and then accessing these registers using the Core Other address space.

The register offsets shown are relative to the offsets listed in [Table 7.14](#).

**Table 7.14 Core-Local Block Register Map**

Register Offset in Block	Name	Type	Description
0x000	CPC Local Command Register ( <i>CPC_CL_CMD_REG</i> )	R/W	Places a new CPC domain state command into this individual domain sequencer. This register is not available within the CM sequencer. Writes to the CM CMD register are ignored while reads will return zero.
0x008	CPC Local Status and Configuration register ( <i>CPC_CL_STAT_CONF_REG</i> )	R/W	Individual domain power status and domain configuration register. Reflects domain micro-sequencer execution. Initiates micro-sequencer after status register programming. Reflects command execution status.
0x010	CPC Core Other Addressing Register ( <i>CPC_CL_OTHER_REG</i> )	R/W R/O for CM2	Used to access local registers of another core.
0x018 0x0F8	CPC Local RESERVED registers	-	For Future Extensions

The register offsets shown are relative to the offsets listed in [Table 7.15](#).

**Table 7.15 Core-Other Block Register Map**

Register Offset in Block	Name	Type	Description
0x000	CPC Local Command Register ( <i>CPC_CO_CMD_REG</i> )	R/W	Places a new CPC domain state command into this individual domain sequencer. This register is not available within the CM sequencer. Writes to the CM CMD register are ignored while reads will return zero.
0x008	CPC Local Status and Configuration register ( <i>CPC_CO_STAT_CONF_REG</i> )	R/W	Individual domain power status and domain configuration register. Reflects domain micro-sequencer execution. Initiates micro-sequencer after status register programming. Reflects command execution status.
0x010	CPC Core Other Addressing Register ( <i>CPC_CO_OTHER_REG</i> )	R/W R/O for CM	Used to access local registers of another core.
0x018 0x0F8	CPC Local RESERVED registers	-	For Future Extensions

CPC Local register are used to set power-down conditions. After setup of conditions, the micro-sequencer can be activated through the command register. The execution of the micro-sequencer can be observed via the status register. Reading the status and configuration register retrieves the last executed command and status flags to reflect on recent commands given.

### 7.3.4.1 Command Register

**Table 7.16 Local Command Register (CPC\_CL[CO]\_CMD\_REG, Offset 0x000)**

Register Fields		Description	Read/Write	Reset State												
Name	Bits															
RESERVED	31:4	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0												
CMD	3:0	Requests a new power sequence execution for this domain. Read value is the last executed command.	R/W Not available in CM domain	0												
		<table border="1"> <thead> <tr> <th>Code</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>4'd1</td> <td> <p><b>ClockOff</b></p> <p>This command causes the domain to cycle into clock-off mode. It disables the clock to this power domain. Only successful if <i>SI_CoherenceEnable</i> and other protocol interlocks are observed. If not, the command remains inactive until the protocol barriers subside. After that, the command is executed.</p> <p>Depending on the current sequencer state, the command either causes power-up of a domain, or a domain leaves active duty to become inactive. A power-up leads to sequencer state U2, which will require the execution of a subsequent Reset or PwrUp command to make this domain operational.</p> </td> </tr> <tr> <td>4'd2</td> <td> <p><b>PwrDown</b></p> <p>this domain using setup values in CPC_STAT_CONF_REG. Only successful if <i>SI_CoherenceEnable</i> inactive and all protocol interlocks are observed. If not, the command remains inactive until the protocol barriers subside. Then, the command is executed.</p> </td> </tr> <tr> <td>4'd3</td> <td> <p><b>PwrUp</b></p> <p>this domain using setup values in CPC_STAT_CONF_REG. Usable only for Core-Others access. It is the software equivalent to <i>SI_PwrUp</i> hardware signal</p> </td> </tr> <tr> <td>4'd4</td> <td> <p><b>Reset</b></p> <p>This domain is reset if in non-coherent mode. After the domain has been reset, the domain becomes operational and the CMD field reads as PwrUp cmd.</p> </td> </tr> <tr> <td>Others</td> <td><b>Reserved</b></td> </tr> </tbody> </table>			Code	Meaning	4'd1	<p><b>ClockOff</b></p> <p>This command causes the domain to cycle into clock-off mode. It disables the clock to this power domain. Only successful if <i>SI_CoherenceEnable</i> and other protocol interlocks are observed. If not, the command remains inactive until the protocol barriers subside. After that, the command is executed.</p> <p>Depending on the current sequencer state, the command either causes power-up of a domain, or a domain leaves active duty to become inactive. A power-up leads to sequencer state U2, which will require the execution of a subsequent Reset or PwrUp command to make this domain operational.</p>	4'd2	<p><b>PwrDown</b></p> <p>this domain using setup values in CPC_STAT_CONF_REG. Only successful if <i>SI_CoherenceEnable</i> inactive and all protocol interlocks are observed. If not, the command remains inactive until the protocol barriers subside. Then, the command is executed.</p>	4'd3	<p><b>PwrUp</b></p> <p>this domain using setup values in CPC_STAT_CONF_REG. Usable only for Core-Others access. It is the software equivalent to <i>SI_PwrUp</i> hardware signal</p>	4'd4	<p><b>Reset</b></p> <p>This domain is reset if in non-coherent mode. After the domain has been reset, the domain becomes operational and the CMD field reads as PwrUp cmd.</p>	Others	<b>Reserved</b>
		Code			Meaning											
		4'd1			<p><b>ClockOff</b></p> <p>This command causes the domain to cycle into clock-off mode. It disables the clock to this power domain. Only successful if <i>SI_CoherenceEnable</i> and other protocol interlocks are observed. If not, the command remains inactive until the protocol barriers subside. After that, the command is executed.</p> <p>Depending on the current sequencer state, the command either causes power-up of a domain, or a domain leaves active duty to become inactive. A power-up leads to sequencer state U2, which will require the execution of a subsequent Reset or PwrUp command to make this domain operational.</p>											
		4'd2			<p><b>PwrDown</b></p> <p>this domain using setup values in CPC_STAT_CONF_REG. Only successful if <i>SI_CoherenceEnable</i> inactive and all protocol interlocks are observed. If not, the command remains inactive until the protocol barriers subside. Then, the command is executed.</p>											
		4'd3			<p><b>PwrUp</b></p> <p>this domain using setup values in CPC_STAT_CONF_REG. Usable only for Core-Others access. It is the software equivalent to <i>SI_PwrUp</i> hardware signal</p>											
4'd4	<p><b>Reset</b></p> <p>This domain is reset if in non-coherent mode. After the domain has been reset, the domain becomes operational and the CMD field reads as PwrUp cmd.</p>															
Others	<b>Reserved</b>															

**Table 7.17 Local Status and Configuration Register (CPC\_CL[CO]\_STAT\_CONF\_REG, Offset 0x008)**

Register Fields		Description	Read/Write	Reset State																								
Name	Bits																											
RESERVED	[31:24]	Reserved.	R	0																								
PWRUP_EVENT	23	The <i>SI_PowerUp</i> pin had been activated and caused the sequencer to cycle into power up state. The event also caused the sequencer to place a PwrUp command into the CMD field. Writing a 0 into the PWRUP_EVENT field will clear this bit.	R/W0	0																								
SEQ_STATE	[22:19]	Current domain sequencer state. State description: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Code</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>4'h0</td> <td>D0 - PwrDwn</td> </tr> <tr> <td>4'h1</td> <td>U0 - VddOK</td> </tr> <tr> <td>4'h2</td> <td>U1 - UpDelay</td> </tr> <tr> <td>4'h3</td> <td>U2 - UCkOff</td> </tr> <tr> <td>4'h4</td> <td>U3 - Reset</td> </tr> <tr> <td>4'h5</td> <td>U4 - ResetDly</td> </tr> <tr> <td>4'h6</td> <td>U5 - nonCoherent execution</td> </tr> <tr> <td>4'h7</td> <td>U6 - Coherent execution</td> </tr> <tr> <td>4'h8</td> <td>D1 - Isolate</td> </tr> <tr> <td>4'h9</td> <td>D3 - ClrBus</td> </tr> <tr> <td>4'ha</td> <td>D2 - DCkOff</td> </tr> </tbody> </table>	Code	State	4'h0	D0 - PwrDwn	4'h1	U0 - VddOK	4'h2	U1 - UpDelay	4'h3	U2 - UCkOff	4'h4	U3 - Reset	4'h5	U4 - ResetDly	4'h6	U5 - nonCoherent execution	4'h7	U6 - Coherent execution	4'h8	D1 - Isolate	4'h9	D3 - ClrBus	4'ha	D2 - DCkOff	R	0
Code	State																											
4'h0	D0 - PwrDwn																											
4'h1	U0 - VddOK																											
4'h2	U1 - UpDelay																											
4'h3	U2 - UCkOff																											
4'h4	U3 - Reset																											
4'h5	U4 - ResetDly																											
4'h6	U5 - nonCoherent execution																											
4'h7	U6 - Coherent execution																											
4'h8	D1 - Isolate																											
4'h9	D3 - ClrBus																											
4'ha	D2 - DCkOff																											
RESERVED	18	Reserved.	R	-																								
CLKGAT_IMPL	17	If set, this domain is implemented with clock tree root gating. If cleared, the CPC will still execute power-down/clock-off sequences if commanded; however, no physical clock gating is performed.	R	IP Configuration Value																								
PWRDN_IMPL	16	If set, this domain is implemented as power-gated. If cleared, the CPC will still execute power-down sequences if commanded; however, no physical power switching is performed.	R	IP Configuration Value																								
EJTAG_PROBE	15	An EJTAG probe connection event has been seen. The domain powers up if required and observes a reset sequence. Thereafter the core transitions into clock-off mode. After a probe has been seen once, the power domain will not assume power-off mode until this bit is written to zero or the CPC experiences a cold reset.	R/W0	0																								
Reserved	14:11	Reserved.	R	0																								
Reserved	10	Reserved.	R/W	1																								

**Table 7.17 Local Status and Configuration Register (CPC\_CL[CO]\_STAT\_CONF\_REG, Offset 0x008)**

Register Fields		Description	Read/Write	Reset State										
Name	Bits													
PWUP_POLICY	[9:8]	<p>Each CPC domain sequencer is hardwired through the <i>SI_ColdPwrUp</i> signal to either power up, remain power-gated, go into clock-off mode, or become operational. To influence the cold start behavior of the domain, three distinct policies can be wired for this domain:</p> <table border="1"> <thead> <tr> <th>Code</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>2'b00</td> <td>This CPU remains powered down after a system cold start. A later PwrUp or Reset command, or <i>SI_PwrUp</i> signal assertion will make this domain operational.</td> </tr> <tr> <td>2'b01</td> <td>Go into Clock-Off mode. Disables domain clock after power-up sequence. Core will wake up through a CPC PwrUp or Reset command or a <i>SI_PwrUp</i> signal assertion. In this Clock-Off mode, the core will not be initialized and its boundary isolation will be maintained.</td> </tr> <tr> <td>2'b10</td> <td>Power up this domain after system cold start. The CPU will be reset and become operational based on its boot vector contents.</td> </tr> <tr> <td>2'b11</td> <td>Reserved</td> </tr> </tbody> </table> <p>Within a processor cluster, CPU zero would power-up, while peer CPU 1-3 remain unpowered until released through a PwrUp commands. The PWUP_POLICY field reflects the hardwired <i>SI_ColdPwrUp</i> bus.</p>	Code	Meaning	2'b00	This CPU remains powered down after a system cold start. A later PwrUp or Reset command, or <i>SI_PwrUp</i> signal assertion will make this domain operational.	2'b01	Go into Clock-Off mode. Disables domain clock after power-up sequence. Core will wake up through a CPC PwrUp or Reset command or a <i>SI_PwrUp</i> signal assertion. In this Clock-Off mode, the core will not be initialized and its boundary isolation will be maintained.	2'b10	Power up this domain after system cold start. The CPU will be reset and become operational based on its boot vector contents.	2'b11	Reserved	R	<p>Hardwired IP Configuration Value</p> <p>CM domain is hard coded to powerUp if any CPU domain is powered up initially.</p>
Code	Meaning													
2'b00	This CPU remains powered down after a system cold start. A later PwrUp or Reset command, or <i>SI_PwrUp</i> signal assertion will make this domain operational.													
2'b01	Go into Clock-Off mode. Disables domain clock after power-up sequence. Core will wake up through a CPC PwrUp or Reset command or a <i>SI_PwrUp</i> signal assertion. In this Clock-Off mode, the core will not be initialized and its boundary isolation will be maintained.													
2'b10	Power up this domain after system cold start. The CPU will be reset and become operational based on its boot vector contents.													
2'b11	Reserved													
RESERVED	[7:5]	Reads zero. Writes ignored	R	0										
IO_TRFFC_EN	[4]	<p>Enable CM for stand alone IOCU traffic. Setting this bit changes the low power state of the CM power domain from PwrDwn to ClkOff. The <i>CM_IOPwrUp</i> signal can be used by an external device to enable the CM to perform IOCU data transfers without CPU activities.</p> <p>Deselecting IO_TRFFC_EN will power down the CM if all CPUs are powered down. In this case, <i>CM_IOPwrUp</i> signal activity is not observed by the CPC.</p> <p>A powered down CM domain will clear all preset CM/IOCU control registers. Powering up due to CPU power-up will send the CM/IOCU through a reset sequence, together with the CPU.</p>	<p>R/O for CPUs, read zero</p> <p>R/W for CM</p>	0										
CMD	3:0	Reflects most recent placed sequencer command. See definition in <i>CPC_CMD_REG</i> Table 7.3.4.1. The sequencer will overwrite the field after a Reset command, or <i>SI_PwrUp</i> signal caused power up of the domain. The command reads then as PwrUp.	R	0										

### 7.3.4.2 Core-Other Addressing Register

This register must be written with the correct CoreNum value before accessing the Core-Other address segment. This register is not available within the CM local domain. Read access to the CM *CPC\_OTHER\_REG* will yield zero. Writes are ignored.

**Table 7.18 Core-Other Addressing Register (CPC\_CL[CO]\_OTHER\_REG Offset 0x010)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:19	Reads as 0. Writes ignored. Must be written with a value of 0x000.	R	0
CORENUM	18:16	CoreNum of the register set to be accessed in the Core-Other address space.	R/W	0x0
RESERVED	15:0	Reads as 0. Writes ignored. Must be written with a value of 0x0000.	R	0

## 7.4 Cluster Power Controller Commands

The CPC provides a set of commands to establish a desired power domain state. CPC commands are:

- ClockOff** - a power domain is brought into ClockOff state as programmed into the *CPC\_CMD\_REG* Table 7.3.4.1. If the domain was powered down before, the power-on sequence is applied according to *CPC\_STAT\_CONF\_REG* settings. If the domain was active before and was in non-coherent operation, the domain is brought into ClockOff state D2. A domain in ClockOff state can be sent into operation using the PwrUp command. A ClockOff command given to a domain in coherent operation will remain inactive until the CPU has left the coherent mode of operation. Sending a ClkOff command to the CPC before a previous command completed will cause the CPC domain target to be redirected towards ClockOff. However, the previous steady state can be observed temporarily before the newly programmed state is reached.
- PwrDwn** - a power domain is powered down into state D0. *CPC\_STAT\_CONF\_REG* and *CPC\_CMD\_REG* settings determine the sequence observed by the CPC. Note, both register settings are observed dynamically. The sequencer will preempt an in flight command at the next steady state to execute the newly given command.
- PwrUp** - the execution of this command depends on the previous domain power state. If the domain is powered down to state D0, a PwrUp command will enable power for the domain and bring the domain into operational state U5. However, if *SL\_CoherenceEnable* is active, the domain will advance into state U6 - coherent operation. Please note, that a set of software initialization needs to complete to safely bring a non-coherent core into coherent state. If the previous power domain state was 'ClkOff', a PwrUp command will raise the domain state to either non-coherent or coherent operation, dependent on the GCR coherence status settings. This will be domain state U5 and U6 respectively.

When bringing a domain up after a PwrDwn command is executed, the Reset command is generally preferable to PwrUp. If the domain did not reach state D0 or was prevented from entering D0 because an EJTAG probe was connected, the CPC may identify that a reset is not required for PwrUp and will simply restart the clocks. This may be fine, but also may cause some problems. One common example where a reset is required is if the core enters an infinite loop after requesting PwrDwn.

A PwrUp command given to an active domain in non-coherent or coherent operation U5/U6 has no effect.

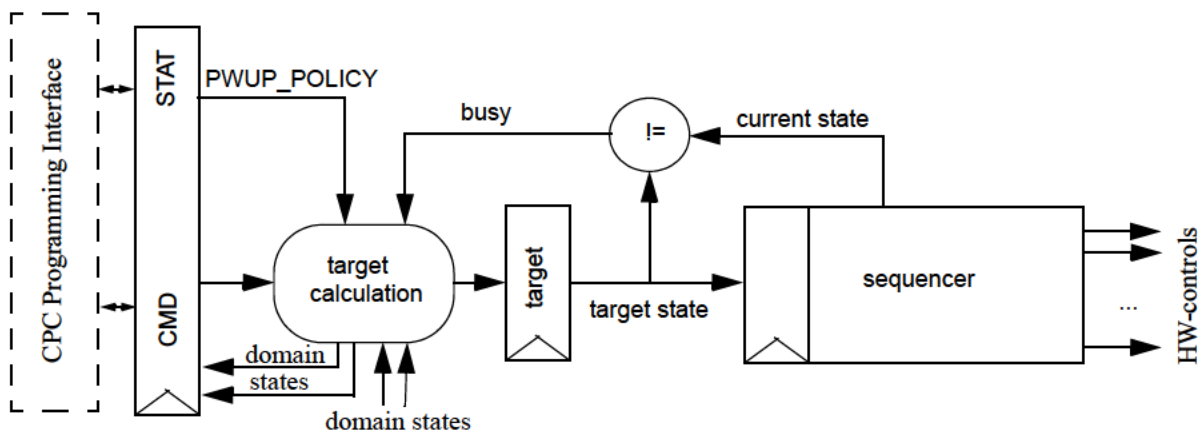
If a PwrUp command is given to the CPC while a previous command is still in flight, the command is placed in the CPC command register and is executed at the earliest possible state, i.e., when the sequencer has reached a non-transitional state.

The hardware *SI\_PwrUp* signal activated for this domain will always bring the core into power-up mode with enabled clocks. The *PWUP\_POLICY* settings of *CPC\_STAT\_CONF\_REG* have no effect on hardware wake-ups. Also, the hardware wake-up has priority over software commands. The *PWRUP\_EVENT* bit of *CPC\_STAT\_CONF\_REG* is set after a hardware power-up has been executed.

- **Reset** - this command allows a domain in non-coherent operation (state U5) to be reset. It also can be sent to a domain in power-down or clock-off mode. The domain will then become active, and a reset sequence is executed which leads to an operational steady state of the domain (U5 or U6, dependent on GCR programming).

Figure 7.5 details the CPC domain command execution. A command given to a CPC power domain will be translated into a domain target state, and the domain sequencer will progress towards this target. A new command is accepted as soon as a suitable state transition is found within the traversed states. Domain sequencer states translate directly to hardware control signals for reset and power gating, as depicted in Figure 7.5.

**Figure 7.5 CPC Command Execution**





## 7.5 interAptiv Core Power Management Options

In addition to the Cluster Power Controller described in the previous sections, MIPS Technologies provides a mechanism for reducing power in the interAptiv core depending on the work load. The conditions under which the interAptiv core is placed in power-down mode are determined by the SOC.

The information in the following sections should be used only when all cores in the system are shut down. The processor and cache states need not be saved for each core shut down as long as their is one core operation. However, once the last core is to be shutdown by the SOC, the following procedure can be used to save the processor state.

There are two basic options for power management in the interAptiv core.

1. Clock gating: Used to stop the clocks and put the core into sleep mode. Refer to [Section 7.6, "interAptiv Core Clock Gating"](#) for more information. In this mode the VDD levels are maintained and power is preserved, so no data is lost.
2. Power gating: Used to shut down power to selected parts of the interAptiv core. In this mode certain elements of the core, such as registers, caches, TLB, etc. are saved, allowing for a more efficient power-up process. Refer to [Section 7.7, "interAptiv Core Power Gating"](#) for more information.

## 7.6 interAptiv Core Clock Gating

Clock gating provides a way for the interAptiv core to shut down the core clock under certain conditions. The mechanism used to suspend and then resume the core clock depends on the power management options selected during the core configuration process. These options include;

- Enabling of ‘top level clock gating’
- Enabling of ‘fine grain clock gating’

### 7.6.1 Designs Implementing Top Level Clock Gating

Top level clock gating is provided as an option during the core configuration process. For designs implementing top level clock gating, there are two ways to place the interAptiv core into sleep mode.

- Instruction-controlled power management
- Register-controlled power management

#### 7.6.1.1 Instruction Controlled Clock Gating

Execution of the WAIT instruction can be used to place the interAptiv core into sleep mode. When the WAIT instruction is executed during normal operation, the interAptiv core completes all outstanding operations, then freezes the pipeline and asserts the SI\_SLEEP signal, indicating to external logic that the interAptiv core has entered sleep mode.

If top level clock gating is enabled, the processor turns off the internal clock to most of the interAptiv core automatically once SI\_SLEEP is asserted. The clock is maintained only for a small amount of logic that waits for an interrupt intended to bring the processor out of sleep mode. In addition to the interrupt logic, the following signals also remain active in sleep mode;

- SI\_INT[5:0]
- SI\_NMI
- SI\_RESET

- EJ\_DEBUGM

Once the clocks are suspended, the entire contents of the processor, including registers, caches, and TLB, are saved. Once the ‘wake’ interrupt is received, the processor restarts its internal clock and can resume normal operation within a few clock cycles. The ‘wake’ interrupt can be any enabled interrupt, NMI, or debug interrupt. This is the fastest and most efficient mechanism to transition the interAptiv core in and out of sleep mode.

Note that the SI\_RESET signal can also be used to exit sleep mode. However, assertion of SI\_RESET causes all internal data to be lost and the registers to revert back to their default values.

### 7.6.1.2 Register Controlled Clock Gating

In addition to instruction controlled clock gating, the MIPS architecture allows for software to initiate entry into sleep mode via the register interface. The *RP* bit in the *CPO Status* register can be set by software to indicate the desire to place the interAptiv core into sleep mode. Once this bit is set, hardware asserts the *SI\_RP* output signal.

On receipt of the *SI\_RP* signal, external logic can then decide whether to suspend or reduce the frequency of the interAptiv core accordingly. Note that this mechanism is different than instruction controlled clock gating in that the core does not determine whether the clock is suspended. Rather, external logic can decide to suspend the clock, reduce the clock frequency from its current level, or take no action.

**Table 7.19 Differences Between Instruction and Register Controlled Power Management**

Type	Trigger	Signal Asserted by Hardware	Clock Suspended by On-Die Hardware	Interrupt Detection During Sleep Mode
Instruction controlled clock gating	WAIT instruction	SI_SLEEP	Yes	Yes
Register controlled power management	Setting RP bit in CPO Status	SI_RP	No	Yes

### 7.6.1.3 Reduction of VDD During Sleep Mode

The information described above deals with clock gating only. In this example, during the time that the clocks are powered down, VDD remains at normal power levels. To obtain the maximum power savings during sleep mode, external logic can reduce the core VDD voltage once the interAptiv core has asserted SI\_SLEEP. This additional step can greatly reduce leakage and consequently power consumption during sleep mode. The minimum VDD voltage that can be used, and still allow the interAptiv core to retain state, is process dependent.

The reduction of VDD can only be controlled by external means. The interAptiv core does not provide a mechanism to reduce VDD internally during sleep mode. Note that if this option is implemented, it will take longer to restart the processor since the VDD must be ramped up to appropriate level before asserting the wake interrupt.

Refer to [Section 7.7 “interAptiv Core Power Gating”](#) for more information.

### 7.6.1.4 Restart Latency Trade-Offs

Once the decision is made to enter sleep mode, some number of clocks are required to place the interAptiv core into sleep mode, and bring the core out of sleep mode. In most designs, once sleep mode is entered, the core must remain in sleep mode for at least 100 clock cycles. Otherwise, the trade-off in time and power savings becomes negligible.

## 7.6.2 Designs Not Implementing Top Level Clock Gating

If top level clock gating was not enabled during the core configuration process, both instruction and register controlled power management can still be used. The main difference is the level of involvement of the interAptiv core in either of these processes.

From an instruction standpoint, the WAIT instruction and SI\_SLEEP signal can still be used to place the interAptiv core into sleep mode. However, since top level clock gating is disabled, it is incumbent upon external logic to suspend the input clock to the processor. If the input clock is suspended, it is suspended to the entire interAptiv core. As a result, the processor has no way to detect a ‘wake’ interrupt. Therefore, the assertion of SI\_RESET is the only way to restart the interAptiv core. Note that if this method is used, all data will be lost and the registers will revert back to their default values.

From a register standpoint, software can still set the RP bit in the CP0 Status register to initiate the transfer to sleep mode. The processor responds by asserting the SI\_RP bit to external logic. At this point, the interAptiv core does not control the clock behavior. It is incumbent upon external logic to provide the following functions:

- Suspend the core clock
- Reduce the core clock frequency
- Implement the interrupt detect function

## 7.6.3 Designs Implementing Fine Grain Clock Gating

Fine grain clock gating allows the interAptiv core to shut down the clocks to individual blocks of logic within the chip. When the ‘fine grain clock gating’ option is selected during build time, separate clock domains are assigned to the various register blocks within the interAptiv core. In the interAptiv core, there is one write enable that is used to write all registers at once. If fine grain clock gating is enabled, the clock can be enabled only to the register block that is being accessed. The write enable for the other blocks is still driven, but no clock is supplied to those blocks not being accessed.

The implementation of fine grain clock gating requires the logic required to implement multiple clock trees within the interAptiv core. Therefore, it works best in ASIC implementations where any number of clock domains can be assigned. It is less useful in FPGA implementations where the number of clock trees may be limited.

## 7.7 interAptiv Core Power Gating

In addition to clock gating, power gating can be used to gain additional power savings. The saving and restoring of processor state can be used when the power savings provided by clock gating alone are not enough. In clock gating, the state of the processor need not be saved externally because even though the clocks are suspended, the power is still applied to the interAptiv core, allowing the processor state to be saved internally.

In power gating, some or all of the power to the interAptiv core can be shut down. This causes all data within the corresponding power domain(s) to be lost once the voltage falls below the retention value as defined by the process vendor. As a result, careful consideration must be taken to save some or all of the processor states before the power is shut down. Some of the logic blocks that can be saved prior to suspending the processor are:

- Registers (GPR, CP0, CP1, and/or CP2)
- Caches (instruction and/or data)
- Translation Lookaside Buffer (TLB)
- Scratch Pad RAM (Instruction and/or Data)

There are two methods that can be used to implement a suspend/resume mechanism in a interAptiv core. These concepts are described in the following subsections.

- Hardware Suspend/Resume
- Software Suspend/Resume

### 7.7.1 Hardware Suspend/Resume

The hardware suspend/resume mechanism in the interAptiv core allows the state of the caches, scratch pad RAM, and TLB to be transferred to memory via hardware using the suspend/resume (BIST) sideband signals that are defined during chip configuration. This process of moving data to and from the interAptiv core is much faster than a pure software implementation. This process is covered in more detail in the interAptiv *Hardware User's Manual*.

### 7.7.2 Software Suspend/Resume

For systems that have not implemented any hardware suspend/resume mechanism as described in the previous section, a software mechanism can be used to save state and power down the interAptiv core. This section describes the tasks that should be performed during the suspend and resume processes.

#### 7.7.2.1 Overview of Suspend/Resume Process

The recommended way of implementing a system suspend/resume in software is having a function that will perform a seamless suspend/resume operation. This means that to the rest of the software it looks like the function was entered and exited like any normal function, while in reality this function self-terminates in the middle of its execution by turning off the power the core, then resumes from where it left off shortly after power is restored.

At a high level, the assembly language skeleton should look like this:

```
/* Entry point to suspend/resume function, including the function prologue. */
suspend_resume:
...
...

/* Here we start the suspend sequence */
suspend:
...
...
...

/* At the end of the suspend sequence we turn off power to the core. The suspend sequence should never reach the
power_is_off label*/
power_is_off:

/* This is the starting point of the resume sequence. We will get here shortly after a warm reset.*/
resume:
...
...
...

/* At the end of the resume sequence we have the function epilogue, which includes a return to the calling function.*/
```


...  
...  
...

```
jr    $31  
nop
```

As one can observe this function is clearly divided into two parts:

- The first part is the function entry (prologue) and the suspend sequence all the way down to the power shutdown. The suspend sequence includes the state saving and other supporting actions which are described in more details in the other sections.
- The second part is the resume sequence followed by the function exit (epilogue) and return to caller. The resume sequence includes state restoring and other actions which are described in more details in other sections.

If we look at the sequence of events on a time line it will look like this:



**Figure 7.6 Suspend/Resume Sequence Time Line**

### 7.7.3 Suspend Process

During a software suspend process, the following tasks are recommended. Each of these tasks is described in the following subsections.

- Save General Purpose Registers (GPR)
- Save some or all CP0 registers
- Flush the L1 data cache dirty lines and L2 cache dirty lines (if applicable)
- Save the return address
- Copy memory power down sequence into cache before switching memory to low-power mode (if applicable)
- Move memory to low-power mode (if applicable)
- Shut down power to the interAptiv core

The GPR and CP0 registers are moved to the memory stack prior so that they can be easily retrieved when power is restored to the interAptiv core. In this example, the registers would be moved to the stack and placed at the following memory offset addresses shown in [Figure 7.7](#).

Memory Stack	
0x74	Wired
0x70	Context
0x6C	Pagemask
0x68	Ebase
0x64	Config3
0x60	Config2
0x5C	Config1
0x58	Config0
0x54	Status
0x50	GPR31
0x4C	GPR30
0x48	GPR29
0x44	GPR28
0x40	GPR27
0x3C	GPR26
0x38	GPR23
0x34	GPR22
0x30	GPR21
0x2C	GPR20
0x28	GPR19
0x24	GPR18
0x20	GPR17
0x1C	GPR16
0x18	GPR7
0x14	GPR6
0x10	GPR5
0x0C	GPR4
0x08	GPR3
0x04	GPR2
0x00	GPR1

**Figure 7.7 GPR and CP0 Register Locations in the Memory Stack**

### 7.7.3.1 Save GPR Registers

MIPS recommends saving those GPR registers shown in the code example below. Note that the register numbers corresponding to the scratch registers are not saved. This includes GPR8 - GPR15, GPR24, and GPR25. For each GPR, a store word (*sw*) instruction is used to move the contents of the GPR register to memory.

```
sw    $1    0x00(sp)
sw    $2    0x04(sp)
sw    $3    0x08(sp)
sw    $4    0x0C(sp)
```

```

sw    $5    0x10(sp)
sw    $6    0x14(sp)
sw    $7    0x18(sp)
sw    $16   0x1C(sp)
sw    $17   0x20(sp)
sw    $18   0x24(sp)
sw    $19   0x28(sp)
sw    $20   0x2C(sp)
sw    $21   0x30(sp)
sw    $22   0x34(sp)
sw    $23   0x38(sp)
sw    $26   0x3C(sp)
sw    $27   0x40(sp)
sw    $28   0x44(sp)
sw    $29   0x48(sp)
sw    $30   0x4C(sp)
sw    $31   0x50(sp)

```

### 7.7.3.2 Save CP0 Registers

In the MIPS architecture the CP0 registers cannot be moved directly to memory. Therefore, they must first be moved to a GPR register. In this example the registers are moved to the k0 scratch pad register, then from the k0 register to memory at the location shown in the corresponding `sw` instruction. Note that the offset addresses for each `sw` instruction correspond to those shown in [Figure 7.7](#).

As shown in the code snippet below, only a partial set of CP0 registers are saved. This is only an example. In some cases additional registers may need to be saved depending on the implementation.

```

mfco  k0,    CP0_STATUS      /*Move from coprocessor 0, CP0_STATUS to k0*/
sw    k0,    0x54(sp)        /*Store word k0 to offset 0x54 in memory*/
mfco  k0,    CP0_CONFIG0     /*Move from coprocessor 0, CP0_CONFIG0 to k0*/
sw    k0,    0x58(sp)        /*Store word k0 to offset 0x58 in memory*/
mfco  k0,    CP0_CONFIG1     /*Move from coprocessor 0, CP0_CONFIG1 to k0*/
sw    k0,    0x5C(sp)        /*Store word k0 to offset 0x5C in memory*/
mfco  k0,    CP0_CONFIG2     /*Move from coprocessor 0, CP0_CONFIG2 to k0*/
sw    k0,    0x60(sp)        /*Store word k0 to offset 0x60 in memory*/
mfco  k0,    CP0_CONFIG3     /*Move from coprocessor 0, CP0_CONFIG3 to k0*/
sw    k0,    0x64(sp)        /*Store word k0 to offset 0x64 in memory*/
mfco  k0,    CP0_EBASE       /*Move from coprocessor 0, CP0_EBASE to k0*/
sw    k0,    0x68(sp)        /*Store word k0 to offset 0x68 in memory*/
mfco  k0,    CP0_PAGEMASK    /*Move from coprocessor 0, CP0_PAGEMASK to k0*/
sw    k0,    0x6C(sp)        /*Store word k0 to offset 0x6C in memory*/
mfco  k0,    CP0_CONTEXT     /*Move from coprocessor 0, CP0_CONTEXT to k0*/
sw    k0,    0x70(sp)        /*Store word k0 to offset 0x70 in memory*/
mfco  k0,    CP0_WIRED       /*Move from coprocessor 0, CP0_WIRED to k0*/
sw    k0,    0x74(sp)        /*Store word k0 to offset 0x74 in memory*/

```

### 7.7.3.3 Flush Dirty Lines in L1 Data Cache

The following routine can be used to flush the dirty lines in a 32 Kbyte, 4-way set associative data cache with a 32-byte line size in preparation for shut-down. In this routine software examines each cache line and performs an invalidate on all non-dirty lines, and a writeback-invalidate on all dirty lines. A similar routine must be applied for L2 dirty lines in systems implementing a level 2 cache.

```

#define INDEX_BASE 0x80000000 // We use KSEG0 address as the base address for cache index access
#define WAY_SIZE 0x2000 // size of one way in a 4-way set associative 32K cache (8K)
#define WAYOFFSET 13 // offset of bits which determine the cache way to access
#define ASSOC 4 // associativity (4 ways)
#define LINE_SIZE 32 // size of each cache line
#define IDX_WB_INV_DC 0x01 // code of index write-back invalidate D-cache operation

```

/\* This macro performs the same cache op on 32 consecutive lines. \*/

```

#define cache32_unroll32(base,op) \
    __asm__ __volatile__( \
        ".set push \n" \
        ".set noreorder \n" \
        ".set mips3 \n" \
        "cache %1, 0x000(%0); cache %1, 0x020(%0)\n" \
        "cache %1, 0x040(%0); cache %1, 0x060(%0)\n" \
        \
        "cache %1, 0x080(%0); cache %1, 0x0a0(%0)\n" \
        "cache %1, 0x0c0(%0); cache %1, 0x0e0(%0)\n" \
        "cache %1, 0x100(%0); cache %1, 0x120(%0)\n" \
        "cache %1, 0x140(%0); cache %1, 0x160(%0)\n" \
        "cache %1, 0x180(%0); cache %1, 0x1a0(%0)\n" \
        "cache %1, 0x1c0(%0); cache %1, 0x1e0(%0)\n" \
        "cache %1, 0x200(%0); cache %1, 0x220(%0)\n" \
        "cache %1, 0x240(%0); cache %1, 0x260(%0)\n" \
        "cache %1, 0x280(%0); cache %1, 0x2a0(%0)\n" \
        "cache %1, 0x2c0(%0); cache %1, 0x2e0(%0)\n" \
        "cache %1, 0x300(%0); cache %1, 0x320(%0)\n" \
        "cache %1, 0x340(%0); cache %1, 0x360(%0)\n" \
        "cache %1, 0x380(%0); cache %1, 0x3a0(%0)\n" \
        "cache %1, 0x3c0(%0); cache %1, 0x3e0(%0)\n" \
        \
        ".set pop \n" \
        : \
        : "r" (base), \
        "i" (op));

```

/\* This function scans a 4-way set associative 32K bytes data cache with 32-byte line size and performs an index write-back invalidate cache operation on each of the cache lines.\*/

```

static void flush_32k_4way_32byteline_dcache(void)
{
    unsigned long start = INDEX_BASE;
    unsigned long end = start + WAY_SIZE;
    unsigned long ws_inc = 1UL << WAYOFFSET;
    unsigned long ws_end = ASSOC << WAYOFFSET;
    unsigned long ws, addr;

    /* For every way (ws = the bits in the address which determine the cache way to access). */
    for (ws = 0; ws < ws_end; ws += ws_inc)
        /* In each way go from start to end address. */
        for (addr = start; addr < end; addr += LINE_SIZE * 32)
            /* Each time we perform the cache op on 32 lines. The address is a

```



```
combination of the cache line offset in side the way (addr) and the way bits (ws).*/
cache32_unroll32(addr|ws, IDX_WB_INV_DC);
```

### 7.7.3.4 Save the Resume Address

This routine takes the starting address of the resume sequence and saves it somewhere on the board, external to the interAptiv core. Later, after power up and reset, the warm boot sequence retrieves that address and jumps to it. This initiates execution of the resume process.

### 7.7.3.5 Copy Memory Power Down Sequence Into Cache

This piece of code loads the remaining instructions of the suspend sequence into the instruction cache. This is done since the memory (e.g. DRAM) is about to be put in low power mode and thus become inaccessible to the core. It is important that all instruction fetches hit in the instruction cache because if they miss the core won't be able to fetch them from memory.

```
*/
    .set noreorder
/* load the start address and end address of the remaining instructions */
    la    $8, mem_to_low_power
    la    $9, post_suspend      /*after power is removed*/
/* Now fill the cache line by line starting from the start address and incrementing the address by a line size in each
iteration until we get beyond the en address.*/
fill_icode:
    cache 0x14, 0($8)
    addiu $8, $8, 32
    bltu  $8, $9, fill_icode
    nop
mem_to_low_power:
```

### 7.7.3.6 Move Memory to Low Power Mode

/\* Here we have a sequence of instructions that will move the memory to low power mode. These instructions used to perform this function are SOC specific depending on the particular way the memory is implemented and addressed.\*/

```
...
...
...
```

/\* The following label comes after the end of the suspend sequence. We should never get here because we are supposed to loose power earlier.\*/

```
post_suspend:
```

### 7.7.3.7 Shut Down Power to the interAptiv Core

Once all of the above tasks have been performed, power to the interAptiv core can be suspended by reducing VDD to 0V. This task is performed by the SOC and is implementation-dependent.

## 7.7.4 Resume Process

During the software resume process, the following tasks are recommended. The tasks are handled in the opposite order in which they were executed during the suspend operation.

- System Wake-up
- Power-Up VDD to the interAptiv core and Assert Power-On Reset
- Warm/Cold Boot Detection
- Exit memory low-power mode
- Initialize caches and TLB
- Jump to resume address
- Restore CP0 registers
- Restore GPR registers

### 7.7.4.1 System Wake-Up

In a typical system the power management (PM) module stays active after the system enters suspend mode. This component will consume very little power but will keep monitoring external signals that may trigger the system to resume normal operation. Once a trigger is detected, the PM block will wake up various system components, one of these being the interAptiv core. Since power to the core was shut down earlier, the core must be powered up and brought to its Reset state.

### 7.7.4.2 Power-Up VDD to the interAptiv Core and Assert Power-On Reset

Once the system logic detects a resume condition, the system power management block must raise the VDD levels of the interAptiv core to their normal operating levels and allow the voltage to stabilize. Once the voltages are stabilized, assert the power-on reset pin to the interAptiv core.

### 7.7.4.3 Warm/Cold Boot Detection

When a processor core goes to its reset state it starts executing instructions from its Reset vector address. We call the initial sequence of instructions "boot" and it typically starts executing off of "boot ROM" memory. At this point the system must distinguish between two boot modes: cold boot and warm boot.

- A cold boot is typically performed when the entire system is powered up and has to initialize all of its hardware components. In this scenario there is typically no (or little) memory of the system's state prior to boot (although some systems will save configuration information in non-volatile memory). After the initial boot the operating system has to go through its own complete boot sequence which takes a relatively long time.
- A warm boot is typically performed to resume a system that was previously suspended for power saving. In this case much of the system state prior to boot is available and can be restored (for example, it was saved into a memory component which did not lose power or otherwise in non-volatile memory). The warm boot sequence is typically short as users expect instant response (from a user point of view the system is available even when it was suspended for power saving). A warm boot does not require the operating system to perform its full boot sequence. For the most part the OS will continue from where it left off.

In the case of a warm boot, the boot software sequence starts from the same place (the Reset vector address) whether it is a cold boot or warm boot condition. However, shortly thereafter it detects its mode whether it is a cold or warm boot. If the system resumes from suspend mode, the boot software will detect this and decide to perform a warm boot. The indication that the system is coming back from suspend mode may be available in the PM block or in some piece of memory. This mechanism is implementation dependent.

Once a decision is made to perform a warm boot and not a cold boot, the warm boot sequence will perform a basic initialization and then jump to the resume address in the suspend/resume function. The resume address will be available in an implementation dependent location where it was saved by the suspendsequence. Then, as discussed earlier, the function will restore some system state and return to its caller as if nothing ever happened. The caller may have no indication that the system was suspended for a while.

Examples of basic core initialization that must be carried out regardless of the boot mode are caches and TLB initialization. Many users will opt not to save and restore their cache and/or TLB states. Note that the interAptiv core caches and TLB wake-up in a random state and must be initialized before data can be written to them.

#### 7.7.4.4 Exit Memory Low-Power Mode

This is an optional system-dependent function. If the external memory devices were placed in low-power mode during the suspend process, the memory must exit its low-power mode before the instructions stored to the stack during the suspend process can be fetched by the interAptiv core.

#### 7.7.4.5 Initialize Caches and TLB

The initialize caches and TLB routines are always performed when reset is asserted to the interAptiv core. This is done to bring the caches to an initial state. This routine would be exactly the same as the one used in the boot example that accompanies the delivery of each interAptiv core. Refer to the boot example associated with the interAptiv core package.

#### 7.7.4.6 Jump to Resume Address

At this point the boot process is done with general initialization process initiated by the assertion of reset and is ready to start the actual resume sequence. It retrieves the starting address of the resume sequence that was saved earlier (as part of the suspend sequence) and jumps to it, thereby initiating execution of the resume sequence.

#### 7.7.4.7 Restore CP0 Registers

In the MIPS architecture the CP0 registers cannot be moved directly from memory. Therefore, they must first be moved to a GPR register. In this example the registers are moved to the k0 scratch pad register, then from the k0 register to memory at the location shown in the corresponding *lw* instruction. Note that the offset addresses for each *lw* instruction correspond to those shown in [Figure 7.7](#).

```
lw    k0,    0x74(sp)           /*Load word k0 from offset 0x74 in memory*/
mtco  k0,    CP0_WIRED         /*Move to coprocessor 0, CP0_WIRED from k0*/
lw    k0,    0x70(sp)           /*Load word k0 from offset 0x70 in memory*/
mtco  k0,    CP0_CONTEXT      /*Move to coprocessor 0, CP0_CONTEXT from k0*/
lw    k0,    0x6C(sp)           /*Load word k0 from offset 0x6C in memory*/
mtco  k0,    CP0_PAGEMASK     /*Move to coprocessor 0, CP0_PAGEMASK from k0*/
lw    k0,    0x68(sp)           /*Load word k0 from offset 0x68 in memory*/
mtco  k0,    CP0_EBASE        /*Move to coprocessor 0, CP0_EBASE from k0*/
lw    k0,    0x64(sp)           /*Load word k0 from offset 0x64 in memory*/
mfco  k0,    CP0_CONFIG3      /*Move to coprocessor 0, CP0_CONFIG3 from k0*/
lw    k0,    0x60(sp)           /*Load word k0 from offset 0x60 in memory*/
mtco  k0,    CP0_CONFIG2      /*Move to coprocessor 0, CP0_CONFIG2 from k0*/
lw    k0,    0x5C(sp)           /*Load word k0 from offset 0x5C in memory*/
mtco  k0,    CP0_CONFIG1      /*Move to coprocessor 0, CP0_CONFIG1 from k0*/
lw    k0,    0x58(sp)           /*Load word k0 from offset 0x58 in memory*/
mtco  k0,    CP0_CONFIG0      /*Move to coprocessor 0, CP0_CONFIG0 from k0*/
lw    k0,    0x54(sp)           /*Load word k0 from offset 0x54 in memory*/
mtco  k0,    CP0_STATUS       /*Move to coprocessor 0, CP0_STATUS from k0*/
```

### 7.7.4.8 Restore GPR Registers

MIPS recommends loading those GPR registers shown in the code example below. Note that the register numbers corresponding to the scratch pad registers are not loaded. This includes GPR8 - GPR15, GPR24, and GPR25. For each GPR, a load word (*lw*) instruction is used to move the contents of the corresponding memory location into the GPR.

```
lw    $31    0x50(sp)|
lw    $30    0x4C(sp)
lw    $29    0x48(sp)
lw    $28    0x44(sp)
lw    $27    0x40(sp)
lw    $26    0x3C(sp)
lw    $23    0x38(sp)
lw    $22    0x34(sp)
lw    $21    0x30(sp)
lw    $20    0x2C(sp)
lw    $19    0x28(sp)
lw    $18    0x24(sp)
lw    $17    0x20(sp)
lw    $16    0x1C(sp)
lw    $7     0x18(sp)
lw    $6     0x14(sp)
lw    $5     0x10(sp)
lw    $4     0x0C(sp)
lw    $3     0x08(sp)
lw    $2     0x04(sp)
lw    $1     0x00(sp)
```

## 7.8 CPC-Basic

This section describes the functionality of the CPC-basic option selected via the GUI at IP configuration time. The pinout for the CPC-basic is different from the standard CPC, but both versions instantiate all signals and each design only uses the signals that it requires. In the CPC-basic, all functions are controlled by hardware. There is no software control via the CPC-basic registers. However, a core's CPC state machine can be placed in "CPC Debug Mode" by the software writing a 1 to the *CPC\_CL\_STAT\_CONF\_REG[4]* register for that core. This allows software to have some control over the CPC-basic power states. Refer to [Section 7.8.7, "Software Control of Power-States"](#).

### 7.8.1 Core Power-up State

After CPC-basic reset, all cores start in the power-down state (D0). Each core stays in D0 until the associated *SI<core>\_PwrReq* pin is asserted. When this pin is asserted, the CPC-basic transitions the core through the power-up sequence ending in state U5 (non-coherent power-up). Once the software places a core into coherent mode, that core transitions into U6 (coherent power-up).

### 7.8.2 CM End Power State Control

The power-up sequence for the CM is similar to a core. After CPC-basic reset, the CM2 starts in the power-down state (D0). The CM2 stays in D0 until the *SI\_CM\_PwrReq* pin is asserted. Once this pin is asserted, the CPC transitions the core through the power-up sequence ending in state U5 (non-coherent power-up).

### 7.8.3 Core and CM Power Up Interlock

When the *SI<core>\_PwrReq* pin is asserted to power up core, the core and CM2 state machines work together to ensure that the CM reaches the power-up state (U5) before any core reaches the power-up state (U5). This is achieved by not allowing a core to transition from U3 (reset) to U4 (resetDly) until the CM2 is in state U5.

### 7.8.4 Core and CM Power Up Sequences

In the existing CPC, Isolate and lpreq and reset all deassert upon transition to the power-on state (U4 -> U5). In CPC-basic, the states U4ER and U4R are added so these signals do not change simultaneously. With these new states, isolate deasserts in U4ER, reset deasserts in U4R and lpreq deasserts when entering U5.

### 7.8.5 Core Power Down Sequence

The power-down of a core is initiated by the deassertion of the *SI<core>\_PwrReq* pin. However, the power-down sequence cannot proceed until both of the following conditions are met.

1. The Core is in the U5 state (software removed the core from the coherence domain).
2. All VPE's on the core have executed WAIT instructions. This is signaled by the assertion of the *SI<core>\_SleepWait* signal. The core transitions out of U5 until the *SI<core>\_SleepWait* signal is asserted.

### 7.8.6 Core and CM Power Down Interlock

The power-down of the CM2 is initiated by the deassertion of the *SI\_CM\_PwrReq* signal. However, the CM2 power-down sequence cannot proceed until all cores are in the power down state (D0). This interlock ensures all cores are powered down before CM2 is powered down.

### 7.8.7 Software Control of Power-States

In normal operation the target power state for the each core is controlled solely by the *SI<core>\_PwrReq* signal. However, a core's CPC state machine can be placed in "CPC Debug Mode" by the software writing a 1 to the *CPC\_CL\_STAT\_CONF\_REG[4]* register for that core. The target power state is controlled solely by value programmed by software into the *CPC\_CL\_CMD\_REG* register. In Debug Mode, the *SI<core>\_PwrReq* signal does not determine the power-state.

Note that the CPC-basic CM2 state machine does not support CPC Debug Mode.



## Coherency Manager

The interAptiv Global Control Registers address space (GCR) contains control/status registers for the entire interAptiv Multiprocessing System cluster (see [Section 8.3 “Global Control Block”](#)), as well as the individual interAptiv cores (see [Section 8.4 “Core-Local and Core-Other Control Blocks”](#)) in the cluster.

The GCR address space has a total size of 32 KBytes, which is divided into 8 KByte blocks as described in [Section 8.1 “Coherence Manager Address Map”](#). The location of the GCR block in the system address map is controlled by the `GCR_BASE` register.

Physically, the registers are located within the GCR block of the Coherence Manager (CM2) and are accessed by the interAptiv cores using 32-bit aligned uncached load/store instructions, or by I/O devices via the I/O Coherence Unit (IOCU), using read/write instructions.

This chapter contains the following sections:

- [Section 8.1 “Coherence Manager Address Map”](#)
- [Section 8.2 “CM2 Programming”](#)
- [Section 8.3 “Global Control Block”](#)
- [Section 8.4 “Core-Local and Core-Other Control Blocks”](#)
- [Section 8.5 “Global Debug Control Block”](#)

### 8.1 Coherence Manager Address Map

[Table 8.1](#) shows the address map of the four, 8-KB GCR blocks relative to the `GCR_BASE` as defined in the `GCR Base Register`. Each of these blocks of registers are described in the following sections.

**Table 8.1 interAptiv Control Space Address Map (Relative to GCR\_BASE[31:15])**

Address Range	Size (bytes)	Description
0x0000 - 0x1FFF	8 KB	<b>Global Control Block.</b> Contains registers pertaining to the global system functionality. All cores can access this block of registers.
0x2000 - 0x3FFF	8 KB	<b>Core-Local Control Block</b> (aliased for each interAptiv core). Contains registers pertaining to the interAptiv core issuing the request. Each core has its own copy of registers within this block.
0x4000 - 0x5FFF	8 KB	<b>Core-Other Control Block</b> (aliased for each interAptiv core). This block of addresses gives each Core a window into another cores Core-Local Control Block. Before accessing this space, the <i>Core-Other Addressing Register</i> in the Local Control Block must be set with the CORENum of the target Core.
0x6000 - 0x7FFF	8 KB	<b>Global Debug Block.</b> Contains global registers useful in debugging the interAptiv MPS.

### 8.1.1 Block Offsets Relative to the Base Address

The block offsets for each of the four blocks listed in [Table 8.1](#) above are relative to a GCR base address and can be located anywhere in physical memory. The base address is a 17-bit value that is programmed into the GCR\_BASE field of the *GCR Base* register located at offset address 0x0000 in the Global Control Block. The MIPS default location for the GCR\_BASE address is 0x1FBF\_8. To determine the physical address of each block using the MIPS default, this value would be added to the GCR block offset to derive the absolute physical address as shown in [Table 8.2](#).

**Table 8.2 Absolute Address of GCR Register Blocks Using the MIPS Default**

MIPS Default Base		GCR Block Offset		Absolute Physical Address	Size (bytes)	Description
0x1FBF_8	+	0x0000 - 0x1FFF	=	0x1FBF_8000 - 0x1FBF_9FFF	8 KB	Global Control Block.
0x1FBF_8	+	0x2000 - 0x3FFF	=	0x1FBF_A000 - 0x1FBF_BFFF	8 KB	Core-Local Control Block
0x1FBF_8	+	0x4000 - 0x5FFF	=	0x1FBF_C000 - 0x1FBF_DFFF	8 KB	Core-Other Control Block
0x1FBF_8	+	0x6000 - 0x7FFF	=	0x1FBF_E000 - 0x1FBF_FFFF	8 KB	Global Debug Block

### 8.1.2 Register Offsets Relative to the Block Offsets

In addition to the block offsets, the register offsets provided in each register description of this chapter are relative to the block offsets shown in [Table 8.2](#) above. To determine the physical address of each register, the MIPS default base address is added to the corresponding GCR block offset plus the actual register offset to derive the absolute physical address as shown in [Table 8.3](#). Note that this example shows only a few selected registers of the Global Control Block.

**Table 8.3 Absolute Address of Individual Global Control Block Registers**

MIPS Default Base		Global Register Block Offset		Global Register Offset		Absolute Physical Address	Global Control Register
0x1FBF_8	+	0x0000	+	0x0000	=	0x1FBF_8000	CM2 Configuration.
0x1FBF_8	+	0x0000	+	0x0008	=	0x1FBF_8008	GCR Base.
0x1FBF_8	+	0x0000	+	0x0010	=	0x1FBF_8010	CM2 Control.
0x1FBF_8	+	0x0000	+	0x0018	=	0x1FBF_8018	CM2 Control2.
0x1FBF_8	+	0x0000	+	0x0020	=	0x1FBF_8020	CM2 Access Privilege.
.....		.....		.....		.....	.....
0x1FBF_8	+	0x0000	+	0x0228	=	0x1FBF_8228	Attribute-Only Region 3 Mask.

The registers within the Core-Local blocks would be accessed in a similar manner as shown in [Table 8.4](#).



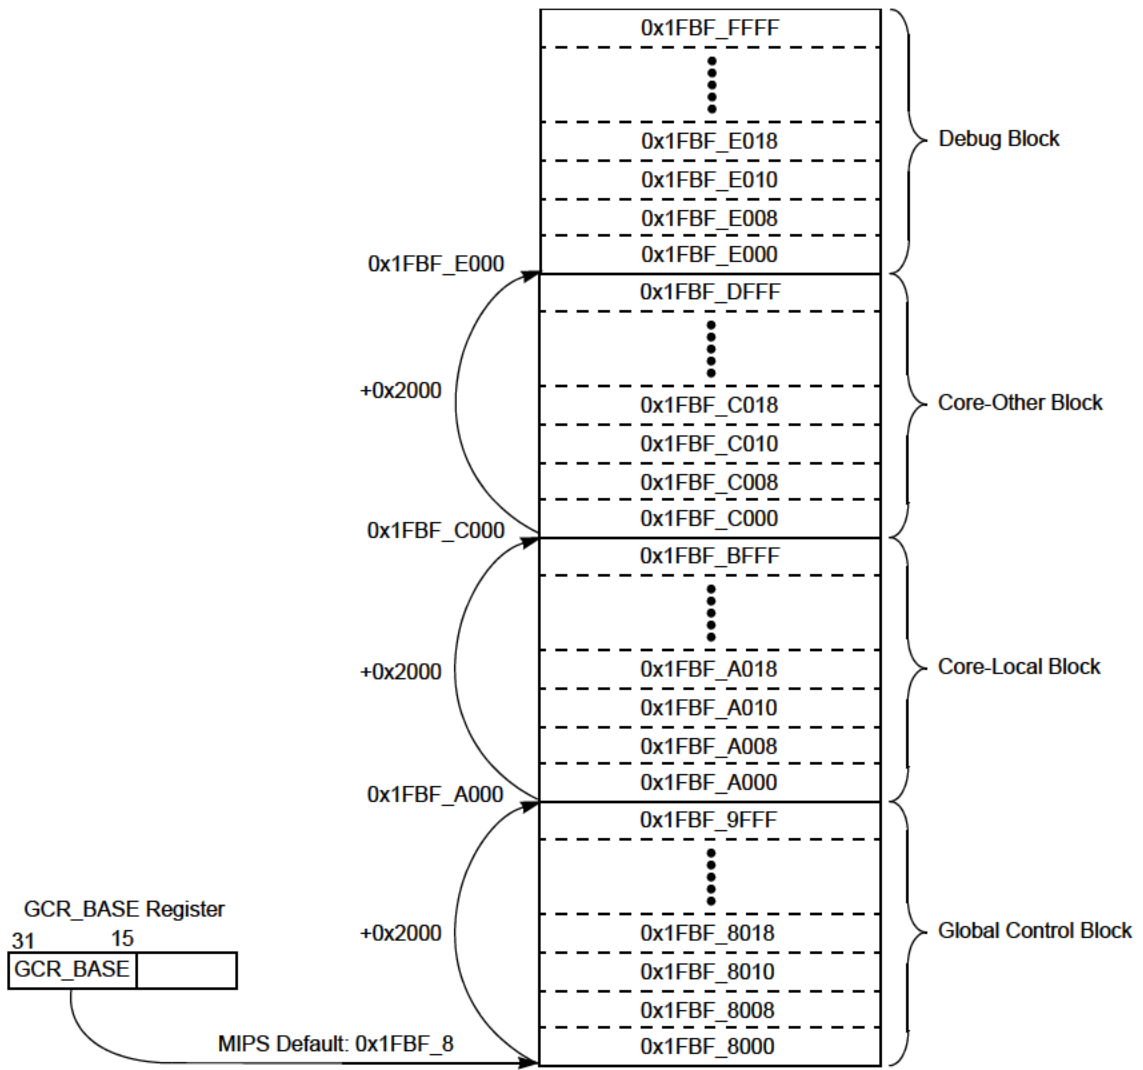
**Table 8.4 Absolute Address of Individual Core-Local Block Registers**

MIPS Default Base		Core-Local Block Offset		Core-Local Register Offset		Absolute Physical Address	Global Control Register
0x1FBB_8	+	0x2000	+	0x0000	=	0x1FBB_A000	Reserved.
0x1FBB_8	+	0x2000	+	0x0008	=	0x1FBB_A008	Core-Local Coherence Control.
0x1FBB_8	+	0x2000	+	0x0010	=	0x1FBB_A010	Core-Local Configuration.
0x1FBB_8	+	0x2000	+	0x0018	=	0x1FBB_A018	Core-Other Addressing.
0x1FBB_8	+	0x2000	+	0x0020	=	0x1FBB_A020	Core-Local Reset Exception Base.
0x1FBB_8	+	0x2000	+	0x0028	=	0x1FBB_A028	Core-Local Identification.
0x1FBB_8	+	0x2000	+	0x0030	=	0x1FBB_A030	Core-Local Reset Exception Extended Base.
0x1FBB_8	+	0x2000	+	0x0040	=	0x1FBB_A040	TCID 0 Priority.
.....	+	.....	+	.....	=	.....	.....
0x1FBB_8	+	0x2000	+	0x0080	=	0x1FBB_A080	TCID 8 Priority.

The Core-Other block would be accessed in the same manner, just with a different (Core-Other) block offset (0x4000).

This concept is described in [Figure 8.1](#) below. For simplicity, the MIPS default value is used for the GCR base address.

Figure 8.1 CM2 Register Addressing Scheme Using the MIPS Default in GCR\_BASE



## 8.2 CM2 Programming

This section provides programming examples based on the capability of the CM2 register set. Some topics described are:

- [Section 8.2.1, "Verifying Overall System Configuration"](#)
- [Section 8.2.2, "Requestor Access to GCR Registers"](#)
- [Section 8.2.3, "CM2 Interface Ports"](#)
- [Section 8.2.4, "Setting the CM2 Register Block Base Address"](#)
- [Section 8.2.5, "Address Regions"](#)
- [Section 8.2.6, "Address Map Programming Example"](#)
- [Section 8.2.7, "Core-Local GCRs"](#)
- [Section 8.2.8, "Core-Other GCRs"](#)
- [Section 8.2.9, "Accessing Another Cores CM2 GCR Registers"](#)
- [Section 8.2.10, "Boot Exception Vector Configuration"](#)
- [Section 8.2.11, "Coherency Domains"](#)
- [Section 8.2.12, "L2-Only SYNC Operation"](#)
- [Section 8.2.13, "Handling of Addresses Not Mapped to a Defined Region"](#)
- [Section 8.2.14, "Setting the Cache Coherency Attributes for Default Memory Transfers"](#)
- [Section 8.2.15, "In-Flight L1 and L2 Cache Operations"](#)
- [Section 8.2.16, "MIPS System Trace"](#)
- [Section 8.2.17, "Error Processing"](#)
- [Section 8.2.18, "Custom GCR Implementation"](#)
- [Section 8.2.19, "Attribute-Only Regions"](#)

### 8.2.1 Verifying Overall System Configuration

At build-time, the developer selects the number of cores in the system, the number of I/O coherency units (IOCU's), and the number of address regions. When the device is built, these values are hardwired into the *Global Configuration* register at offset address 0x0000. Reading this register provides the following information:

- Bits 7:0 — Number of cores in the system (1, 2, 3, or 4)
- Bits 11:8 — Number of IOCU's (0, 1, or 2)
- Bits 19:16 — Number of address regions

### 8.2.2 Requestor Access to GCR Registers

The CM2 allows up to eight requestor's in a system. A requestor can be either a core or an IOCU. The interAptiv core allows up to 6 requestors in a multiprocessing system; four cores and two IOCU's.

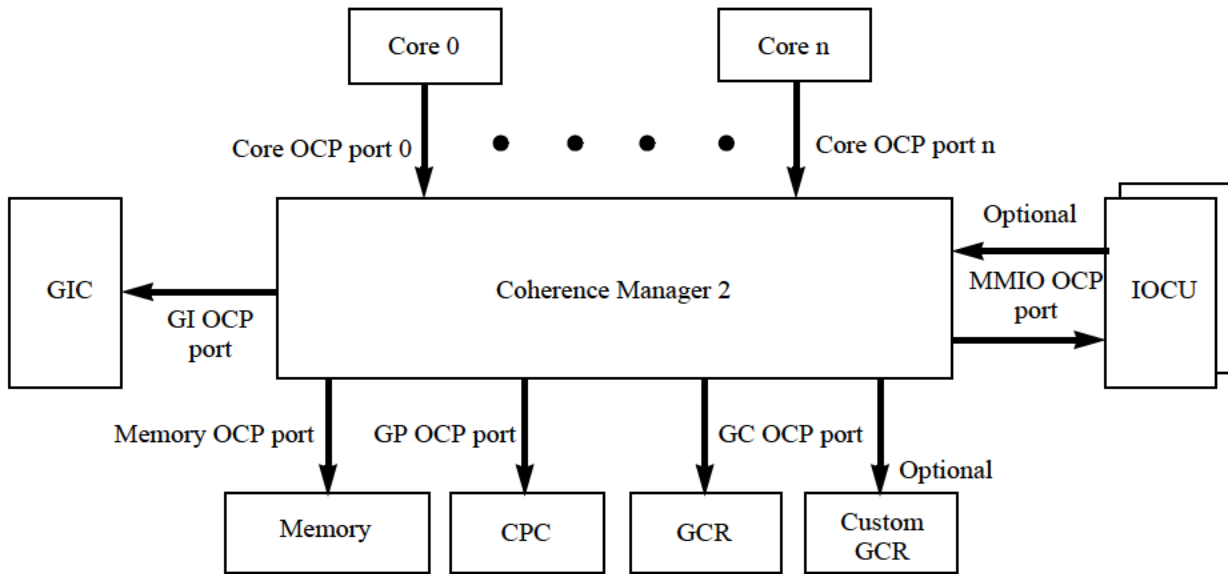
The requestor's may not have unrestricted access to the CM2 registers. During boot time, software determines which requestor's are provided access to the CM2 registers by programming the *CM2\_ACCESS\_EN* field of the *Global CSR Access Privilege* register located at offset 0x0020. Each bit in this field corresponds to a specific requestor.

The MIPS default for this field is 0xFF, meaning that all requestor's in the system have access to the CM2 register set. To disable access to the registers for a particular requestor, software need only clear the corresponding bit of this field to zero and all write requests to the CM2 registers by that requestor will be ignored.

### 8.2.3 CM2 Interface Ports

The CM2 contains numerous ports that allow the various system peripherals to communicate with the CM2. The ports connected to the CM2 are shown in Figure 8.2. The interAptiv Multiprocessing System can have up to 4 cores.

**Figure 8.2 Interface Ports of the CM2**



### 8.2.4 Setting the CM2 Register Block Base Address

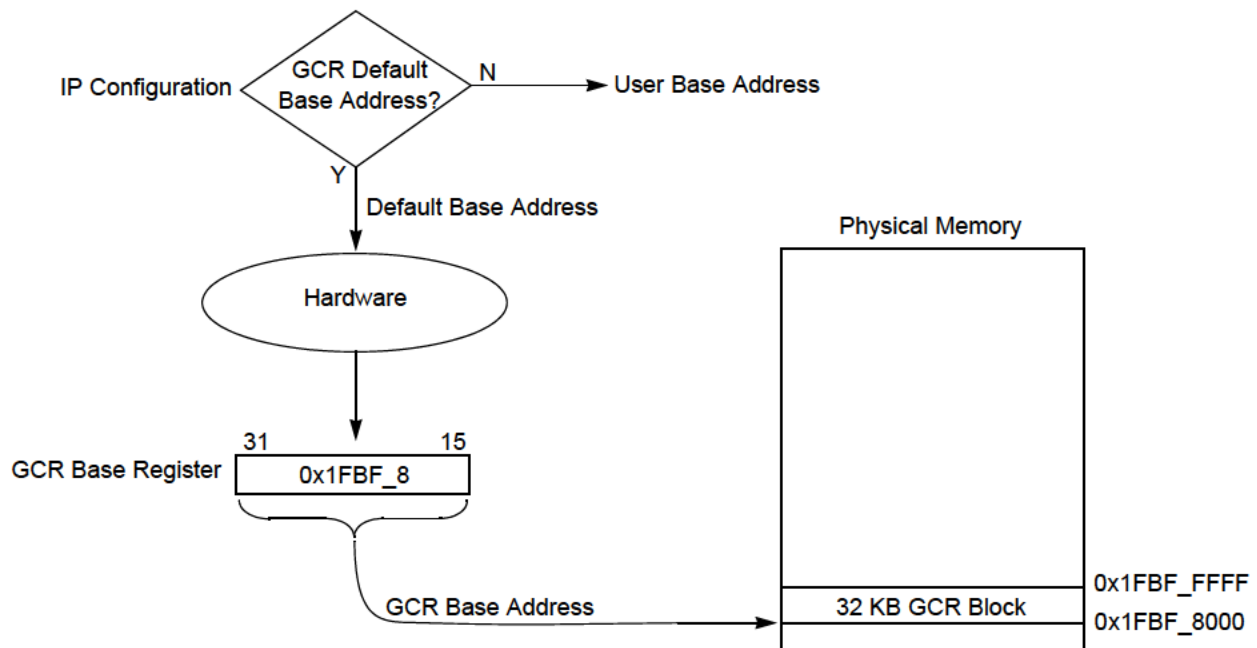
As shown in Table 8.1 above, the CM2 register map contains four contiguous 8K blocks and can be located anywhere within physical memory. During IP configuration, the user can select the option to use the MIPS default base address of 0x1FBF\_8, or they can select any 32 KB location in memory to locate the CM2 registers.

This decision determines how the 17-bit GCR\_BASE field is programmed. If the MIPS default base address option is selected, a value of 0x1FBF\_8 is loaded into this field. If the user selects their own base address, then that address is programmed into the GCR\_BASE field. Refer to Section 8.3.2.2, "GCR Base Register (GCR\_BASE Offset 0x0008)" for more information. In addition to the value in the GCR\_BASE field, the user can also select whether this field is R/W or RO during IP configuration.

The following example shows the assignment of the CM2 GCR registers in memory using the MIPS default address. Note that the physical address is shown in this diagram. During actual programming, the programmer may use the virtual address associated with a physical address of 0x1FBF\_8 to address the GCR block. The virtual address is pro-

vided prior to address translation and will be different from the resulting physical address. Refer to Chapter 3 of this manual for more information on virtual to physical address translation.

**Figure 8.3 Mapping the CM2 Registers in Physical Memory Using the MIPS Default Value**



## 8.2.5 Address Regions

The CM2 divides the address space into two types of regions:

- Fixed-size regions
- Variable-size regions

### 8.2.5.1 Fixed-Size Regions

Fixed-size regions are those that have a fixed size in memory. These include:

- GCR Base; contains the global, core-local, core-other, and debug register blocks, fixed at 32 KB.
- GIC (global interrupt controller) address space, fixed at 128 KB
- CPC (cluster power controller) address space, fixed at 32 KB
- Custom GCR address space, fixed at 64 KB

The 32 KB GCR Base region is further divided into four 8 KB blocks as described in [Table 8.1](#). Refer to [Section 8.2.4, "Setting the CM2 Register Block Base Address"](#) for more information on setting the base address in memory for the CM2 register block.

The GIC region is fixed at 128 KB. Refer to [Section 8.3.3.1, "Global Interrupt Controller Base Address Register \(GCR\\_GIC\\_BASE Offset 0x0080\)"](#) for more information on programming the base address for the GIC interface.

The CPC region is fixed at 32 KB. Refer to [Section 8.3.3.2, "Cluster Power Controller Base Address Register \(GCR\\_CPC\\_BASE Offset 0x0088\)"](#) for more information on programming the base address for the CPC interface.

The Custom GCR region is fixed at 64 KB. Refer to [Section 8.3.2.11, "GCR Custom Base Register \(GCR\\_CUSTOM\\_BASE Offset 0x0060\)"](#) for more information on programming the base address for the Custom GCR interface.

### 8.2.5.2 Variable-Size Regions

The interAptiv multiprocessing system may provide four programmable variable size address regions for mapping the IOCU's and memory. The number of regions is determined at IP configuration time. If an IOCU is not present, then the regions registers are not used. The number of regions implemented is determined as follows.

**Table 8.5 Setting the Number of Regions**

ADDR_REGIONS Field	Number of Regions	Region Assignments
0x0	0	None (typically used when there is no IOCU).
0x4	4	4 standard regions.
0x6	6	4 standard regions and 2 attribute-only regions.
0x8	8	4 standard regions and 4 attribute-only regions.

For more information, refer to the ADDR\_REGIONS field in bits 19:16 of the [Section 8.3.2.1, "Global Config Register \(GCR\\_CONFIG Offset 0x0000\)"](#). For more information on the attribute-only regions, refer to [Section 8.2.19](#).

Each region is controlled by a corresponding base and mask register as described below. These registers are used to determine not only the location and size of the memory space, but also whether this space is mapped to an IOCU or to memory. In addition, the cache coherency attributes (CCA) for each region can be defined as described in [Section 8.2.5.6, "Setting the Cache Coherency Attributes for Region Memory Transfers"](#).

In a MIPS core, mapped addresses are processed by the memory management unit (MMU) and the cache coherency attributes for a given memory page are determined. In this case, the CCA corresponds to both the L1 and L2 caches. In some situations it may be advantageous to have the CCA of the L2 different from that of the L1 cache. In this case, software can use the *CCA\_Override\_Value* field of each *Region Address Mask* register to set the CCA for the L2 cache. This changes the attributes of the cache from what was originally assigned by the core.

The CM2 provides four base address and four address mask registers for controlling variable-size address regions 0 through 3. These regions control how some transactions are routed by the CM2. The possible routing options for requests that map to these variable-size regions are:

- To/From Memory via the CM2's system memory OCP port
- To/From the IOCU's via the CM2's MMIO OCP port for Memory-Mapped I/O (in hardware I/O coherent systems only)

Refer to [Section 8.3.3.3, "CM2 Region \[0 - 3\] Base Address Register \(GCR\\_REGn\\_BASE Offsets 0x0090, 0x00A0, 0x00B0, 0x00C0\)"](#) and [Section 8.3.3.4, "CM2 Region \[0 - 3\] Address Mask Register \(GCR\\_REGn\\_MASK Offsets 0x0098, 0x00A8, 0x00B8, 0x00C8\)"](#) for more information on these registers.

### 8.2.5.3 Address Region Priorities

The priority for the region decode is as follows:

1. GCR (highest priority)
2. Custom GCR
3. CPC
4. GIC
5. Programmed MMIO regions
6. Programmed memory regions
7. *CM2\_DEFAULT\_TARGET* (lowest priority)

The above priority allows for large memory regions to be defined with small IOCU regions carved out. Note that these regions can overlap as described in [Section 8.2.5.8, "Overlapping Regions"](#).

### 8.2.5.4 Defining the Base Address Location and Size for Each Region

The address map is programmable through a set of registers located in the GCR as summarized below. Up to 8 variable-size programmable regions can be implemented. When an IOCU is present (i.e., hardware I/O Coherence is implemented), these regions determine if requests are routed to memory or to the IOCU via the CM2's MMIO port. The regions can also be used with or without an IOCU for the CCA Override feature as described in [Section 8.2.14 "Setting the Cache Coherency Attributes for Default Memory Transfers"](#).

- The *GCR Base Register* defines the address base of the GCR region. The GCR region has a fixed size of 32 KB (see [Table 8.19](#)), hence no corresponding Mask register is required. Note that this region must reside on a 32 KB boundary.
- The *Cluster Power Controller Base Address Register* defines the address base of the CPC address region. This CPC region may be disabled via the *CPC\_EN* bit in that register. When enabled, the CPC address region has a fixed size of 32 KB (see [Table 8.32](#)), hence no corresponding Mask register is required. Note that this region must reside on a 32KB boundary.
- The *Global Interrupt Controller Base Address Register* defines the address base of the GIC address region. This GIC region may be disabled via the *GIC\_EN* bit in that register. When enabled, the GIC address region has a fixed size of 128 KB (see [Table 8.31](#)), hence no corresponding Mask register is required. Note that this region must reside on a 128 KB boundary.
- The *CM2 Region [0-3] Base Address Registers* define the address base for each of the four programmable regions. The regions have a programmable base address and a programmable size that is selected via the corresponding Mask register.
- The *CM2 Region [0-3] Address Mask Registers* define the size for each of the four programmable regions. These registers work in conjunction with the corresponding *CM2 Region [0-3] Base Address Registers* to configure a given region.
- The *Custom GCR Base Register* defines the address base of the Custom GCR region. This region defines the location of registers that are implemented by the user. This region may be disabled via the *GGU\_EN* bit in the *Custom GCR Base Register*. When enabled, the Custom GCR region has a fixed size of 64 KB (see

Table 8.28), hence no corresponding Mask register is required. Note that this region must reside on a 64 KB boundary.

As described above, the base of each region is defined in the corresponding *CM2 Region [0,1,2,3] Address Base Register* (see Table 8.33), and the size of the region is defined in the corresponding *CM2 Region [0,1,2,3] Address Mask Register* (see Table 8.34). Because a base/mask scheme is used, the base must be located on a boundary of its size. A region can be sized from 64K to the entire 32-bit address space.

**Table 8.6 Setting the Base Address for the CM2 Peripheral Devices**

Block	Register Name	Offset Address	Field Name	Bits	Description
GCR	GCR_BASE	0x0008	GCR_BASE_ADDR	31:15	Sets the base address of the GCR registers. This field has a fixed size of 32 KB.
Custom GCR	GCR_CUSTOM_BASE	0x0060	CUSTOM_BASE	31:16	Sets the base address of the Customer GCR registers. This field has a fixed size of 64 KB.
GIC	GCR_GIC_BASE	0x0080	GIC_BASE_ADDR	31:17	Sets the base address of the GIC. This field has a fixed size of 128 KB.
CPC	GCR_CPC_BASE	0x0088	CPC_BASE_ADDR	31:15	Sets the base address of the CPC. This field has a fixed size of 32 KB.
Region 0	GCR_REG0_BASE	0x0090	REGION0_BASE_ADDR	31:16	Sets the base address of region 0 in memory. Minimum size is 64 KB.
	GCR_REG0_MASK	0x0098	REGION0_BASE_MASK	31:16	Sets the size of region 0 in memory.
Region 1	GCR_REG1_BASE	0x00A0	REGION1_BASE_ADDR	31:16	Sets the base address of region 1 in memory. Minimum size is 64 KB.
	GCR_REG1_MASK	0x00A8	REGION1_BASE_MASK	31:16	Sets the size of region 1 in memory.
Region 2	GCR_REG2_BASE	0x00B0	REGION2_BASE_ADDR	31:16	Sets the base address of region 2 in memory. Minimum size is 64 KB.
	GCR_REG2_MASK	0x00B8	REGION2_BASE_MASK	31:16	Sets the size of region 2 in memory.
Region 3	GCR_REG3_BASE	0x00C0	REGION3_BASE_ADDR	31:16	Sets the base address of region 3 in memory. Minimum size is 64 KB.
	GCR_REG3_MASK	0x00C8	REGION3_BASE_MASK	31:16	Sets the size of region 3 in memory.

As described above, some of the blocks are a fixed size, hence there is no corresponding Mask register. Since the GCR, GIC, and CPC blocks each contain a dedicated Base Address register, the Region 0 - 3 registers are used to access the memory and IOCU peripherals.

### 8.2.5.5 Defining the Target Device

Each *CM2 Region Address Mask* register contains a field that determines how the CM2 routes requests whose address matches the corresponding region. As defined in the *CM2\_REGION\_TARGET* field, the transaction may be routed to memory or to an I/O device via the CM2's MMIO port and IOCU. A region may be disabled by setting the *CM2\_REGION\_TARGET* in the corresponding *CM2 Region Address Mask* register to 0.

The *CM2\_DEFAULT\_TARGET* field in the *GCR Base Register* determines how to route the requests that don't match any of the defined regions. Refer to Section 8.2.13, "Handling of Addresses Not Mapped to a Defined Region" for more information.



### 8.2.5.6 Setting the Cache Coherency Attributes for Region Memory Transfers

As described in [Section 8.2.4 “Setting the CM2 Register Block Base Address”](#), the interAptiv core provides a CCA override capability that allows the CCA’s for the L2 cache to be different from those of the L1 data cache.

This capability can be achieved via the CCA override feature in the CM2 Region Address Map Registers listed in [Table 8.6](#). Software can establish up to 4 address map regions by programming the *CM2 Region Base Register 0-3* and *CM2 Region Mask Register 0-3*.

#### **Programming the CCA**

Each region has the *CCA\_Override\_Enable* and *CCA\_Override\_Value* fields which can be used to set the CCA for transactions on the system memory OCP port. If the *CCA\_Override\_Enable* field is set to 1 for a given region and the corresponding *CM2\_TARGET* field in bits 1:0 is set to memory (0x1), then transactions that map to that region and proceed to the system memory port will have a CCA value set to the corresponding *CCA\_Override\_Value* for that region. This field also determines the CCA value driven to system memory.

Any valid CCA value can be programmed into *CCA\_Override\_Value*, but because the L2 does not process coherent CCAs, a value of CWB (5) or CWBE (4) is automatically changed to WB (3) by the CM2 before being driven on the system memory OCP port. The encoding of the *CCA\_Override\_Value* field is identical to that shown in [Table 8.8](#).

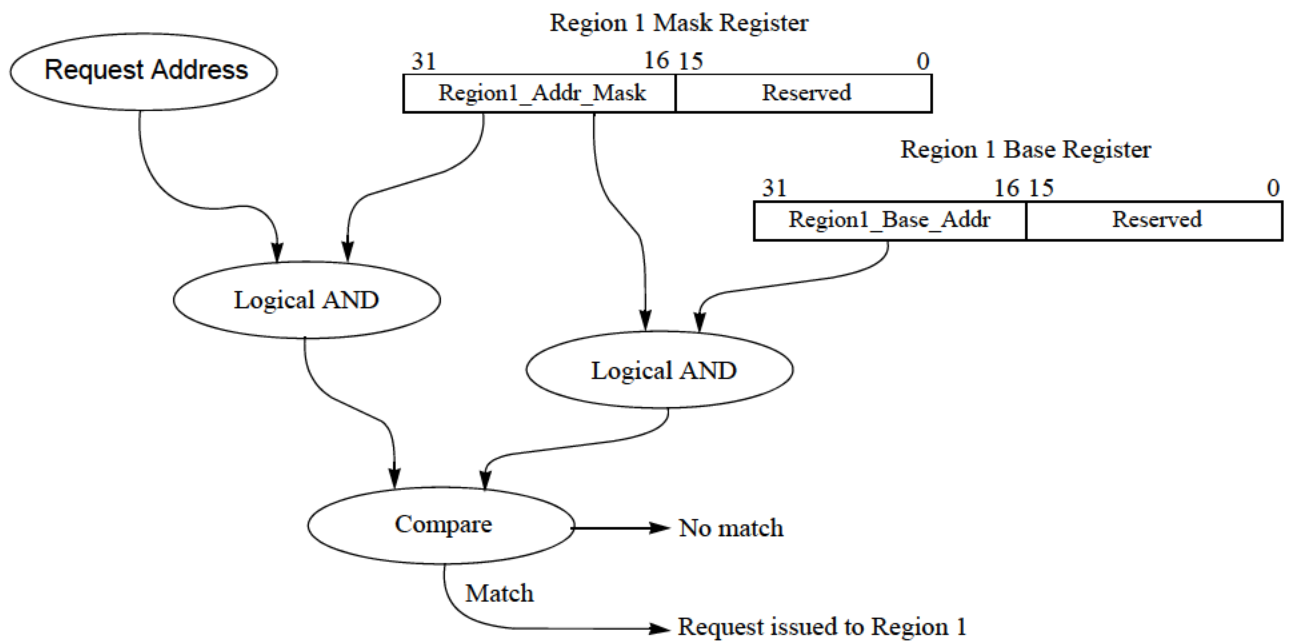
### 8.2.5.7 Issue Request Protocol and Region Masking

The CM2 contains four region mask registers used to set the size of a given region. These mask registers work in conjunction with their corresponding base address registers as shown in [Table 8.6](#). The requesting address is logically ANDed with the value in the selected *Region Address Mask* register. At the same time, the value in the corresponding *REGION\_BASE\_ADDR* field is compared to the value in the *Region Address Mask* register. If both outputs match, the request is routed to this region.

For example, if the requesting address is compared to the value in the *CM2\_REGION1\_BASE\_ADDR* and the *CM2\_REGION1\_ADDR\_MASK* registers and there is a match, then the requesting address is routed to region 1. This concept is shown in [Figure 8.4](#).

The only allowed values in this register are contiguous sets of leading 0x1’s. An 0x1 preceded by a 0x0 is not allowed (e.g., the value of 0xFFFF0 is allowed, but the value 0xFFEF is not allowed).

**Figure 8.4 Mapping a Request to Region 1**

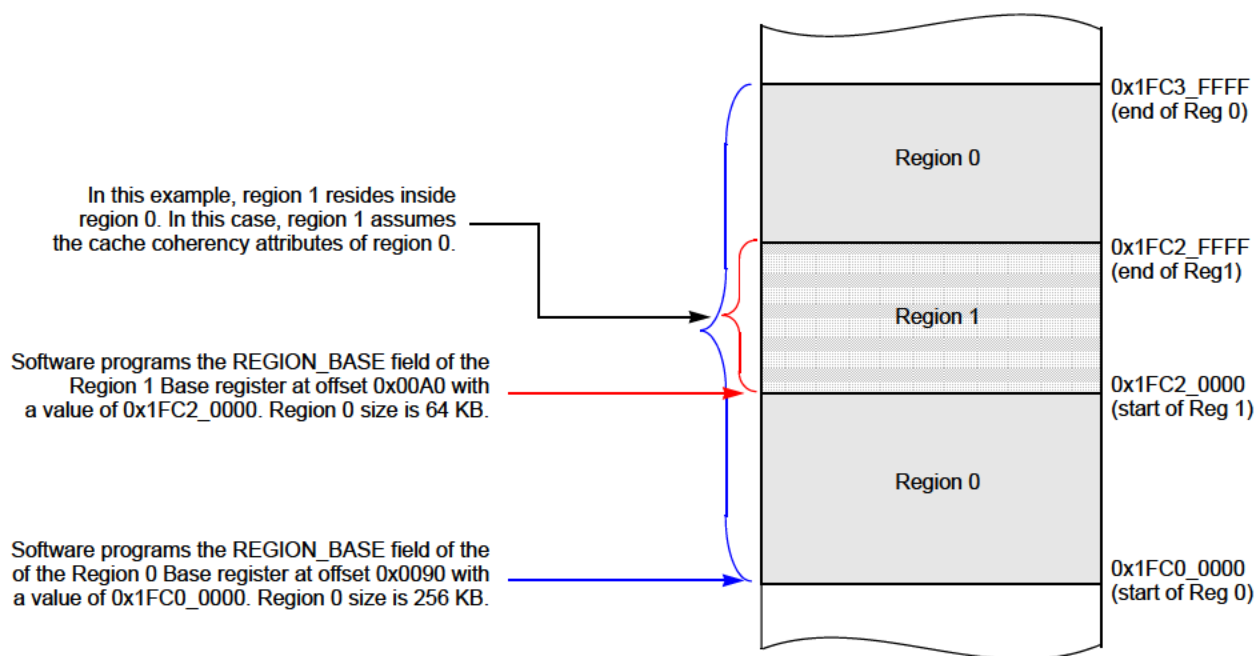


### 8.2.5.8 Overlapping Regions

Since overlapping regions are supported, it is possible that an address maps to more than one region. In this case, the CCA override enable and value are used from the lowest numbered region mapped to memory. For example, if an address matches both *CM2 Region Base/Mask Register 0* and *CM2 Region Base/Mask Register 1*, and both regions 0 and 1 are mapped to Memory (*CM2\_REGION\_TARGET* is set to 1 in both *CM2 Region Mask Register 0* and *1*), then the values of *CCA\_Override\_Enable* and *CCA\_Override\_value* in *CM2 Region Mask Register 0* will be used to determine the CCA value driven on the system memory OCP Port.

This concept is shown in [Figure 8.5](#). In this example, region 1 is a 64 KB space located inside the larger 256 KB region 0.

**Figure 8.5 Example of Overlapping Regions**

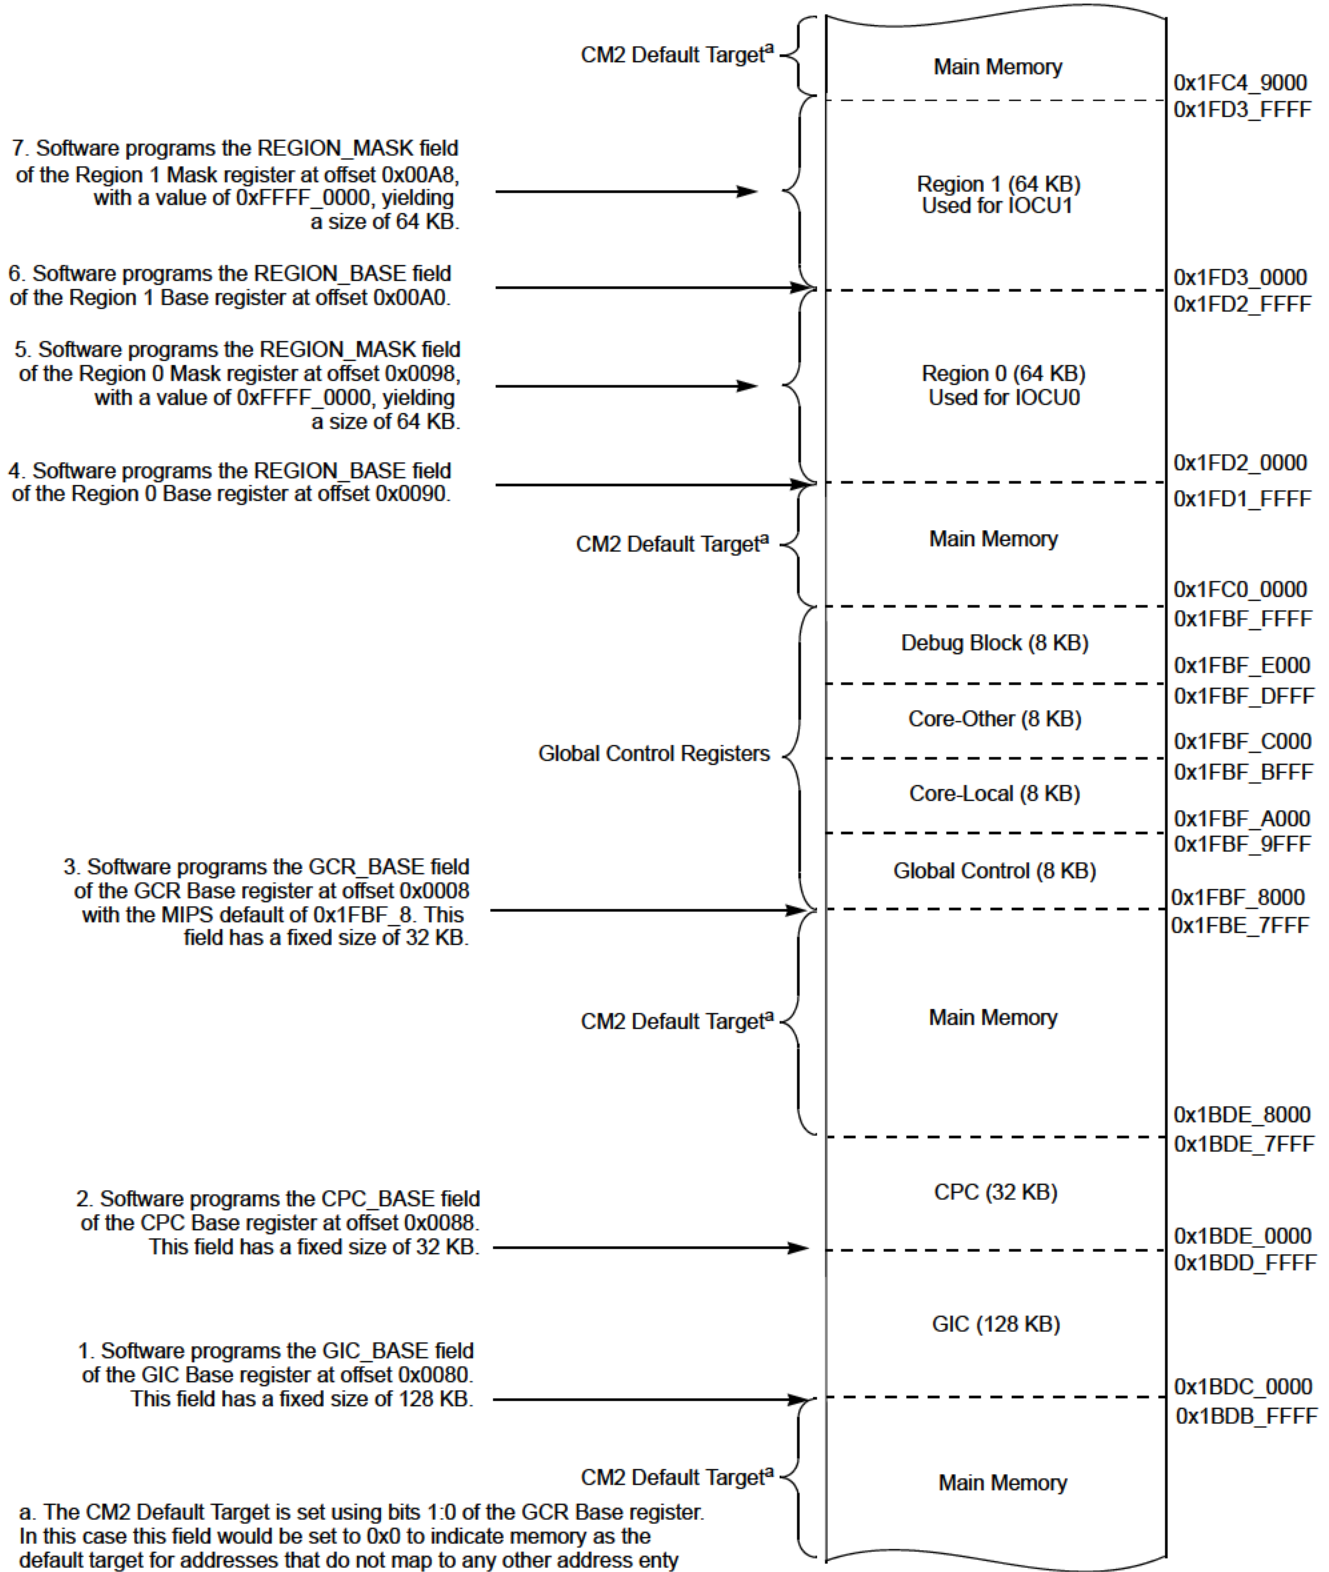


When overriding a CCA value, only the CCA driven to the system memory OCP is affected. Otherwise, the functionality of the transaction within the CM2 is based on the original CCA. When the CM2 is programmed to override the CCAs for an address region, all accesses to that region including speculative reads and writebacks (explicit or implicit) from the L1 are overridden. Transactions that are never mapped to regions, such as Legacy Syncs, CohCompletionSyncs or L2/L3 CacheOps are unaffected by the CCA override functionality.

## 8.2.6 Address Map Programming Example

This subsection provides an example of memory mapping for all of the aforementioned regions at different locations using the MIPS default base address. The memory map for this example is shown in [Figure 8.6](#).

**Figure 8.6 Address Map Programming Example**



The following programming sequence is used to configure the memory map as shown in [Figure 8.6](#) above.

1. Software programs the *GIC\_BASE* field of the *GIC Base* register located at offset 0x0080 with a value of 0x1BDC. This sets the base address of the GIC registers. This block has a fixed size of 128 KB. Refer to bits 31:17 in [Section 8.3.3.1, "Global Interrupt Controller Base Address Register \(GCR\\_GIC\\_BASE Offset 0x0080\)"](#) for more information. Note that this block must reside on a 128 KB boundary.
2. Software programs the *CPC\_BASE* field of the *CPC Base* register located at offset 0x0088 with a value of 0x1BDE\_0. This sets the base address of the CPC registers. This block has a fixed size of 32 KB. Refer to bits 31:15 in [Section 8.3.3.2, "Cluster Power Controller Base Address Register \(GCR\\_CPC\\_BASE Offset 0x0088\)"](#) for more information. Note that this block must reside on a 32 KB boundary.
3. Software programs the *GCR\_BASE* field of the *GCR Base* register located at offset 0x0008 with a value of 0x1FBF\_8. This sets the base address of the 32 KB block of GCR registers. This block is divided into four 8 KB subblocks that contain the Global, Core-Local, Core-Other, and Debug register blocks. Note that if the MIPS default address of 0x1FBF\_8 is selected for the base address of the GCR registers during IP configuration, this field becomes read-only. In this case, hardware writes the default value of 0x1FBF\_8 to this field. Refer to bits 31:15 in [Section 8.3.2.2, "GCR Base Register \(GCR\\_BASE Offset 0x0008\)"](#) for more information.
4. Software programs the *REGION\_BASE\_ADDR* field of the *CM2 Region 0 Base* register located at offset 0x0090 with a value of 0x1FD2. This sets the base address of region 0 to 0x1FD2\_0000. Refer to bits 31:16 in [Section 8.3.3.3, "CM2 Region \[0 - 3\] Base Address Register \(GCR\\_REGn\\_BASE Offsets 0x0090, 0x00A0, 0x00B0, 0x00C0\)"](#) for more information.
5. Software programs the *REGION\_ADDR\_MASK* field of the *CM2 Region 0 Address Mask* register located at offset 0x0098 with a value of 0xFFFF\_0000. This sets the size of region 0 to 64 KB. Refer to bits 31:16 in [Section 8.3.3.4, "CM2 Region \[0 - 3\] Address Mask Register \(GCR\\_REGn\\_MASK Offsets 0x0098, 0x00A8, 0x00B8, 0x00C8\)"](#) for more information. Other values for this field could be 0xFFFE (128 KB), 0xFFFC (256 KB), etc.
6. Software programs the *REGION\_BASE\_ADDR* field of the *CM2 Region 1 Base* register located at offset 0x00A0 with a value of 0x1FD3. This sets the base address of region 1 to 0x1FD3\_0000. Refer to bits 31:16 in [Section 8.3.3.3, "CM2 Region \[0 - 3\] Base Address Register \(GCR\\_REGn\\_BASE Offsets 0x0090, 0x00A0, 0x00B0, 0x00C0\)"](#) for more information.
7. Software programs the *REGION\_ADDR\_MASK* field of the *CM2 Region 1 Address Mask* register located at offset 0x00A8 with a value of 0xFFFF\_0000. This sets the size of region 1 to 64 KB. Refer to bits 31:16 in [Section 8.3.3.4, "CM2 Region \[0 - 3\] Address Mask Register \(GCR\\_REGn\\_MASK Offsets 0x0098, 0x00A8, 0x00B8, 0x00C8\)"](#) for more information. Other values for this field could be 0xFFFE (128 KB), 0xFFFC (256 KB), etc.
8. Software programs the *CM2\_DEFAULT\_TARGET* field of the *GCR Base* register with a value of 2'b00, indicating that memory is the target device for addresses that do not map to any of the address blocks shown in [Figure 8.6](#). Refer to bits 1:0 in [Section 8.3.2.2, "GCR Base Register \(GCR\\_BASE Offset 0x0008\)"](#) for more information.
9. Software programs the *CM2\_TARGET* field of the *CM2 Region 0 Address Mask* register located at offset 0x0098 with a value of 2'b10. This maps region 0 to IOCU0.
10. Software programs the *CM2\_TARGET* field of the *CM2 Region 1 Address Mask* register located at offset 0x00A8 with a value of 2'b11. This maps region 1 to IOCU1.

## 8.2.7 Core-Local GCRs

The Core-Local GCR block contains the configuration and status registers for a given core. Each core has its own copy of Core-Local registers. A core can access its own Core-Local block to determine the programmable parameters for that core. Parameters include base address assignments for cache coherency attributes, reset exception base, boot exception vector mask, etc.

## 8.2.8 Core-Other GCRs

The Core-Other GCR block is a single block that all of the cores have access to, and provides a way for one core to access the Core-Local registers of another core. Before a core can access the Core-Other space, the *Core-Other Addressing* register in that core's own Core-Local Control Block must be set with the core number (CORENUM) of the target core. In this case, a particular core would program the *Core-Other Addressing* register in its own Core-Local block with the core number to be accessed. The core would then write the contents of the register to be accessed into the Core-Other address space.

## 8.2.9 Accessing Another Cores CM2 GCR Registers

As shown in [Table 8.1](#), the CM2 provides two blocks of registers.

- Core-Local (offset range 0x2000 - 0x3FFF)
- Core-Other (offset range 0x4000 - 0x5FFF)

Each core contains a copy of these registers. The Core-Local address space contains the GCR registers for that core. The Core-Other address space allows a core to access the GCR registers for another core's Core-Local GCR block.

As described in [Section 8.2.4](#), these registers can be located anywhere in physical memory if this option is selected during IP configuration. If this option is not selected, the location of these registers are located at the MIPS default address of 0x1FBF\_8000. Refer to [Section 8.1 "Coherence Manager Address Map"](#) and related subsection for more information on use of the MIPS default memory location.

The Core-Local block represents registers corresponding to that core. If a core wishes to modify the contents of its own set of CM2 GCR registers, it writes to the Core-Local block located at the address range shown in [Table 8.1](#). If a core wishes to program the GCR registers of another core, it selects the core number and writes this value into the Core-Other Addressing register in its own Core-Local block at offset address 0x0018. The actual register in the other core to be written would use the corresponding offset in the Core-Other block shown in [Table 8.1](#).

In a multiprocessor system, it is common for one core to boot up first, then have that core boot the other cores in the system. In the following example, assume core 0 is booted up first. Then core 0 is used to program the GCR registers in core 1. This example examines how core 0 would program the boot exception vector location for core 1. Note that this example uses the MIPS default addressing scheme. The programming sequence would be as follows:

1. Core 0 writes a value of 0x0001 to the *CORENUM* field (bits 31:16) of the *Core-Other Addressing* register located in its own Core-Local block at offset 0x0018 (physical address of 0x1FBF\_A018 in [Table 8.3](#)). This indicates that the register to be programmed corresponds to core 1. Refer to [Section 8.4.1.3, "Core-Other Addressing Register"](#) for more information.
2. Core 0 writes the appropriate value into the *BEVEXCBase* field (bits 31:12) of the *Reset Exception Base* register located in the Core-Other block at offset 0x0020 (physical address of 0x1FBF\_C020 in [Table 8.4](#)). Because core 0 is setting the BEV base value for core 1, as opposed to its own core, the write is done to the Core-Other address block. Refer to [Section 8.4.1.4, "Core Local Reset Exception Base Register \(GCR\\_Cx\\_RESET\\_BASE Offset 0x0020\)"](#) for more information.



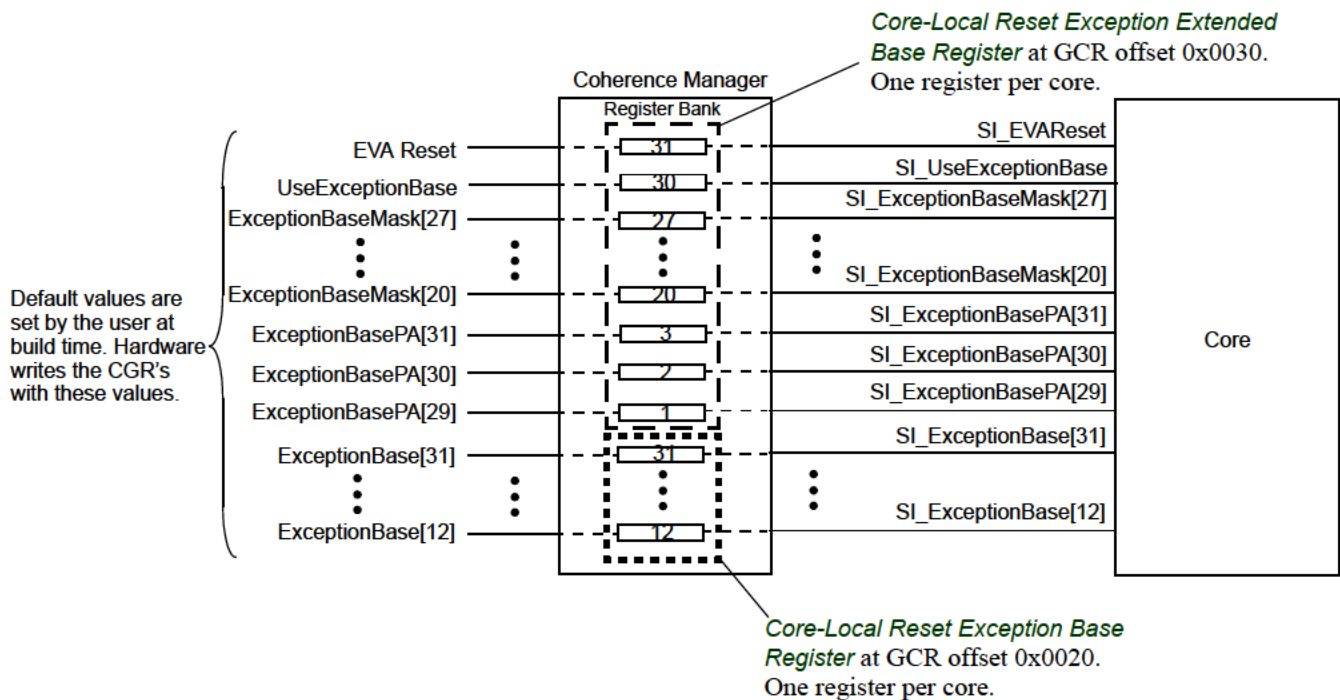


core to be programmed in a different manner and independently from one another. Refer to [Section 8.4.1.4, "Core Local Reset Exception Base Register \(GCR\\_Cx\\_RESET\\_BASE Offset 0x0020\)"](#) and [Section 8.4.1.6, "Core Local Reset Exception Extended Base Register \(GCR\\_Cx\\_RESET\\_EXT\\_BASE Offset 0x0030\)"](#) for more information on these two registers.

The CM2 drives these values to the interAptiv cores at reset. Note that the two CGR registers are loaded only on a cold boot and are programmed with the values selected by the user at build time. Each of these pins is described in [Table 8.21](#).

[Figure 8.8](#) shows the boot exception vector pins for a single interAptiv core. Each additional core would have an identical set of CM2 registers and set of BEV related pins shown in the figure. For more information, refer to [Section 3.7](#) in the MMU chapter.

**Figure 8.8 Registered Boot Exception Vector Relocation Pins — One Core**



As noted in the figure above, there is one pair of GCR registers for each core. This allows each interAptiv core to be powered up in a different memory mode and independently from one another.



The boot exception vector relocation pins are described in [Table 8.7](#).

**Table 8.7 interAptiv Boot Exception Vector Pins**

Pin Name	Field Size in Bits	CM2 GCR Register Mapping	Description
SI_EVAReset	1	Bit 31 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	<p>If this pin is asserted at reset, the interAptiv core comes up in the EVA configuration. In this case the <i>CONFIG5.K</i> bit becomes read-only with a fixed value of 1 to indicate EVA as the addressing scheme. In addition, the <i>SegCtl0 - SegCtl2</i> registers are configured with values that correspond to the EVA mapping.</p> <p>If this pin is not asserted at reset, the interAptiv core comes up in the legacy setting. In this case the <i>CONFIG5.K</i> bit becomes read-write with an initial value of 0 to indicate legacy mode. This bit is modified by software when switching from legacy mode to EVA mode.</p> <p>This pin is used in both the legacy and EVA settings. There is one <i>SI_EVAReset</i> pin per core.</p>
SI_UseExceptionBase	1	Bit 30 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	<p>In the legacy configuration, if the <i>SI_UseExceptionBase</i> pin is not asserted, then the BEV location defaults to 0xBFC0_0000.</p> <p>If the <i>SI_UseExceptionBase</i> pin is asserted, address bits <i>SI_ExceptionBase[31:30]</i> are forced to a value of 2'b10 to force the BEV location into the KSEG0/ KSEG1 space.</p> <p>This pin is only used in the legacy configuration. There is one <i>SI_UseExceptionBase</i> pin per core.</p>
SI_ExceptionBaseMask[27:20]	8	Bits 27:20 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	Used to determine the size of the boot exception vector overlay region from 1 MB to 256 MB in powers of two. These pins are used in both the legacy and EVA configurations. There is one set of <i>SI_ExceptionBaseMask</i> pins per core.
SI_ExceptionBasePA[31:29]	3	Bits 3:1 of the <i>Core-Local Reset Exception Extended Base Register</i> (offset = 0x0030)	Upper physical address bits. The size of the overlay region defined by <i>SI_ExceptionBaseMask[27:20]</i> is remapped to a location in physical address space pointed to by the <i>SI_ExceptionBasePA[31:29]</i> pins. This allows the overlay region to be placed into one of the 512 MB segments in physical memory. These pins are used in both the legacy and EVA configurations. There is one set of <i>SI_ExceptionBasePA</i> pins per core.

**Table 8.7 interAptiv Boot Exception Vector Pins (continued)**

Pin Name	Field Size in Bits	CM2 GCR Register Mapping	Description
SI_ExceptionBase[31:12]	20	Bits 31:12 of the <i>Core-Local Reset Exception Base Register</i> (offset = 0x0020)	The <i>SI_ExceptionBase[31:12]</i> pins define the boot address in virtual address space which is used to define the overlay region. These pins, along with the <i>SI_ExceptionBaseMask[27:20]</i> pins, determine the size and location of the BEV region within virtual address space. Note that the <i>CONFIG5_K</i> CP0 register bit is used to determine which pins of the <i>SI_ExceptionBase[31:12]</i> address are used to calculate the overlay. These pins are used in the EVA setting and can also be used in the legacy setting. There is one set of <i>SI_ExceptionBase</i> pins per core.

## 8.2.11 Coherency Domains

The CM2 provides the *COH\_DOMAIN\_EN* field in *Core-Local Coherence Control* register at offset 0x0008 for managing the coherency aspects of each requestor in the system. There is one register per core. A requestor can be either a core or an IOCU.

In the 8-bit *COH\_DOMAIN\_EN* field, each bit corresponds to one requestor. Setting a given bit in the *COH\_DOMAIN\_EN* field for the GCR local register corresponding to a given core puts that core into coherent mode. If the same bit in the *COH\_DOMAIN\_EN* is 0 for the GCR local register corresponding to a given core, then that core is not in coherence mode and will never issue a coherent request.

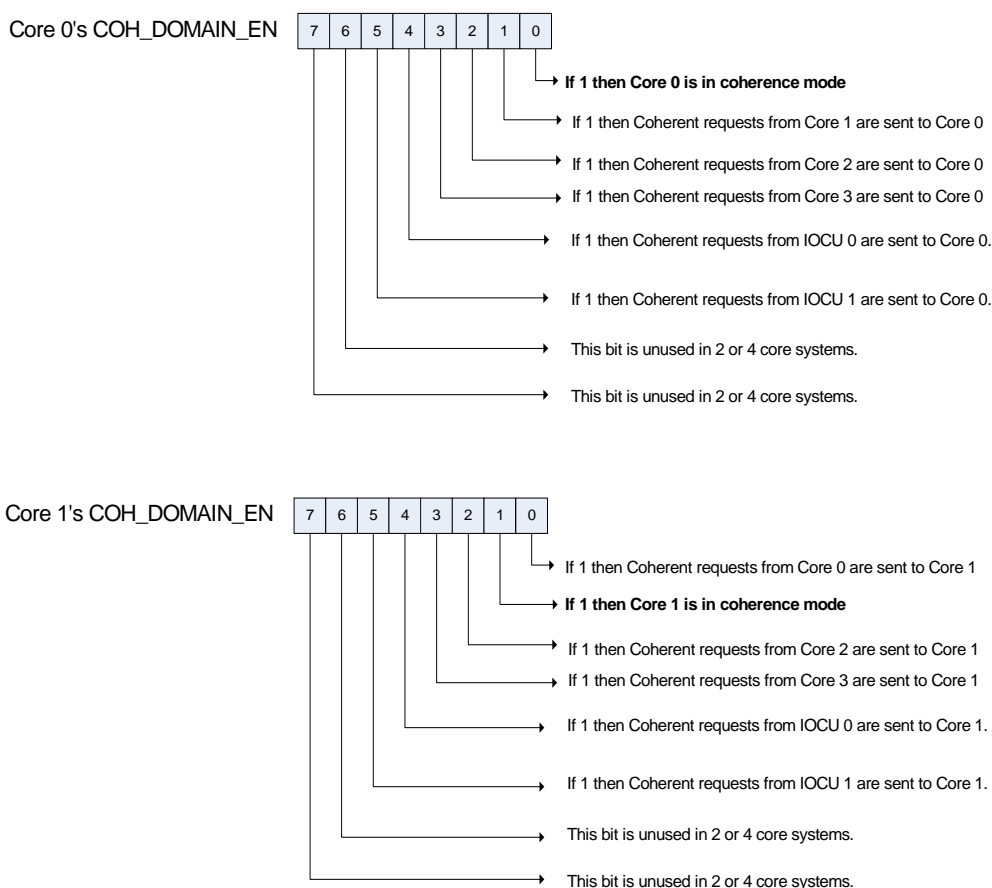
For example, if bit 1 of this field is set, then interventions from core 1 to core 0 are enabled and can occur. Note that changing the coherence mode for a local core from 0x1 to 0x0 can only be done after flushing and invalidating all the cache lines in the core; otherwise, the system behavior is UNDEFINED.

Also note that if bit 1 of the *COH\_DOMAIN\_EN* field is set for the GCR local register corresponding to core 0, then software should also set bit 0 of the *COH\_DOMAIN\_EN* field for the GCR local register corresponding to core 1.

There is no need to program *COH\_DOMAIN\_EN* for the GCR local register corresponding to IOCU.

[Section 7.1.2, "Operating Level Transitions"](#) in Chapter 7 of this manual provides examples of how this field is used to transition between coherency domains.

**Figure 8.9 Encoding of COH\_DOMAIN\_EN Field — 2 or 4 Core Package**



## 8.2.12 L2-Only SYNC Operation

In previous generation MIPS processors, the execution of a SYNC instruction would cause the entire core pipeline to stall until all read/write requests were completed. This included the L2 pipeline. After all instructions had been completed, a signal was sent to the L2 cache to continue. This caused a sometimes unnecessary stalling of the L2 cache.

The interAptiv core provides a way to perform a SYNC operation on only the L2 cache. The core defines a fixed 4 KB address space for performing L2 only SYNC operations. The base address for the location of this fixed 4 KB segment is programmed using bits 31:12 of the *L2-Only Sync Base* register located at offset 0x0070.

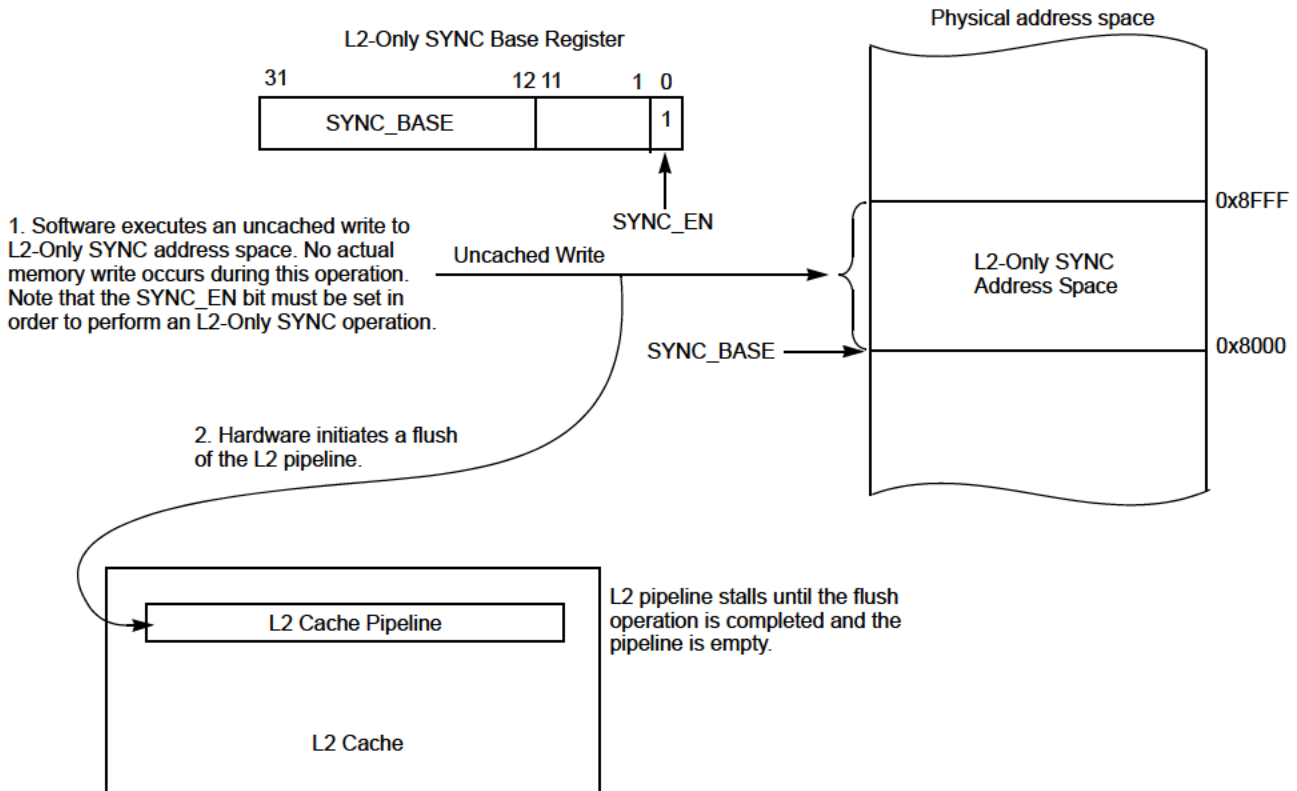
Bit 0 of the *L2-Only Sync Base* register enabled the L2-only SYNC function. If this bit is set, the CM2 treats an uncached write to anywhere within the 4 KB block as an L2-only SYNC. This operation does not write anything to memory, but rather just initiates the L2-only SYNC.

The L2-only SYNC provides a way for the software to ensure that subsequent uncached loads and stores from a core will not pass previous L2 cache operations, such as L2 cacheops.

Note that the L2-Only SYNC is not required, but it can be useful for optimizing performance. Since the L2-Only SYNC operation does not synchronize to the L1 caches, care should be taken to ensure correct system functionality.

As an example of how this operation works, assume the 4 KB block is located at offset address 0x8000 as shown in Figure 8.10.

**Figure 8.10 Example of an L2-Only SYNC Operation**



### 8.2.13 Handling of Addresses Not Mapped to a Defined Region

The CM2 handles transactions between the core and several devices as described in Figure 8.2.

For addresses that do not map to any of the defined address regions, these transactions can be mapped to either memory or one of the IOCU's as determined by the *CM2\_DEFAULT\_TARGET* field in bits 1:0 of the *GCR Base* register located at offset 0x0008. The default state of this field is determined by the value of the *SI\_CM\_Default\_Target[1:0]* pins at reset, but can be changed by software at any point. Refer to Section 8.3.2.2, "GCR Base Register (GCR\_BASE Offset 0x0008)" for more information on the *CM2\_TARGET* field.

Because programmable regions of the address map are disabled at reset, the value of *SI\_CM\_Default\_Target[1:0]* determines whether the initial boot code upon power-up is fetched from the L2/Memory port or the MMIO port. For systems without an IOCU, *SI\_CM\_Default\_Target[1:0]* should be set to 0 (memory) so that all non-coherent requests are routed to memory.

### 8.2.14 Setting the Cache Coherency Attributes for Default Memory Transfers

In previous generation MIPS processors, the cache coherency attributes (CCA) for the L1 and L2 caches were configured as one, and the CCA for the L2 cache could not be different from the CCA for the L1 data cache. The interAptiv

core provides a CCA override capability that allows the CCA's for the L2 cache to be different from those of the L1 data cache. For example, it may be useful to treat a line as cached in the L1, but uncached in the L2.

The default region determined by the *GCR Base Address* register described in [Section 8.2.4](#) above contains a mechanism for modifying the cache coherency attributes of the base region relative to that of the L1 cache. The attributes are programmed using the *CCA\_Override\_Enable* (bit 4) and *CCA\_Override\_Value* (bits 7:5) fields in the *CM2 GCR Base Address Register*. Addresses that do not map to any other region are mapped to the default region.

Any valid CCA value can be programmed into *CCA\_Override\_Value*, but because the L2 does not process coherent CCAs, a value of CWB (0x5) or CWBE (0x4) is automatically changed to WB (0x3) by the CM2 before being driven on the system memory OCP port.

The various coherency options are shown in [Table 8.8](#). Note that the CCA overrides shown below only affect the L2 cache and not the L1 cache.

**Table 8.8 Cache Coherency Attributes**

Encoding	Name	Descriptions
0x0	WT	Write through.
0x1	—	Reserved.
0x2	UC	Uncached.
0x3	WB	Writeback, cacheable, non-coherent.
0x4	CWBE	Coherent writeback exclusive. Since the CM2 does not process coherent CCA's, this encoding automatically maps to WB (0x3).
0x5	CWB	Coherent writeback. Since the CM2 does not process coherent CCA's, this encoding automatically maps to WB (0x3).
0x6	—	Reserved.
0x7	UCA	Uncached accelerated.

The *CCA\_Override\_Enable* (bit 4) must be set in order for the *CCA\_Override\_Value* field to have meaning.

When overriding a CCA value, the CCA used withing the L2 cache and driven to the system memory OCP interface is affected. Otherwise, the functionality of the transaction within the CM2 is based on the original CCA. Transactions that are not routed to the system memory OCP port, such as accesses to GCRs, GIC, CPC, or MMIO are also unaffected by the CCA Override.

## 8.2.15 In-Flight L1 and L2 Cache Operations

A core has the ability to issue a steady stream of cache operations and can potentially saturate the CM2 resources. To mitigate the possibility of this happening, the CM2 provides a mechanism to limit the number of successive cache transactions by a particular core. This limits a single core from issuing cache operations in rapid succession. The CM2 provides limits for both the L1 cache and the L2 cache via the *Global CM2 Control2* register located at offset address 0x0018. The default limit for successive L2 cache operations is four, meaning that a given core can execute a maximum of four cache operations (bits 19:16). For the L1 cache the limit is six cache operations (bits 3:0).

Setting a value of 0x0 in either of these fields disables this limitation. In this case the CM2 will not limit the number of successive cache operations that can be issued by a single core.

## 8.2.16 MIPS System Trace

The MIPS System trace is a new feature to the interAptiv Multiprocessing System and allows the SoC designer to place signals from their non-probe SoC logic directly into the trace funnel for PDTrace to capture. The logic and registers that controls System Trace are handled by the CM2. Refer to Chapter 8 of the *interAptiv Multiprocessing System Hardware User's Manual* for more information on MIPS System Trace.

## 8.2.17 Error Processing

The CM2 detects, reports, and handles several types of errors that may be caused by errant software or hardware soft or hard errors. [Table 8.9](#) lists the errors detected by the CM2. The first 7 errors are invalid requests to the GCR, GIC, or MMIO. There are two errors for invalid intervention responses due to inconsistent L1 cache states. And there are 3 errors due to L2 RAM parity errors.

When an error is detected, information that may be useful in debugging the error is captured in the *Global CM2 Error Cause Register* and *Global CM2 Error Address Register*. Refer to [Section 8.3.2.8, "Global CM2 Error Cause Register \(GCR\\_ERROR\\_CAUSE Offset 0x0048\)"](#) and [Section 8.3.2.9, "Global CM2 Error Address Register \(GCR\\_ERROR\\_ADDR Offset 0x0050\)"](#) for more information.

If these registers already have valid error information and a second error is detected, the error type of the second error is captured in the *CM2 Error Multiple Register*. However, an L2 ram correctable error is overwritten by a 2nd error that is not a second L2 ram correctable error. Refer to [Section 8.3.2.10, "Global CM2 Error Multiple Register \(GCR\\_ERROR\\_MULT Offset 0x0058\)"](#) for more information. Note that for the second error, only the error type is captured, not the associated error address.

When the *Global CM2 Error Cause Register* is loaded, an interrupt may be generated if the corresponding bit for that type of error is set in the *Global CM2 Error Mask Register* (see [Table 8.24](#)). If the error was generated by a request that requires a response and the corresponding *Global CM2 Error Mask Register* bit is 0, then the CM2 issues an ERROR response. However, if the corresponding *Global CM2 Error Mask Register* bit is 1, then the CM2 issues a normal response and an interrupt will be generated instead.

**Table 8.9 CM2 Error Types**

CM2_ERROR_TYPE	Error Name	Description	Action
0	-	Reserved	-
1	<i>GC_WR_ERR</i>	Non-Coherent Write of length > 1 to GCR or GIC	Drop Write Signal Interrupt if <i>CM_ERROR_MASK[1]</i> = 1
2	<i>GC_RD_ERR</i>	Non-Coherent Read of length > 1 to GCR or GIC	No GCR access Return SResp = ERROR if <i>CM_ERROR_MASK[2]</i> = 0 Signal Interrupt if <i>CM2_ERROR_MASK[2]</i> = 1
3	<i>COH_WR_ERR</i>	Coherent Writeback, Cacheop, or CohWriteInvalidate to GIC, GCR, MMIO	Intervention occurs Signal Interrupt if <i>CM_ERROR_MASK[3]</i> = 1
4	<i>COH_RD_ERR</i>	Coherent Read to GIC, GCR, MMIO	Intervention occurs After intervention, return SResp = ERROR to the original requestor if <i>CM_ERROR_MASK[4]</i> = 0 Signal Interrupt if <i>CM_ERROR_MASK[4]</i> = 1

**Table 8.9 CM2 Error Types (continued)**

CM2_ERROR_TYPE	Error Name	Description	Action
5	<i>MMIO_WR_ERR</i>	Write to MMIO from the IOCU (only occurs if <i>CM_DISABLE_MMIO_LIMIT</i> = 0)	Drop Write Signal Interrupt if <i>CM_ERROR_MASK[5]</i> = 1
6	<i>MMIO_RD_ERR</i>	Write to MMIO from the IOCU (only occurs if <i>CM_DISABLE_MMIO_LIMIT</i> = 0)	Return SResp = ERROR if <i>CM_ERROR_MASK[6]</i> = 0 Signal Interrupt if <i>CM_ERROR_MASK[6]</i> = 1
17	<i>INTVN_WR_ERR</i>	Request does not require a response and: One core responded with M and one or more cores responded with E, or S or One core responded with E and one or more cores responded with S or Multiple cores responded with data	If multiple M or E responses then data from core with lowest port ID is used.  Signal Interrupt if <i>CM_ERROR_MASK[17]</i> = 1
18	<i>INTVN_RD_ERR</i>	Request requires a response and: One core responded with M and one or more cores responded with E, or S or One core responded with E and one or more cores responded with S or Multiple cores responded with data	If multiple M or E responses then data from core with lowest port ID is used. Return SResp = ERROR if <i>CM_ERROR_MASK[18]</i> = 0 Signal Interrupt if <i>CM_ERROR_MASK[18]</i> = 1
24	<i>L2_RD_UNCORR</i>	Request requires a response and: an uncorrectable parity/ECC error occurred during an access to an L2 RAM	Signal Interrupt if <i>CM_ERROR_MASK[24]</i> = 1
25	<i>L2_WR_UNCORR</i>	Request does not require a response and: an uncorrectable parity/ECC error occurred during an access to an L2 RAM	Signal Interrupt if <i>CM_ERROR_MASK[25]</i> = 1
26	<i>L2_CORR</i>	A correctable parity/ECC error occurred during an access to an L2 RAM	Signal Interrupt if <i>CM_ERROR_MASK[26]</i> = 1

When an error occurs, hardware updates the read-only CM2\_ERROR\_TYPE field in bits 31:27 of the Global Config register with one of the values listed in [Table 8.9](#) above. Refer to [Section 8.3.2.1 “Global Config Register \(GCR\\_CONFIG Offset 0x0000\)”](#) for more information. When this field is written, hardware also updates the 27-bit ERROR\_INFO field that provides additional information about the error. The organization of this field varies depending on the value in the CM2\_ERROR\_TYPE field.

### 8.2.17.1 Error Codes 1 - 15

If the decimal value in the CM2\_ERROR\_TYPE field is between 1 and 15, the ERROR\_INFO field in the *Global CM2 Error Cause* register is organized as shown in [Table 8.10](#).

**Table 8.10 State of ERROR\_INFO Field for Error Types 1 through 15**

Bits	Meaning
26:18	Reserved.
17:15	CCA
14:12	Target Region (0: MEM, 1:GCR, 2: GIC, 3: MMIO, 5: CPC)
11:7	OCp MCmd (see <a href="#">Table 8.11</a> )
6:3	Source TagID
2:0	Source Port

As shown in the above table, the *OCp MCmd* field in bits 11:7 is further encoded as shown in [Table 8.11](#) below.

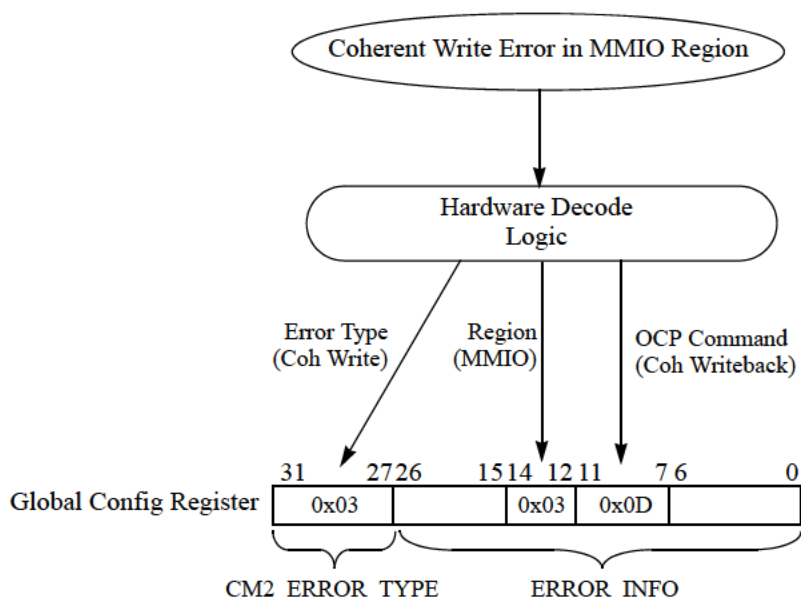
**Table 8.11 MCmd (Bits 11:7) Encoding for CM2\_ERROR\_INFO**

MCmd Encoding	Description
0x01	Legacy Write
0x02	Legacy Read
0x08	Coherent Read Own
0x09	Coherent Read Share
0x0A	Coherent Read Discard
0x0B	Coherent Ready Share Always
0x0C	Coherent Upgrade
0x0D	Coherent Writeback
0x10	Coherent Copyback
0x11	Coherent Copyback Invalidate
0x12	Coherent Invalidate
0x13	Coherent Write Invalidate
0x14	Coherent Completion Sync

Consider the example where a coherent write error occurs to the MMIO region during a coherent writeback operation. In this case, the *Global Config* register would be programmed by hardware as follows:



**Figure 8.11 Example of a Coherent Write Error to MMIO**



### 8.2.17.2 Error Codes 16 - 23

If the decimal value in the CM2\_ERROR\_TYPE field is between 16 and 23, the ERROR\_INFO field in the *Global Config* register is organized as shown in Table 8.12.

**Table 8.12 State of ERROR\_INFO Field for Error Types 16 through 23**

Bit	Meaning
26:21	Reserved
20:19	Coherent state from core 3 (see Table 8.13)
18	Intervention SResp from core 3 (see Table 8.14)
17:16	Coherent state from core 2 (see Table 8.13)
15	Intervention SResp from core 2 (see Table 8.14)
14:13	Coherent state from core 1 (see Table 8.13)
12	Intervention SResp from core 1 (see Table 8.14)
11:10	Coherent state from core 0 (see Table 8.13)
9	Intervention SResp from core 0 (see Table 8.14)
8	Request was from a Store Conditional
7:3	OCP MCmd (see Table 8.11)
2:0	Source port

Note that for each of the coherent state errors in [Table 8.12](#) (bits 20:19, 17:16, 14:13, and 11:10), the encoding for these fields is shown in [Table 8.13](#).

**Table 8.13 Coherent State Values for Error Types 16 through 23**

Encoding	Meaning
0	Invalid
1	Shared
2	Modified
3	Exclusive

For each of the Intervention SResp errors in [Table 8.12](#) (bits 18, 15, 12, and 9), the encoding for these bits is shown in [Table 8.14](#).

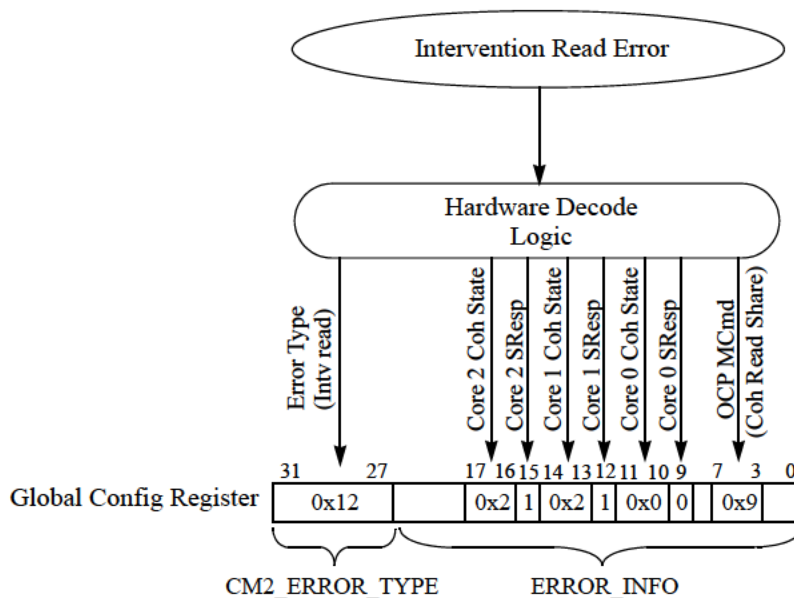
**Table 8.14 Intervention SResp Values for Error Type 16 to 23**

Encoding	Meaning
0	OK
1	Data (DVA)

Bits 7:3 of the ERROR\_INFO field are encoded the same as those shown in [Table 8.11](#).

Consider the example where a core issues a coherent read, and both cores 1 and 2 respond with modified data. In this case, the *Global Config* register would be programmed by hardware as follows:

**Figure 8.12 Example of a Intervention Read Error to MMIO**



### 8.2.17.3 Error Codes 24 - 26

If the decimal value in the *CM2\_ERROR\_TYPE* field is between 24 and 26, the *ERROR\_INFO* field in the *Global Config* register is organized as shown in [Table 8.15](#).

**Table 8.15 State of ERROR\_INFO Field for Error Types 24 to 26**

Bit	Meaning
26:24	Reserved (zero)
23	Multiple Uncorrectable
22:18	Instruction[4:0] associated with the error see <a href="#">Table 8.16</a>
17:16	Array type[1:0]: 00 = None 01 = Tag RAM single/double ECC error 10 = Data RAM single/double ECC error 11 = WS RAM uncorrectable dirty parity
15:12	DWord[3:0] with error, Array type = 2 only
11:9	Way[2:0] associated with the error
8	Multi-way error for Tag or WS RAM
7:0	Syndrome associated with Tag or WS way, or Syndrome associated with Data DWord

For each of the errors types 24 - 26 listed in [Table 8.9](#), the instruction associated with the error is encoded into bits 22:18 of the *ERROR\_INFO* field as shown in [Table 8.15](#). The encoding for these bits is shown in [Table 8.16](#) below.

**Table 8.16 Instructions for Error Type 24 to 26**

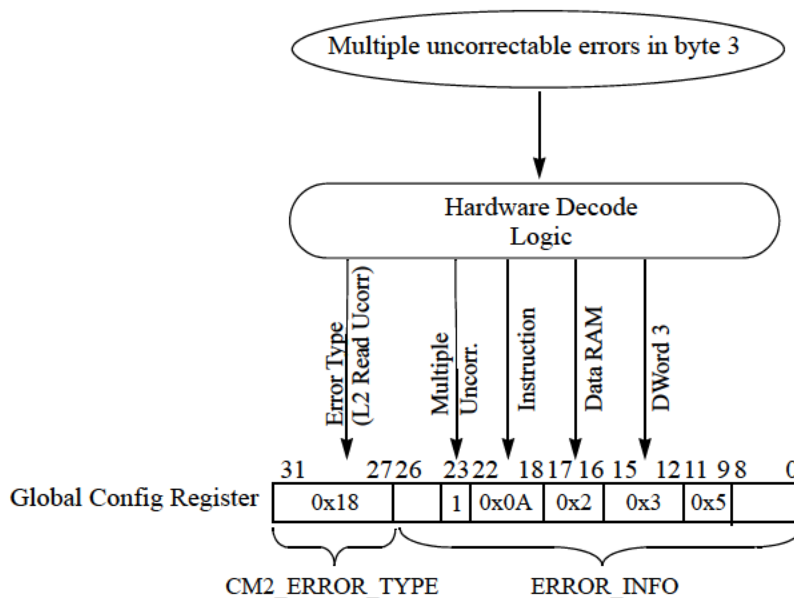
Bit	Meaning
0x00	L2_NOP
0x01	L2_ERR_CORR
0x02	L2_TAG_INV
0x03	L2_WS_CLEAN
0x04	L2_RD_MDYFY_WR
0x05	L2_WS_MRU
0x06	L2_EVICT_LN2
0x08	L2_EVICT
0x09	L2_REFL
0x0A	L2_RD
0x0B	L2_WR
0x0C	L2_EVICT_MRU
0x0D	L2_SYNC
0x0E	L2_REFL_ERR
0x10	L2_INDX_WB_INV
0x11	L2_INDX_LD_TAG

**Table 8.16 Instructions for Error Type 24 to 26 (continued)**

Bit	Meaning
0x12	L2_INDX_ST_TAG
0x13	L2_INDX_ST_DATA
0x14	L2_INDX_ST_ECC
0x18	L2_FTCH_AND_LCK
0x19	L2_HIT_INV
0x1A	L2_HIT_WB_INV
0x1B	L2_HIT_WB

Consider the example of multiple uncorrectable errors in DWord 3, way 5 of the Data RAM during an *L2 Read* instruction. In this case, the *Global Config* register would be programmed by hardware as follows:

**Figure 8.13 Multiple Uncorrectable Errors to Byte 3 of the Data RAM During an L2 Hit Writeback Instruction**



### 8.2.18 Custom GCR Implementation

The CM2 provides the ability for the user to implement a 64 KB block of custom registers that can be used to control system level functions. These registers are defined by the user and then instantiated into the design. The CM2 provides two global registers to handle the implementation of customer registers: the *Global Custom Base* register at offset 0x0060, and the *Global Custom Status* register located at offset 0x0068.

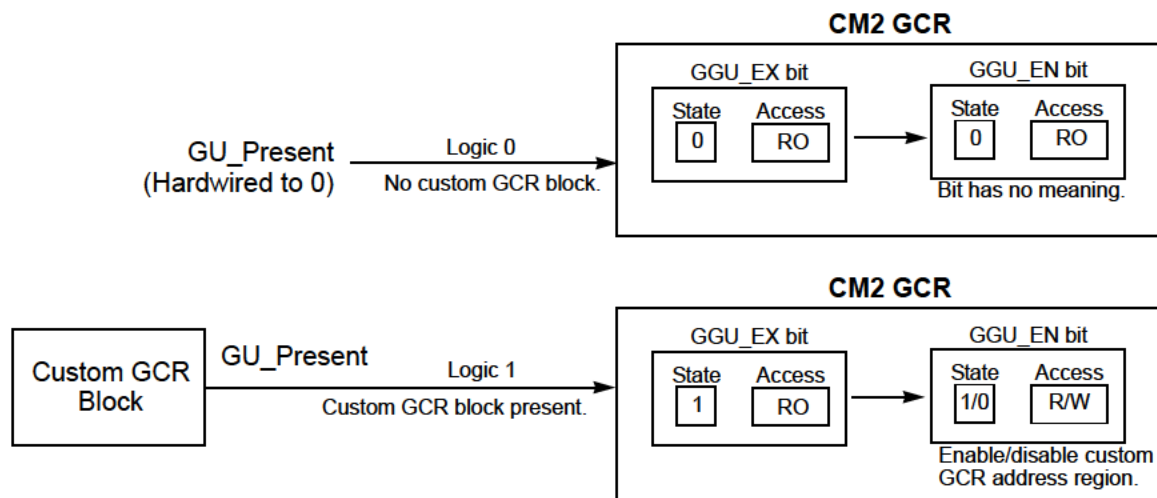
The existence of a custom GCR implementation in the system is selected during IP Configuration. If this option is selected, custom GCR hardware must drive the internal *GU\_Present* pin to the CM2. The state of this pin is loaded into the *GGU\_EX* bit in the *Global Custom Status* register. This bit indicates that a custom GCR block is connected to the CM2. Note that *GU\_Present* is an internal signal that is an output of the Custom GCR and is connected to the CM2 logic.

If a custom block is implemented, the starting address in memory of the 64 KB block is determined using the 16-bit `CUSTOM_BASE` field in the *Global Custom Base* register. Note that unlike the configuration of the CM2 Global control registers described in [Section 8.2.4](#), the `CUSTOM_BASE` field does not have a default base address and this field is undefined at reset. Therefore, it is software's responsibility to program the base address into this field during boot time if a custom GCR block is implemented.

In addition, the selected address region where the registers will reside must be enabled by setting the `GGU_EN` bit in the *Global Custom Base* register. Note that the accessibility of this bit by software depends on the state of the the `GGU_EX` bit described above. If `GGU_EX` is cleared (zero), indicating that no custom GCR is connected to the CM2, then the `GGU_EN` bit becomes RO and is not accessible by software. If this bit is set, indicating that a custom GCR is connected to the CM2, then the `GGU_EN` bit becomes R/W and is accessible by software.

This concept is described in [Figure 8.14](#) below.

**Figure 8.14 Relationship Between the `CM_Present` Signal and the `GGU_EX` and `GGU_EN` Bits at Reset**



Note that, depending on the user's implementation, the custom GCR may handle 64-bit reads/writes (unlike the normal GCR which only handles 32-bit accesses). For more information on this feature, contact MIPS Customer Support.

### 8.2.19 Attribute-Only Regions

The CM2 provides four standard variable-size regions as described in [Section 8.2.5, "Address Regions"](#), as well as four additional attribute-only regions. The attribute only regions allows the cache coherency attributes for that region to be modified, but they cannot be used to select between memory and I/O as the target.

In a situation where all of the standard variable size regions have been allocated, the attribute-only regions can be used to override the cache coherency attributes for that memory region. For example, all four attribute-only regions can be mapped to a single IOCU.

The CM2 uses four sets of base/mask registers to manage up to four attribute-only regions. The Base registers described in [Section 8.3.5.1, "CM2 Attribute-Only Region \[0 - 3\] Base Address Registers \(GCR\\_REGn\\_ATTR\\_BASE Offsets 0x0190, 0x01A0, 0x0210, 0x0220\)"](#) contain the base address in memory for each region. The Mask registers described in [Section 8.3.5.2, "CM Attribute-Only Region\[0 - 3\] Address Mask](#)

Registers (*GCR\_REGn\_ATTR\_MASK* Offsets 0x0198, 0x1A8, 0x218, 0x228)" contain the size of the region and the CCA override information.

These registers are shown starting at offset address 0x0190 in [Table 8.17](#) below:

## 8.3 Global Control Block

### 8.3.1 Global Control Block Address Map

All registers in the Global Control Block are 32 bits wide and should only be accessed using 32-bit uncached load/stores. Reads from unpopulated registers in the GCR address space return 0x0, and writes to those locations are silently dropped without generating any exceptions.

**Table 8.17 Global Control Block Register Map (Relative to Global Control Block offset)**

Register Address	Name	Type	Description
0x0000	Global Config Register ( <i>GCR_CONFIG</i> )	R	Indicates the number of Processor cores, number of interrupts, number of IOcUs, etc.
0x0008	GCR Base Register ( <i>GCR_BASE</i> )	R/W	Base of the Control Register Space
0x0010	Global CM2 Control Register ( <i>GCR_CONTROL</i> )	R/W	Control bits for the Coherence Manager
0x0018	Global CM2 Control2 Register ( <i>GCR_CONTROL2</i> )	R/W	More Control bits for the Coherence Manager
0x0020	Global CSR Access Privilege Register ( <i>GCR_ACCESS</i> )	R/W	Controls which Cores can modify the GCR Registers
0x0030	GCR Revision Register ( <i>GCR_REV</i> )	R	RevisionID of the GCR hardware
0x0040	Global CM2 Error Mask Register ( <i>GCR_ERROR_MASK</i> )	R/W	Controls what Errors are reported as Interrupts
0x0048	Global CM2 Error Cause Register ( <i>GCR_ERROR_CAUSE</i> )	R/W	Captures info when an Error occurs within the CM2
0x0050	Global CM2 Error Address Register ( <i>GCR_ERROR_ADDR</i> )	R/W	Captures address which caused the CM2 error.
0x0058	Global CM2 Error Multiple Register ( <i>GCR_ERROR_MULT</i> )	R/W	Captures information for subsequent CM2 errors.
0x0060	GCR Custom Base Register ( <i>GCR_CUSTOM_BASE</i> )	R/W	Base address of the custom user-defined 64KB control register space.
0x0068	GCR Custom Status Register ( <i>GCR_CUSTOM_STATUS</i> )	R/W	Existence and status of the custom user-defined GCR
0x0070	Global L2 only Sync Register ( <i>GCR_L2_ONLY_SYNC_BASE</i> )	R/W	Base address of the L2 only Sync 4KB address space
0x0080	Global Interrupt Controller Base Address Register ( <i>GCR_GIC_BASE</i> )	R/W	GIC Base Address

**Table 8.17 Global Control Block Register Map (Relative to Global Control Block offset)**

Register Address	Name	Type	Description
0x0088	Cluster Power Controller Base Address Register ( <i>GCR_CPC_BASE</i> )	R/W	CPC Base Address
0x0090	CM2 Region0 Base Address Register ( <i>GCR_REG0_BASE</i> )	R/W	Address Region0 Base Address This register is present only when the IOCU is present
0x0098	CM2 Region0 Address Mask Register ( <i>GCR_REG0_MASK</i> )	R/W	Address Region0 Size and Destination This register is present only when the IOCU is present
0x00A0	CM2 Region1 Base Address Register ( <i>GCR_REG1_BASE</i> )	R/W	Address Region1 Base Address This register is present only when the IOCU is present
0x00A8	CM2 Region1 Address Mask Register ( <i>GCR_REG1_MASK</i> )	R/W	Address Region1 Size and Destination This register is present only when the IOCU is present
0x00B0	CM2 Region2 Base Address Register ( <i>GCR_REG2_BASE</i> )	R/W	Address Region2 Base Address This register is present only when the IOCU is present
0x00B8	CM2 Region2 Address Mask Register ( <i>GCR_REG2_MASK</i> )	R/W	Address Region2 Size and Destination This register is present only when the IOCU is present
0x00C0	CM2 Region3 Base Address Register ( <i>GCR_REG3_BASE</i> )	R/W	Address Region3 Base Address This register is present only when the IOCU is present
0x00C8	CM2 Region3 Address Mask Register ( <i>GCR_REG3_MASK</i> )	R/W	Address Region3 Size and Destination This register is present only when the IOCU is present
0x00D0	Global Interrupt Controller Status Register ( <i>GCR_GIC_STATUS</i> )	R	Existence and status of GIC
0x00E0	Cache Revision Register ( <i>GCR_CACHE_REV</i> )	R	Revision of cache attached to the coherent Cluster.
0x00F0	Cluster Power Controller Status Register ( <i>GCR_CPC_STATUS</i> )	R	Existence and status of CPC.
0x0160	CM Arbiter Priority Register ( <i>GCR_ARB_PRI</i> )	R/W	Allows arbitration logic to give cores priority over IOCU's..
0x0190	CM Attribute-Only Region0 Base Address Register ( <i>GCR_REG0_ATTR_BASE</i> )	R/W	Attribute Only Region.
0x0198	CM Attribute-Only Region0 Address Mask Register ( <i>GCR_REG0_ATTR_MASK</i> )	R/W	Attribute Only Region.
0x01A0	CM Attribute-Only Region1 Base Address Register ( <i>GCR_REG0_ATTR_BASE</i> )	R/W	Attribute Only Region.
0x01A8	CM Attribute-Only Region1 Address Mask Register ( <i>GCR_REG1_ATTR_MASK</i> )	R/W	Attribute Only Region.
0x0200	IOCU Revision Register ( <i>GCR_IOCU1_REV</i> )	R	Revision of IOCU
0x0210	CM Attribute-Only Region2 Base Address Register ( <i>GCR_REG2_ATTR_BASE</i> )	R/W	Attribute Only Region.

**Table 8.17 Global Control Block Register Map (Relative to Global Control Block offset)**

Register Address	Name	Type	Description
0x0218	CM Attribute-Only Region2 Address Mask Register ( <i>GCR_REG2_ATTR_MASK</i> )	R/W	Attribute Only Region.
0x0220	CM Attribute-Only Region3 Base Address Register ( <i>GCR_REG3_ATTR_BASE</i> )	R/W	Attribute Only Region.
0x0228	CM Attribute-Only Region3 Address Mask Register ( <i>GCR_REG3_MASK</i> )	R/W	Attribute Only Region.
All Others	RESERVED	-	For Future Extensions

### 8.3.2 CM2 Configuration Registers

This section describes the CM2 configuration registers, including control, error and mask, revision, and custom-GCR registers.

#### 8.3.2.1 Global Config Register (GCR\_CONFIG Offset 0x0000)

This register provides information on the overall system configuration. These fields are read-only and their reset state is determined at IP configuration time. Refer to [Section 8.2.5, "Address Regions"](#) for more information on how the address regions are used.

**Figure 8.15 Global Configuration Register Format**



**Table 8.18 Global Config Register Descriptions**

Name	Bits	Description	Read/Write	Reset State										
RESERVED	31:20	Reserved, Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	-										
<i>ADDR_REGIONS</i>	19:16	Number of address regions. Total number of CM2 Address Regions. Note: only 0, 4, 6, or 8 address regions are currently supported. All other encoded values not listed below are reserved.  <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>0 Address Regions - no IOCU</td> </tr> <tr> <td>0x4</td> <td>4 Address Regions - standard</td> </tr> <tr> <td>0x6</td> <td>6 Address Regions - 4 standard + 2 Attribute Only</td> </tr> <tr> <td>0x8</td> <td>8 Address Regions - 4 standard + 4 Attribute Only</td> </tr> </tbody> </table>	Encoding	Meaning	0x0	0 Address Regions - no IOCU	0x4	4 Address Regions - standard	0x6	6 Address Regions - 4 standard + 2 Attribute Only	0x8	8 Address Regions - 4 standard + 4 Attribute Only	R	IP Configuration Value
Encoding	Meaning													
0x0	0 Address Regions - no IOCU													
0x4	4 Address Regions - standard													
0x6	6 Address Regions - 4 standard + 2 Attribute Only													
0x8	8 Address Regions - 4 standard + 4 Attribute Only													



**Table 8.18 Global Config Register Descriptions**

Name	Bits	Description	Read/ Write	Reset State										
RESERVED	15:12	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	-										
<i>NUMIOCU</i>	11:8	<p>Total number of IOCU's in the system. Note: only 0, 1, or 2 IOCU's are currently supported.</p> <table border="1" data-bbox="565 491 982 684"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>0 IOCU</td> </tr> <tr> <td>0x1</td> <td>1 IOCU's</td> </tr> <tr> <td>0x2</td> <td>2 IOCU's</td> </tr> <tr> <td>0x3 - 0xF</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	0x0	0 IOCU	0x1	1 IOCU's	0x2	2 IOCU's	0x3 - 0xF	Reserved	R	IP Configuration Value
Encoding	Meaning													
0x0	0 IOCU													
0x1	1 IOCU's													
0x2	2 IOCU's													
0x3 - 0xF	Reserved													
<i>PCORES</i>	7:0	<p>Total number of interAptiv cores in the system <i>not</i> including the IOCU's. All values not shown are reserved.</p> <table border="1" data-bbox="565 806 998 999"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x00</td> <td>1 core</td> </tr> <tr> <td>0x01</td> <td>2 cores</td> </tr> <tr> <td>0x02</td> <td>3 cores</td> </tr> <tr> <td>0x03</td> <td>4 cores</td> </tr> </tbody> </table>	Encoding	Meaning	0x00	1 core	0x01	2 cores	0x02	3 cores	0x03	4 cores	R	IP Configuration Value
Encoding	Meaning													
0x00	1 core													
0x01	2 cores													
0x02	3 cores													
0x03	4 cores													




**Table 8.19 GCR Base Register Descriptions (continued)**

Name	Bits	Description	Read/Write	Reset State										
<i>CM2_DEFAULT_TARGET</i>	1:0	<p>Determines the target device for addresses which do not match any address map entry.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Memory</td> </tr> <tr> <td>1</td> <td>Reserved</td> </tr> <tr> <td>2</td> <td>IOCU 0</td> </tr> <tr> <td>3</td> <td>IOCU 1</td> </tr> </tbody> </table> <p>Only used for hardware I/O-Coherent systems.</p>	Encoding	Meaning	0	Memory	1	Reserved	2	IOCU 0	3	IOCU 1	R/W	Value of signal <i>SI_CM_Default_Target[1:0]</i>
Encoding	Meaning													
0	Memory													
1	Reserved													
2	IOCU 0													
3	IOCU 1													

**8.3.2.3 Global CM2 Control Register (GCR\_CONTROL Offset 0x0010)**

**Figure 8.17 Global CM2 Control Register Format**



**Table 8.20 Global CM2 Control Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:17	Read as 0x0. Must be written with a value of 0x0.	-	0x0
<i>SYNCCTL</i>	16	<p>Determines SYNC behavior when a SYNC level 0x0 is executed by a core.</p> <p><i>SyncCtl</i> = 1 means Sync0 generates a memory sync</p> <p><i>SyncCtl</i> = 0 means Sync0 generates an intervention sync</p>	RW	0x0
RESERVED	15:12	Read as 0x0. Must be written with a value of 0x0.	R	0x0
<i>MMIO_SYNCDIS_1</i>	11	<p>SYNC transmit disable for IOCU 1 MMIO. Set to 1 to disable the propagation of SYNC transactions on the IOCU 1 MMIO port. This has the same effect as deasserting the external <b>SI1_CMP_IOC_SyncTxEn</b> signal.</p> <p>Clearing this bit makes the propagation of SYNC transactions on the IOCU 1 MMIO port dependent solely on the state of the external <b>SI1_CMP_IOC_SyncTxEn</b> signal.</p> <p>Refer to the pin descriptions chapter in the interAptiv Hardware User's Manual for more information on this pin. This bit only has an effect when configured with two IOCU's (<i>GCR_CONFIG.NUMIOCU</i> = 2).</p>	RW	1

**Table 8.20 Global CM2 Control Register Descriptions (continued)**

Name	Bits	Description	Read/ Write	Reset State
<i>MMIO_SYNCDIS_0</i>	10	<p>SYNC transmit disable for IOCU 0 MMIO. Set to 1 to disable the propagation of SYNC transactions on the IOCU 0 MMIO port. This has the same effect as deasserting the external <b>SI0_CMP_IOC_SyncTxEn</b> signal.</p> <p>Clearing this bit makes the propagation of SYNC transactions on the IOCU 1 MMIO port dependent solely on the state of the external <b>SI0_CMP_IOC_SyncTxEn</b> signal.</p> <p>Refer to the pin descriptions chapter in the interAptiv Hardware User's Manual for more information on this pin. This bit only has an effect when configured with no IOCU's (<i>GCR_COFIG.NUMIOCU = 0</i>)</p>	RW	1
RESERVED	9:8	Read as 0x0. Must be written with a value of 0x0.	R	0x0
UNUSED	7:6	These bits are currently unused. When writing to this register, software should assign a value of 2'b00 to this field.	R/W	0x0
<i>SYNCDIS</i>	5	<p>SYNC transmit disable. Set to 1 to disable the propagation of SYNC transactions on the system memory port. This has the same effect as deasserting <i>SI_SyncTxEn</i>. Setting to 0 makes the propagation of SYNC transactions on the system memory port dependent solely on the state of <i>SI_SyncTxEn</i>. Refer to the pin descriptions chapter in the interAptiv Hardware User's Manual for more information on this pin.</p>	RW	0x0
<i>IVU_EN</i>	4	<p>Stall until interventions are completed.</p> <p>Set to 1 to stall serialization when a core's clock is stopping or is being powered down by the CPC until all previous interventions are complete.</p> <p>Set to 0 for no stalling of serialization when a core is going offline.</p>	RW	0x0
<i>SHST_EN</i>	3	<p>Force coherent read data to shared state in L1 data cache.</p> <p>If set to 1 then Coherent Read Data is always installed in the Level 1 cache of the requesting interAptiv core in the SHARED state.</p> <p>If set to 0 then Coherent Read Data may be installed in the Level 1 cache in the SHARED state (if the data coexists in other Level 1 caches) or EXCLUSIVE (if the data does not coexist in other Level 1 caches).</p>	RW	0x0


**Table 8.20 Global CM2 Control Register Descriptions (continued)**

Name	Bits	Description	Read/ Write	Reset State
<i>PARK_EN</i>	2	<p>I/O port parking enable.</p> <p>If set to 1 and the <i>Sl&lt;iocu&gt;_CMP_IOC_ParkEn</i> signal is 1, then I/O Port Parking is enabled for the corresponding IOCU. I/O Port parking is a mechanism where the CM2 only serializes requests from the IOCU for some period of time.</p> <p>If set to 0 or <i>Sl&lt;iocu&gt;_CMP_IOC_ParkEn</i> signal is 0, then the I/O Port Parking is disabled for the corresponding IOCU.</p> <p>This bit has no effect in systems without an IOCU (i.e., they are not hardware I/O coherent). In addition, I/O Parking is automatically disabled when the CM2 arbiter is configured with Priority support (<i>GCR_ARB_PRI.ARB_PRI_PRESENT</i> = 1) and at least one port is configured as high priority (<i>GCR_ARB_PRI.ARB_PRI_MASK</i> != 0).</p>	RW	0x0
<i>MMIO_LIMIT_DIS</i>	1	<p>Limit requests to memory-mapped I/O.</p> <p>If set to 0, the CM2 avoids deadlock in systems with hardware I/O coherence by limiting requests issued to Memory-Mapped I/O. An MMIO request will be selected for serialization only if the previous request and write data (if applicable) has been accepted by the IOCU.</p> <p>If set to 1, MMIO requests are not limited and therefore deadlock may occur in systems with hardware I/O coherence unless avoided by some other mechanism.</p> <p>This bit has no effect in systems without an IOCU (i.e., they are not hardware I/O coherent) because there are no MMIO ports and therefore the limit does not apply.</p>	RW	0x0
<i>SPEC_READ_EN</i>	0	<p>Speculative coherent read enable.</p> <p>If set to 1, the CM2 may speculatively read memory for a coherent read before the intervention for that read has completed. Performance is improved by reading memory in parallel with the intervention.</p> <p>If set to 0, the CM2 will never issue speculative reads to memory.</p>	R/W	0x1

**8.3.2.4 Global CM2 Control2 Register (GCR\_CONTROL2 Offset 0x0018)**

This register sets limits on how many consecutive cache operations are allowed to the L1 and L2 caches. Refer to [Section 8.2.15, "In-Flight L1 and L2 Cache Operations"](#) for more information on how this register is used.

**Figure 8.18 Global CM2 Control2 Register Format**



**Table 8.21 Global CM2 Control2 Register**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:20	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	-	0x0
<i>L2_CACHEOP_LIMIT</i>	19:16	<p>L2 CacheOp transaction limit.</p> <p>The total number of L2 CacheOp transactions allowed by the CM2 serialization arbiter to be simultaneously in-flight. An L2 CacheOp is defined as any transaction with MAddrSpace = 0b001 or 0b010. In this context, an L2 CacheOp transaction is considered in-flight when it is selected for serialization by the CM2 until the request is issued on the CM2's system memory OCP Port.</p> <p>Setting a value of 0x0 disables the limit (i.e., the CM2 serialization arbiter will not explicitly limit the number of in-flight L12 CacheOps).</p> <p>Setting a value of 0x1 allows only a single in-flight L2 CacheOp. Setting a value of 0x2 allows two in-flight L2 CacheOps, etc...</p> <p>The purpose of this limit is to avoid the case where one or more cores substantially impact the performance of other cores by issuing a rapid succession of L2 CacheOps.</p>	R/W	0x4
RESERVED	15:4	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	-	0x0


**Table 8.21 Global CM2 Control2 Register (continued)**

Name	Bits	Description	Read/Write	Reset State
<i>L1_CACHEOP_LIMIT</i>	3:0	<p>L1 CacheOp transaction limit.</p> <p>The total number of L1 CacheOp transactions allowed by the CM2 serialization arbiter to be simultaneously in-flight. A L1 CacheOp is defined as a transaction with MAddrSpace = 0b011 or 0b1xx. In this context, a transaction is considered in-flight when it is selected for serialization by the CM2 until its intervention response is processed by the CM2 (if the cacheOp did not receive a DVA intervention response) or until all intervention data has been received (if the cacheOp received a DVA intervention response).</p> <p>Setting a value of 0x0 disables the limit (i.e., the CM2 serialization arbiter will not explicitly limit the number of in-flight L1 CacheOps). Setting a value of 0x1 allows only a single in-flight L1 CacheOp. Setting a value of 0x2 allows two in-flight L1 CacheOps, etc...</p> <p>The purpose of this limit is to avoid the case where one or more cores substantially impact the performance of other cores by issuing a rapid succession of L1 CacheOps that receive an intervention response of DVA.</p>	R/W	0x6

**8.3.2.5 Global CSR Access Privilege Register (GCR\_ACCESS Offset 0x0020)**

A request can be initiated by either a core or an IOCU. The CM2 allows for a maximum of six requestors. However, these requestors do not have unrestricted access to the CM2 register set and must be granted permission by software via this register. Refer to [Section 8.2.2, "Requestor Access to GCR Registers"](#) for more information on how this register is used.

**Figure 8.19 Global CSR Access Privilege Register Format**



**Table 8.22 Global CSR Access Privilege Register Descriptions**


Name	Bits	Description	Read/Write	Reset State
RESERVED	31:8	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0000_00

**Table 8.22 Global CSR Access Privilege Register Descriptions(continued)**

Name	Bits	Description	Read/Write	Reset State
<i>CM2_ACCESS_EN</i>	7:0	<p>Requester access to global control registers. Each bit in this field represents a coherent requester.</p> <p>If the bit is set, that requester is able to write to the GCR registers (this includes all registers within the Global, Core-Local, Core-Other, and Global Debug control blocks. The GIC is always writable by all requestors).</p> <p>If the bit is clear, any write request from that requester to the GCR registers (Global, Core-Local, Core-Other, or Global Debug control blocks) will be dropped.</p>	R/W	0xFF

**8.3.2.6 CM2 Revision Register (GCR\_REV Offset 0x0030)**

**Figure 8.20 GCR Revision Register Format**



**Table 8.23 GCR Revision Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:16	Reads as 0x0. Must be written with a value of 0x0.	R	0x0000
<i>MAJOR_REV</i>	15:8	<p>CM2 Major revision number.</p> <p>This field reflects the major revision of the GCR block. A major revision might reflect the changes from one product generation to another.</p> <p>This value changes based on the processor revision. Refer to the errata sheet of the interAptiv core for the exact value of this field.</p>	R	Preset
<i>MINOR_REV</i>	7:0	<p>CM2 Minor revision number.</p> <p>This field reflects the minor revision of the GCR block. A minor revision might reflect the changes from one release to another.</p> <p>This value changes based on the processor revision. Refer to the errata sheet of the interAptiv core for the exact value of this field.</p>	R	Preset



### 8.3.2.7 Global CM2 Error Mask Register (GCR\_ERROR\_MASK Offset 0x0040)

This register is used in conjunction with the *Global CM2 Error Cause* and *Global CM2 Error Address* registers to determine the type of error and the address which caused the error. Refer to [Section 8.2.17, "Error Processing"](#) for more information on how this register is used.

**Figure 8.21 Global CM2 Error Mask Register Format**




**Table 8.24 Global CM2 Error Mask Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>CM2_ERROR_MASK</i>	31:0	<p>CM2 Error Mask field.</p> <p>Each bit in this field represents an Error Type. If the bit is set, an interrupt is generated if an error of that type is detected.</p> <p>If the bit is set, the transaction for Read-Type Errors completes with OK response to avoid double reporting of the error.</p> <p>The Error Types that can be captured are implementation-specific.</p>	R/W	0x000A_002A (write errors cause interrupts; read errors provide error response)

### 8.3.2.8 Global CM2 Error Cause Register (GCR\_ERROR\_CAUSE Offset 0x0048)

This register is used in conjunction with the *Global CM2 Error Mask* and *Global CM2 Error Address* registers to determine the type of error and the address which caused the error. Refer to [Section 8.2.17, "Error Processing"](#) for more information on how this register is used.

**Figure 8.22 Global CM2 Error Cause Register Format**



**Table 8.25 Global CM2 Error Cause Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>CM2_ERROR_TYPE</i>	31:27	<p>Indicates type of error detected.</p> <p>When <i>CM2_ERROR_TYPE</i> is zero, no errors have been detected. When <i>CM2_ERROR_TYPE</i> is non-zero, another error will not be reloaded until a power-on reset or this field is written to 0.</p>	R/W	0


**Table 8.25 Global CM2 Error Cause Register Descriptions(continued)**

Name	Bits	Description	Read/Write	Reset State
<i>ERROR_INFO</i>	26:0	Information about the error. If <i>CM2_ERROR_TYPE</i> = 1 through 15, see <a href="#">Table 8.10</a> if <i>CM2_ERROR_TYPE</i> = 16 through 23, see <a href="#">Table 8.12</a> if <i>CM2_ERROR_TYPE</i> = 24 through 26, see <a href="#">Table 8.15</a>	R/W	Undefined

### 8.3.2.9 Global CM2 Error Address Register (*GCR\_ERROR\_ADDR* Offset 0x0050)

This register is used in conjunction with the *Global CM2 Error Cause* and *Global CM2 Error Mask* registers to determine the type of error and the address which caused the error. Refer to [Section 8.2.17, "Error Processing"](#) for more information on how this register is used.

**Figure 8.23 Global CM2 Error Address Register Format**




**Table 8.26 Global CM2 Error Address Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>CM2_ERROR_ADDR</i>	31:0	Request address which caused error. Loaded when the <i>Global Error Cause Register</i> is loaded. Bits 2:0 should always be 0.	R/W	Undefined

### 8.3.2.10 Global CM2 Error Multiple Register (*GCR\_ERROR\_MULT* Offset 0x0058)

The *Global CM2 Error Cause*, *Global CM2 Error Address*, and *Global CM2 Error Mask* registers described above provide information on the type of error, and the address which caused the error. In addition to this information, the inter-Aptiv core also provides a way to determine the type of error should a secondary error occur. However, for the secondary error, only the type of error is logged, not the associated address. This register is used to log the type of secondary error. Refer to [Section 8.2.17, "Error Processing"](#) for more information on how this register is used.

**Figure 8.24 Global CM2 Error Multiple Register Format**




**Table 8.27 Global CM2 Error Multiple Register**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:5	Reads as 0x0. Must be written with a value of 0x0.	R	0x0000_000
<i>CM2_ERROR_2ND</i>	4:0	Type of second error. Loaded when the <i>Global CM2 Error Cause Register</i> has valid error information and a second error is detected.	R/W	5'b0

### 8.3.2.11 GCR Custom Base Register (GCR\_CUSTOM\_BASE Offset 0x0060)

This register allows for the implementation of custom registers that are designed by the customer and instantiated into the design at build time. Refer to [Section 8.2.18, "Custom GCR Implementation"](#) for more information on how this register is used.

**Figure 8.25 Global Custom Base Register Format**



**Table 8.28 GCR Custom Base Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>CUSTOM_BASE</i>	31:16	This field sets the base address of the 64KB GCR custom user-defined block of the interAptiv Multiprocessing System.	R/W	Undefined
RESERVED	15:1	Reads as 0x0. Must be written with a value of 0x0.	R	0x0000
<i>GGU_EN</i>	0	If this bit is set, the address region for the Custom GCR is enabled. This bit cannot be set to 1 if <i>GGU_EX</i> = 0, indicating that a custom GCR is not attached to the CM.	R/W (if <i>GGU_EX</i> = 1) R (if <i>GGU_EX</i> = 0)	0

### 8.3.2.12 GCR Custom Status Register (GCR\_CUSTOM\_STATUS Offset 0x0068)

Refer to [Section 8.2.18, "Custom GCR Implementation"](#) for more information on how this register is used.

**Figure 8.26 Global Custom Status Register Format**



**Table 8.29 GCR Custom Status Register Descriptions**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:1	Reads as 0x0. Must be written with a value of 0x0.	R	0x0
<i>GGU_EX</i>	0	If this bit is set, the Custom GCR is connected to the CM2. The state of this bit is set based on whether or not this block is implemented at build time as determined by the state of the <i>GU_Present</i> signal.  If a Custom GCR block is not present, the <i>GU_Present</i> pin is driven to 0. If there is a custom GCR block present, then the user must drive <i>GU_Present</i> = 1 inside their custom GCR module.	R	Build time option

### 8.3.2.13 L2-Only Sync Base Register (GCR\_L2\_ONLY\_SYNC\_BASE Offset 0x0070)

The interAptiv core provides a mechanism to execute a SYNC operation to only the L2 cache, without affecting the core. Refer to [Section 8.2.12, "L2-Only SYNC Operation"](#) for more information on how this register is used.

**Figure 8.27 L2-Only Sync Base Register Format**




**Table 8.30 L2-Only Sync Base Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>SYNC_BASE</i>	31:12	L2-only SYNC base address.  This field sets the base address of the 4KB GCR L2 only Sync of the interAptiv MPS.	R/W	Undefined
RESERVED	11:1	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0
<i>SYNC_EN</i>	0	L2-only SYNC enable.  If this bit is set, the CM2 treats an uncached write request as an L2 only Sync.  If set to 0, the CM2 treats the uncached write as a regular uncached request.	R/W	0x0

## 8.3.3 CM2 Region Address Map Registers

### 8.3.3.1 Global Interrupt Controller Base Address Register (GCR\_GIC\_BASE Offset 0x0080)

**Figure 8.28 Global Interrupt Controller Base Address Register Format**



**Table 8.31 Global Interrupt Controller Base Address Register Descriptions**


Name	Bits	Description	Read/Write	Reset State
<i>GIC_BASE_ADDR</i>	31:17	Global Interrupt Controller Base Address. This field sets the base address of the 128KB Global Interrupt Controller.	R/W	Undefined

**Table 8.31 Global Interrupt Controller Base Address Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
RESERVED	16:1	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
<i>GIC_EN</i>	0	Global Interrupt Controller Enable. If this bit is set, the address region for the GIC is enabled. This bit can not be set to 1 if <i>GIC_EX</i> = 0, indicating that a GIC is not attached to the CM2.	R/W (if <i>GIC_EX</i> = 1)  R (if <i>GIC_EX</i> = 0)	0

**8.3.3.2 Cluster Power Controller Base Address Register (GCR\_CPC\_BASE Offset 0x0088)**

**Figure 8.29 Cluster Power Controller Base Address Register Format**




**Table 8.32 Cluster Power Controller Base Address Register**

Name	Bits	Description	Read/Write	Reset State
<i>CPC_BASE_ADDR</i>	31:15	This field sets the base address of the 32K Cluster Power Controller.	R/W	Undefined
RESERVED	14:1	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
<i>CPC_EN</i>	0	If this bit is set, the address region for the CPC is enabled. This bit can not be set if 1 <i>CPC_EX</i> = 0, indicating that a CPC is not attached to the CM2.	R/W (if <i>CPC_EX</i> = 1)  R (if <i>CPC_EX</i> = 0)	0

**8.3.3.3 CM2 Region [0 - 3] Base Address Register (GCR\_REGn\_BASE Offsets 0x0090, 0x00A0, 0x00B0, 0x00C0)**

Some or all of these registers may be removed during IP configuration. When an IOCU is present, there may be 4 CM2 Address Mask Registers implemented. When no IOCU is present, there may be 0 or 4 CM2 Address Mask Registers. When a register is not present, it is defined as Reserved and Read-Only of 0.

**Figure 8.30 CM2 Region [0 - 3] Base Address Register Format**




**Table 8.33 CM2 Region [0 - 3] Base Address Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>CM2_REGION_BASE_ADDR</i>	31:16	CM2 region base address. This field sets the base physical address of the memory region.	R/W	Undefined
RESERVED	15:0	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0

**8.3.3.4 CM2 Region [0 - 3] Address Mask Register (GCR\_REGn\_MASK Offsets 0x0098, 0x00A8, 0x00B8, 0x00C8)**

Some or all of these registers may be removed during IP configuration. When an IOCU is present, there may be 4 CM2 Address Mask Registers implemented. When no IOCU is present, there may be 0 or 4 CM2 Address Mask Registers. When a register is not present, it is defined as Reserved and Read-Only of 0.

**Figure 8.31 CM2 Region [0-3] Address Mask Register Format**



**Table 8.34 CM2 Region [0 - 3] Address Mask Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>CM2_REGION_ADDR_MASK</i>	31:16	This field is used to set the size of the CM2 Region. This field is used along with its equivalent <i>CM2 Region Base Address Register</i> . The request address is logically ANDed with the value of this register. The value of the associated <i>Base Address Register</i> is also logically ANDed with the value of this register. If both outputs match, then the request is routed to the CM2 region. The only allowed values in this register are contiguous sets of leading 0x1's. An 0x1 preceded by a 0x0 is not allowed (e.g., the value of 0xFFF0 is allowed, but the value 0xFFEF is not allowed).	R/W	Undefined
RESERVED	15:8	Reads as 0x0. Must be written with a value of 0x0.	R	0

**Table 8.34 CM2 Region [0 - 3] Address Mask Register Descriptions (continued)**

Name	Bits	Description	Read/Write	Reset State																										
<i>CCA_Override_Value</i>	7:5	Used with <i>CCA_Override_Enable</i> to force the Cache Coherence Attribute (CCA) value for transactions on the system memory OCP. See <i>CCA_Override_Enable</i> field. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Name</th> <th>CCA</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>WT</td> <td>Write Through</td> </tr> <tr> <td>0x1</td> <td>-</td> <td>Reserved</td> </tr> <tr> <td>0x2</td> <td>UC</td> <td>Uncached</td> </tr> <tr> <td>0x3</td> <td>WB</td> <td>WriteBack cacheable, non-coherent,</td> </tr> <tr> <td>0x4</td> <td>CWBE</td> <td rowspan="2">Mapped to WB</td> </tr> <tr> <td>0x5</td> <td>CWB</td> </tr> <tr> <td>0x6</td> <td>-</td> <td>Reserved</td> </tr> <tr> <td>0x7</td> <td>UCA</td> <td>Uncached Accelerated</td> </tr> </tbody> </table>	Encoding	Name	CCA	0x0	WT	Write Through	0x1	-	Reserved	0x2	UC	Uncached	0x3	WB	WriteBack cacheable, non-coherent,	0x4	CWBE	Mapped to WB	0x5	CWB	0x6	-	Reserved	0x7	UCA	Uncached Accelerated	R/W	0
Encoding	Name	CCA																												
0x0	WT	Write Through																												
0x1	-	Reserved																												
0x2	UC	Uncached																												
0x3	WB	WriteBack cacheable, non-coherent,																												
0x4	CWBE	Mapped to WB																												
0x5	CWB																													
0x6	-	Reserved																												
0x7	UCA	Uncached Accelerated																												
<i>CCA_Override_Enable</i>	4	If <i>CCA_Override_Enable</i> is set and the <i>CM2_TARGET</i> field is set to Memory (0x1), then transactions with addresses that map to this region will have a CCA value set to <i>CCA_Override_Value</i> when driven to system memory.	R/W	0																										
<i>Reserved</i>	3	Reads as 0x0. Must be written with a value of 0x0.	R	0																										
<i>DROP_L2</i>	2	Drop L2 CacheOp write. If this bit is set, the CM2 drops the L2 CacheOp write after it has been serialized. If this bit is cleared, the L2 CacheOp writes behave like a regular L2 CacheOp request.	R/W	0																										
<i>CM2_TARGET</i>	1:0	Maps this region to the specified device. The IOCU can only be mapped to regions 0 - 3, while memory can be mapped to all regions. . <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>Disabled</td> </tr> <tr> <td>0x1</td> <td>Memory</td> </tr> <tr> <td>0x2</td> <td>IOCU 0</td> </tr> <tr> <td>0x3</td> <td>IOCU 1</td> </tr> </tbody> </table>	Encoding	Meaning	0x0	Disabled	0x1	Memory	0x2	IOCU 0	0x3	IOCU 1	R/W	0																
Encoding	Meaning																													
0x0	Disabled																													
0x1	Memory																													
0x2	IOCU 0																													
0x3	IOCU 1																													

### 8.3.4 CM2 Status and Revision Registers

This section contains the status registers for the GIC and CPC, and the revision information for the L2 cache.

#### 8.3.4.1 Global Interrupt Controller Status Register (GCR\_GIC\_STATUS Offset 0x00D0)

**Figure 8.32 Global Interrupt Controller Status Register Format**




**Table 8.35 Global Interrupt Controller Status Register**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:1	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
<i>GIC_EX</i>	0	GIC to CM2 connection. If this bit is set, the GIC is connected to the CM2.	R	1

**8.3.4.2 Cache Revision Register (GCR\_CACHE\_REV Offset 0x00E0)**

**Figure 8.33 Cache Revision Register Format**




**Table 8.36 Cache Revision Register**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:16	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0
<i>MAJOR_REV</i>	15:8	This field reflects the major revision of the Cache block inside the CM2.	R	Preset
<i>MINOR_REV</i>	7:0	This field reflects the minor revision of the Cache block inside the CM2.	R	Preset

**8.3.4.3 Cluster Power Controller Status Register (GCR\_CPC\_STATUS Offset 0x00F0)**

**Figure 8.34 Cluster Power Controller Status Register Format**



**Table 8.37 Cluster Power Controller Status Register Descriptions**


Name	Bits	Description	Read/Write	Reset State
RESERVED	31:1	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
<i>CPC_EX</i>	0	This bit is always 1 in the interAptiv core as the CPC is always connected to the CM2.	R	1



### 8.3.4.4 CM Arbiter Priority Register (GCR\_ARB\_PRI Offset 0x0160)

The CM2 Serialization arbiter has a new feature that allows requesters (either cores or IOcUs) to be assigned as a high priority port. Agents assigned as high priority are selected ahead of ports not assigned as high priority. The assignment of high priority is provided via the new GCR\_ARB\_PRI register. The arbiter also has a time-out counter that ensures that low priority agents are not permanently starved.

**Figure 8.35 CM Arbiter Priority Register Format**




**Table 8.38 CM Arbiter Priority Register Descriptions**

Name	Bits	Description	Read/Write	Reset State														
R	31:30	Reads as 0x0. Must be written with a value of 0x0.	R	0														
<i>ARB_PRI_MASK</i>	29:24	Setting a bit of ARB_PRI_MASK field to 1 indicates that requests received on the corresponding port are considered high priority. High priority requests are serviced before low priority, unless the priority timer has expired.  <table border="0" style="margin-left: 40px;"> <tr> <td style="text-align: center;"><u>Register Bit</u></td> <td style="text-align: center;"><u>Requestor</u></td> </tr> <tr> <td>GCR_ARB_PRI[29]</td> <td>IOCU 1 (if present)</td> </tr> <tr> <td>GCR_ARB_PRI[28]</td> <td>IOCU 0 (if present)</td> </tr> <tr> <td>GCR_ARB_PRI[27]</td> <td>CORE 3 (if present)</td> </tr> <tr> <td>GCR_ARB_PRI[26]</td> <td>CORE 2 (if present)</td> </tr> <tr> <td>GCR_ARB_PRI[25]</td> <td>CORE 1 (if present)</td> </tr> <tr> <td>GCR_ARB_PRI[24]</td> <td>CORE 0 (if present)</td> </tr> </table>	<u>Register Bit</u>	<u>Requestor</u>	GCR_ARB_PRI[29]	IOCU 1 (if present)	GCR_ARB_PRI[28]	IOCU 0 (if present)	GCR_ARB_PRI[27]	CORE 3 (if present)	GCR_ARB_PRI[26]	CORE 2 (if present)	GCR_ARB_PRI[25]	CORE 1 (if present)	GCR_ARB_PRI[24]	CORE 0 (if present)	RW	0x00
<u>Register Bit</u>	<u>Requestor</u>																	
GCR_ARB_PRI[29]	IOCU 1 (if present)																	
GCR_ARB_PRI[28]	IOCU 0 (if present)																	
GCR_ARB_PRI[27]	CORE 3 (if present)																	
GCR_ARB_PRI[26]	CORE 2 (if present)																	
GCR_ARB_PRI[25]	CORE 1 (if present)																	
GCR_ARB_PRI[24]	CORE 0 (if present)																	
<i>ARB_PRI_PRESENT</i>	23	Indicates that the CM2 arbiter supports priority. If this value is 0, then all other fields in this register have no effect.	R	Build Time Option														
R	22:21	Reads as 0x0. Must be written with a value of 0x0.	R	0														
<i>LIMITSEL</i>	20	This bit selects either the number of clocks or the number of requests in the STARVE_CNT_LIMIT field. This bit is encoded as follows: 0: The STARVE_CNT_LIMIT field refers to the number of clock cycles. 1: The STARVE_CNT_LIMIT field refers to the number of requests.	RW	0														
<i>STARVE_CNT_LIMIT</i>	19:0	Starvation Count Limit. If there is at least one low priority request waiting to be serviced for either the number of clocks or requests indicated by the Starvation Count Limit, the CM arbiter will temporarily ignore the priority until all waiting low priority requests have been serviced. This field can contain either a value in clock cycles of a value in requests as specified by the LIMITSEL field	RW	x00000														

### 8.3.4.5 IOCU Revision Register (GCR\_IOCU1\_REV Offset 0x0200)

This register gives the existence and revision information for an IOCU which might be connected to the CM2.

**Figure 8.36 IOCU Revision Register Format**



**Table 8.39 IOCU Revision Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:16	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0
MAJOR_REV	15:8	This field reflects the major revision of the IOCU attached to the CM2. A major revision might reflect the changes from one product generation to another. The value of 0x0 means that no IOCU is attached.	R	Preset
MINOR_REV	7:0	This field reflects the minor revision of the IOCU attached to the CM2. A minor revision might reflect the changes from one release to another.	R	Preset


### 8.3.5 CM2 Attribute-Only Region Address Map Registers

This section contains the base address and address mask registers for CM2 attribute-only regions 0 through 3. These register have the same functionality as the normal region registers, except they can not be used to map to MMIO vs. memory.

#### 8.3.5.1 CM2 Attribute-Only Region [0 - 3] Base Address Registers (GCR\_REGn\_ATTR\_BASE Offsets 0x0190, 0x01A0, 0x0210, 0x0220)

Some or all of these registers may be removed during IP configuration. These registers are similar to the CM2 Region Address Register except the attribute-only regions can not be used to determine if a request is routed to memory or the IOCU.

**Figure 8.37 CM2 Attribute-Only Region [0 - 3] Register Format**



**Table 8.40 CM2 Attribute-Only Region [0 - 3] Base Address Register Format**

Name	Bits	Description	Read/Write	Reset State
CM2_REGION_BASE_ADDR	31:16	This field sets the base physical address of the memory region.	R/W	Undefined
RESERVED	15:0	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0

### 8.3.5.2 CM Attribute-Only Region[0 - 3] Address Mask Registers (GCR\_REGn\_ATTR\_MASK Offsets 0x0198, 0x1A8, 0x218, 0x228)

These registers may be removed during IP Configuration. These registers are similar to the CM Region Address Mask registers except they may not be used to route requests to memory or the IOCU.

**Figure 8.38 CM2 Attribute Only Region [0-3] Address Mask Register Format**

31	15 14	8 7	5	4	3	2	1 0
CM2_REGION_ADDR_MASK	R	CCA_Override_Value	CCA_Override_EN	R	DROP_L2	R	

**Table 8.41 CM Attribute-Only Region [0 - 3] Address Mask Register Descriptions**

Register Fields		Description	Read/Write	Reset State																										
Name	Bits																													
<i>CM2_REGION_ADDR_MASK</i>	31:16	This field is used to set the size of the CM Region. This field is used along with its equivalent CM Region Base Address Register. The request address is logically ANDed with the value of this register. The value of the associated Base Address Register is also logically ANDed with the value of this register. If both outputs match, then the request is routed to the CM region. The only allowed values in this register are contiguous sets of leading 0x1's. An 0x1 preceded by a 0x0 is not allowed (e.g., the value of 0xfff0 is allowed, but the value 0xffef is not allowed).	R/W	Undefined																										
<i>RESERVED</i>	15:8	Reads as 0x0. Must be written with a value of 0x0.	R	0																										
<i>CCA_Override_Value</i>	7:5	Used with CCA_Override_Enable to force the Cache Coherence Attribute (CCA) value for transactions on the system memory OCP. See CCA_Override_Enable field. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Name</th> <th>CCA</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>WT</td> <td>Write Through</td> </tr> <tr> <td>0x1</td> <td>-</td> <td>Reserved</td> </tr> <tr> <td>0x2</td> <td>UC</td> <td>Uncached</td> </tr> <tr> <td>0x3</td> <td>WB</td> <td>WriteBack cacheable, non-coherent</td> </tr> <tr> <td>0x4</td> <td>CWBE</td> <td rowspan="2">Mapped to WB</td> </tr> <tr> <td>0x5</td> <td>CWB</td> </tr> <tr> <td>0x6</td> <td>-</td> <td>Reserved</td> </tr> <tr> <td>0x7</td> <td>UCA</td> <td>Uncached Accelerated</td> </tr> </tbody> </table>	Encoding	Name	CCA	0x0	WT	Write Through	0x1	-	Reserved	0x2	UC	Uncached	0x3	WB	WriteBack cacheable, non-coherent	0x4	CWBE	Mapped to WB	0x5	CWB	0x6	-	Reserved	0x7	UCA	Uncached Accelerated	R/W	0
Encoding	Name	CCA																												
0x0	WT	Write Through																												
0x1	-	Reserved																												
0x2	UC	Uncached																												
0x3	WB	WriteBack cacheable, non-coherent																												
0x4	CWBE	Mapped to WB																												
0x5	CWB																													
0x6	-	Reserved																												
0x7	UCA	Uncached Accelerated																												
<i>CCA_Override_Enable</i>	4	If set CCA_Override_Enable is set to 1 and CM_TARGET is set to Memory, then transactions with addresses that map to this region will have a CCA value set to CCA_Override_Value when driven to system memory.	R/W	0																										
<i>RESERVED</i>	3	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0																										

**Table 8.41 CM Attribute-Only Region [0 - 3] Address Mask Register Descriptions (continued)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>DROP_L2</i>	2	Set to 1 for the CM to drop L2 CacheOp writes after it has been serialized. If set to 0, the L2 CacheOp writes behaves like a regular L2 CacheOp request.	R/W	0x0
<i>RESERVED</i>	1:0	Reads as 0x0. Must be written with a value of 0x0. Since the attribute-only registers can not be used to map to MMIO vs. memory, this field is not needed and is reserved.	R/W	0x0

## 8.4 Core-Local and Core-Other Control Blocks

### 8.4.1 Core-Local and Core-Other Control Blocks Address Map

A set of these registers exists for each core in the interAptiv MPS. These registers can also be accessed from other cores by first writing the *Core Other Addressing Register* (in the Core-Local Control Block) with the proper core number and then accessing these registers using the Core Other Register block.

All registers are 32 bits wide and should only be accessed using 32-bit uncached load/stores. Reads from unpopulated registers in the GCR address space return 0x0, and writes to those locations are silently dropped without generating any exceptions.

**Table 8.42 Core Local and Core Other Block Register Map (Relative to Core-Local/Core-Other CB Offset)**

Register Offset	Name	Type	Description
0x0000	Reserved	-	Reserved
0x0008	Core Local Coherence Control Register ( <i>GCR_CL_COHERENCE</i> <i>GCR_CO_COHERENCE</i> )	R/W	Controls which coherent intervention transactions apply to the local core.
0x0010	Core Local Config Register ( <i>GCR_CL_CONFIG</i> <i>GCR_CO_CONFIG</i> )	R	Contains configuration parameters for the Core-Local address space.
0x0018	Core Other Addressing Register ( <i>GCR_CL_OTHER</i> <i>GCR_CO_OTHER</i> )	R/W	Used to access the registers of another core.
0x0020	Core Local Reset Exception Base Register ( <i>GCR_CL_RESET_BASE</i> <i>GCR_CO_RESET_BASE</i> )	R/W	Sets the Reset Exception Base for the local core.
0x0028	Core Local Identification Register ( <i>GCR_CL_ID</i> <i>GCR_CO_ID</i> )	R	Indicates the ID number of the local core.
0x0030	Core Local Reset Exception Extended Base ( <i>GCR_CL_RESET_EXT_BASE</i> <i>GCR_CO_RESET_EXT_BASE</i> )	R/W	Extends the capabilities of the Core Local Reset Exception Base Register.

**Table 8.42 Core Local and Core Other Block Register Map (Relative to Core-Local/Core-Other CB Offset)(continued)**

Register Offset	Name	Type	Description
0x0040	Core Local TCID_0_PRIORITY Register (GCR_CL_TCID_0_PRIORITY GCR_CO_TCID_0_PRIORITY)	R/W	TCID 0 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0048	Core Local TCID_1_PRIORITY Register (GCR_CL_TCID_1_PRIORITY GCR_CO_TCID_1_PRIORITY)	R/W	TCID 1 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0050	Core Local TCID_2_PRIORITY Register (GCR_CL_TCID_2_PRIORITY GCR_CO_TCID_2_PRIORITY)	R/W	TCID 2 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0058	Core Local TCID_3_PRIORITY Register (GCR_CL_TCID_3_PRIORITY GCR_CO_TCID_3_PRIORITY)	R/W	TCID 3 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0060	Core Local TCID_4_PRIORITY Register (GCR_CL_TCID_4_PRIORITY GCR_CO_TCID_4_PRIORITY)	R/W	TCID 4 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0068	Core Local TCID_5_PRIORITY Register (GCR_CL_TCID_5_PRIORITY GCR_CO_TCID_5_PRIORITY)	R/W	TCID 5 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0070	Core Local TCID_6_PRIORITY Register (GCR_CL_TCID_6_PRIORITY GCR_CO_TCID_6_PRIORITY)	R/W	TCID 6 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0078	Core Local TCID_7_PRIORITY Register (GCR_CL_TCID_7_PRIORITY GCR_CO_TCID_7_PRIORITY)	R/W	TCID 7 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
0x0080	Core Local TCID_8_PRIORITY Register (GCR_CL_TCID_8_PRIORITY GCR_CO_TCID_8_PRIORITY)	R/W	TCID 8 Priority value (2 bits) if IOCU_TYPE=0 in GCR_Cx_CONFIG.
All Others	RESERVED	-	Reserved for future expansion.

#### 8.4.1.1 Core Local Coherence Control Register (GCR\_Cx\_COHERENCE Offset 0x0008)

This register allows each core to respond to intervention requests from only a subset of the coherent masters within the interAptiv Multiprocessing System (MPS). Software can control entry and exit from the coherence domain by setting the *COH\_DOMAIN\_EN* bit in this register for:

- Initialization during (asynchronous) boot
- Power control for shutting down and bringing up a core

**Table 8.43 Core Local Coherence Control Register**

Name	Bits	Description	Read/Write	Reset State
RESERVED	31:8	Reads as 0. Writes ignored. Must be written with a value of 0x0.	W	0x0

**Table 8.43 Core Local Coherence Control Register (continued)**

Name	Bits	Description	Read/ Write	Reset State
<i>COH_DOMAIN_EN</i>	7:0	<p>Each bit in this field represents a coherent requester within the MPS. Setting a bit within this field will enable interventions to this Core from that requester. This field is encoded as follows:</p> <ul style="list-style-type: none"> <li>Bit 0: core 0</li> <li>Bit 1: core 1</li> <li>Bit 2: core 2</li> <li>Bit 3: core 3</li> <li>Bit 4: IOCU 0</li> <li>Bit 5: IOCU 1</li> <li>Bit 6-7: Reserved</li> </ul> <p>The requestor bit which represents the local core is used to enable or disable coherence mode in the local core. Changing the coherence mode for a local core from 0x1 to 0x0 can only be done after flushing and invalidating all the cache lines in the core. Otherwise the system behavior is UNDEFINED. Refer to <a href="#">Section 8.2.11, "Coherency Domains"</a> for more information on the encoding of this field.</p>	R/W	0x0

### 8.4.1.2 Core Local Config Register

Figure 8.39 Core Local Config Register Format




Table 8.44 Core Local Config Register (GCR\_Cx\_CONFIG Offset 0x0010)

Name	Bits	Description	Read/Write	Reset State																		
RESERVED	31:12	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	-																		
<i>IOCU_TYPE</i>	11:10	<table border="1" style="width: 100%; border-collapse: collapse; margin-bottom: 10px;"> <thead> <tr> <th style="width: 15%;">Encoding</th> <th style="width: 85%;">Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>This is a interAptiv core and not an IOCU<sup>1</sup>. Only the interAptiv core can access priority values in the GCR_Cx_TCID_n_PRIORITY registers.</td> </tr> <tr> <td>0x1</td> <td>This is a non-caching IOCU (no intervention port). The IOCU does not access the GCR_Cx_TCID_n_PRIORITY registers.</td> </tr> <tr> <td>0x2</td> <td>This is a caching IOCU (not currently implemented by MIPS).</td> </tr> <tr> <td>0x3</td> <td>Reserved</td> </tr> </tbody> </table> <p>1. Note that the first encoding is redundant information for convenience. It is possible for the system to determine if a core is an IOCU or not by reading the <i>Global Config</i> register.</p>	Encoding	Meaning	0x0	This is a interAptiv core and not an IOCU <sup>1</sup> . Only the interAptiv core can access priority values in the GCR_Cx_TCID_n_PRIORITY registers.	0x1	This is a non-caching IOCU (no intervention port). The IOCU does not access the GCR_Cx_TCID_n_PRIORITY registers.	0x2	This is a caching IOCU (not currently implemented by MIPS).	0x3	Reserved	R	IP Configurable Value								
Encoding	Meaning																					
0x0	This is a interAptiv core and not an IOCU <sup>1</sup> . Only the interAptiv core can access priority values in the GCR_Cx_TCID_n_PRIORITY registers.																					
0x1	This is a non-caching IOCU (no intervention port). The IOCU does not access the GCR_Cx_TCID_n_PRIORITY registers.																					
0x2	This is a caching IOCU (not currently implemented by MIPS).																					
0x3	Reserved																					
<i>PVPE</i>	9:0	<p>Number of VPE's in the system.</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-bottom: 10px;"> <thead> <tr> <th style="width: 15%;">Encoding</th> <th style="width: 85%;">Meaning</th> </tr> </thead> <tbody> <tr><td>0x0</td><td>1 VPE</td></tr> <tr><td>0x1</td><td>2 VPE's</td></tr> <tr><td>0x2</td><td>3 VPE's</td></tr> <tr><td>0x3</td><td>4 VPE's</td></tr> <tr><td>0x4</td><td>5 VPE's</td></tr> <tr><td>0x5</td><td>6 VPE's</td></tr> <tr><td>0x6</td><td>7 VPE's</td></tr> <tr><td>0x7</td><td>8 VPE's</td></tr> </tbody> </table>	Encoding	Meaning	0x0	1 VPE	0x1	2 VPE's	0x2	3 VPE's	0x3	4 VPE's	0x4	5 VPE's	0x5	6 VPE's	0x6	7 VPE's	0x7	8 VPE's	R	IP Configurable Value
Encoding	Meaning																					
0x0	1 VPE																					
0x1	2 VPE's																					
0x2	3 VPE's																					
0x3	4 VPE's																					
0x4	5 VPE's																					
0x5	6 VPE's																					
0x6	7 VPE's																					
0x7	8 VPE's																					

### 8.4.1.3 Core-Other Addressing Register

This register must be written with the correct core number before accessing the Core-Other address segment.

**Figure 8.40 Core Local Config Register Format**




**Table 8.45 Core-Other Addressing Register (GCR\_Cx\_OTHER Offset 0x0018)**

Name	Bits	Description	Read/Write	Reset State
<i>CORENUM</i>	31:16	Core number of the register set to be accessed in the Core-Other address space.	R/W	0x0
RESERVED	15:0	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.-	R	-

### 8.4.1.4 Core Local Reset Exception Base Register (GCR\_Cx\_RESET\_BASE Offset 0x0020)

This register is used to drive the *SI\_ExceptionBase[31:12]* input to the local core. The value is used for placing the exception vectors within the virtual address map during core boot-up time (e.g., when *COP0 Status<sub>BEV</sub>* = 1). The value in this register is reset only on Cold Reset (not Warm Reset).

**Figure 8.41 Core Local Reset Exception Base Register Format**



**Table 8.46 Core Local Reset Exception Base Register**

Name	Bits	Description	Read/Write	Cold Reset State
<i>BEVEXCBase</i>	31:12	Bits [31:12] of the virtual address that the local core will use as the exception base in the boot environment ( <i>COP0 Status<sub>BEV</sub></i> =1).	R/W	IP Configuration Value. MIPS Default Value is 0xBFC00
RESERVED	11:0	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	-

For Core 0, the user can configure the reset location at IP configuration.

Core 0 can write the register to force any of the other cores to use a different reset vector. This register write is done before releasing the other core from reset.

This allows a subset of the processor cores to boot one operating system while another subset of the processor cores boot a different operating system.



### 8.4.1.5 Core Local Identification Register (GCR\_Cx\_ID Offset 0x0028)

The aliased memory scheme is normally invisible to software when accessing GCR registers within the Core-Local control block. What actually happens is that an offset is used to make a subset of the GCR registers appear in the Core-Local addressing window.

This register reports the core number that is used as the addressing offset for the Core-Local control block.

**Figure 8.42 Core Local Identification Register Format**



**Table 8.47 Core Local Identification Register**

Name	Bits	Description	Read/Write	Reset State
<i>CORENUM</i>	31:0	This number is used as an index to the registers within the GCR when accessing the Core-local control block for this core.	R	-

### 8.4.1.6 Core Local Reset Exception Extended Base Register (GCR\_Cx\_RESET\_EXT\_BASE Offset 0x0030)

This register is an extension to the Core-Local Reset Exception Base Register (see Section 8.4.1.4 “Core Local Reset Exception Base Register (GCR\_Cx\_RESET\_BASE Offset 0x0020)”). It also is used to drive the *SI\_ExceptionBase* input to the local core. The value is used for placing the exception vectors within the virtual address map during core boot-up time (e.g., when *COP0 Status<sub>BEV</sub>*=1). The value in this register is reset only on Cold Reset (not Warm Reset).

**Figure 8.43 Core Local Exception Extended Base Register Format**

31	30	29	28	27	20	19	8	7	1	0
EVAReset	UEB	R	BEVExceptionBaseMask				R	BEVExceptionBasePA		PRESENT

**Table 8.48 Core Local Reset Exception Extended Base Register**

Name	Bits	Description	Read/Write	Cold Reset State
<i>EVAReset</i>	31	Assertion of this bit indicates to the core to come up in the EVA configuration at reset. This bit is originally set based on the state of the <i>EVAReset</i> pin during reset.	R/W	IP Configuration Value. MIPS Default Value is 0
<i>UseExceptionBase</i>	30	UseExceptionBase address. This bit reflects the state of the <i>SI_UseExceptionBase</i> pin at reset.  In the legacy configuration, if the <i>SI_UseExceptionBase</i> pin is not asserted, then the BEV location defaults to 0xBFC0_0000.  If the <i>SI_UseExceptionBase</i> pin is asserted, address bits <i>SI_ExceptionBase</i> [31:30] are forced to a value of 2'b10 to force the BEV location into the KSEG0/ KSEG1 space.  Refer to Section 3.7.2 in Chapter 3 for more information. This pin is only used in the legacy configuration. There is one <i>SI_UseExceptionBase</i> pin per core.	R/W	IP Configuration Value. MIPS Default Value is 1
RESERVED	29:28	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	-
<i>BEVExceptionBaseMask</i>	27:20	This field is used to determine the size of the boot exception vector overlay region from 1 MB to 256 MB in powers of two. This field reflects the state of the <i>SI_ExceptionBaseMask</i> [27:20] pins at reset.  This field is used to mask bits [27:20] of the virtual address that the local core will use as the exception base in the boot environment ( <i>COP0 Status<sub>BEV</sub></i> = 1). These pins are used in both the legacy and EVA configurations. There is one set of <i>SI_ExceptionBaseMask</i> pins per core.  Refer to Section 3.7.2 in Chapter 3 for more information.	R/W	IP Configuration Value. MIPS Default Value is 0x00
RESERVED	19:8	Reads as 0x0. Must be written with a value of 0x0.	R	-

**Table 8.48 Core Local Reset Exception Extended Base Register (continued)**

Name	Bits	Description	Read/ Write	Cold Reset State
<i>BEVExceptionBasePA</i>	7:1	<p>BEV exception base physical address. This field contains the upper bits of the physical address that the local core will use as the exception base in the boot environment (<i>COP0 Status<sub>BEV</sub></i> = 1), and reflects the state of the <i>SI_ExceptionBasePA[31:29]</i> pins at reset.</p> <p>The size of the overlay region defined by <i>SI_ExceptionBaseMask[27:20]</i> is remapped to a location in physical address space pointed to by the <i>SI_ExceptionBasePA[31:29]</i> pins. This allows the overlay region to be placed into one of the 512 MB segments in physical memory. These pins are used in both the legacy and EVA configurations. There is one set of <i>SI_ExceptionBasePA</i> pins per core.</p> <p>Note that the bits of this register correspond to upper address bits 35:29. However, in the interAptiv core only the lower three bits (31:29) are used, which correspond to bits 3:1 of this field. The upper four bits are reserved for future cores which implement a 36-bit address. This bit should always be driven with a value of 0x0.</p> <p>Refer to Section 3.7.2 in Chapter 3 for more information.</p>	R/W	IP Configuration Value. MIPS Default Value is 0x00.
PRESENT	0	Reads as 0x1. Writes are ignored	R	1


### 8.4.1.7 Core Local TCID Registers (GCR\_Cx\_TCID\_n\_PRIORITY)

The CM2 contains nine TCID registers, one for each thread, located at the following offset addresses:

**Table 8.49 TCID Priority Offset Addresses**

Register	Offset
Core Local TCID Register 0	0x0040
Core Local TCID Register 1	0x0048
Core Local TCID Register 2	0x0050
Core Local TCID Register 3	0x0058
Core Local TCID Register 4	0x0060
Core Local TCID Register 5	0x0068
Core Local TCID Register 6	0x0070
Core Local TCID Register 7	0x0078
Core Local TCID Register 8	0x0080

**Figure 8.44 Core Local TCID Register Format**



**Table 8.50 Core Local TCID Register Description**

Name	Bits	Description	Read/Write	Reset State
Reserved	31:2	Reads as 0x0. Must be written with a value of 0x0.	R	0x0000_000
<i>TCID_PRIORITY</i>	1:0	TCID priority. This 2-bit value contains the thread context priority level and is encoded as follows: 00: Lowest priority .... 11: Highest priority	R	0x0

## 8.5 Global Debug Control Block

### 8.5.1 Global Debug Control Block Address Map

This block holds registers which are used for debugging the CM2 and software which uses the coherence features supplied by the CM2. The registers associated with PDTrace are reset upon assertion of the TAP controller reset. The other registers in this block are reset when the CM2 is reset. TAP reset occurs when *PB\_EJ\_TRST\_N* is asserted or the Test-Logic-Reset TAP state is entered.

**Table 8.51 Global Debug Block Register Map (Relative to Global Debug Block Offset)**

Register Offset	Name	Type	Reset Source	Description
0x0008	PDTrace TCBControlB Register ( <i>GCR_DB_TCBCONTROLB</i> )	R/W	TAP	Controls how the TCB deals with the trace information. This register only exists if the CM2 is configured with PDTrace.
0x0010	CM2 PDTrace TCBControlD Register ( <i>GCR_DB_TCBCONTROLD</i> )	R/W	TAP	Controls CM2 PDTrace. This register only exists if the CM2 is configured with PDTrace.
0x0020	PDTrace TCBControlE Register ( <i>GCR_DB_TCBCONTROLE</i> )	R/W	TAP	Controls how the TCB deals with trace information. This register only exists if the CM2 is configured with PDTrace.
0x0028	PDTrace TCB Config Register ( <i>GCR_DB_TCBConfig</i> )	R/W	TAP	Contains trace control block configuration information such as probe width, on-trace memory size, and trace clock ratios.
0x0040	PDTrace TCBSYS Register ( <i>GCR_DB_TCBSYS</i> )	R/W	TAP	Controls how external logic uses the System Trace interface. Bit 31 is a PRESENT bit and bits [30:0] are completely user defined. The output of this register is available on the TC_Sys_UserCtl pins. This register only exists if the CM2 is configured with PDTrace.
0x0100	CM2 Performance Counter Control Register ( <i>GCR_DB_PC_CTL</i> )	R/W	CM2	Controls starting/stopping of Performance Counters.
0x0108	PDTrace Trace Word Read Pointer Register ( <i>GCR_DB_TCBRDP</i> )	R/W	TAP	Pointer into the On-Chip Trace Buffer memory for reads from <i>GCR_DB_TCBTW_LO</i> and <i>GCR_DB_TCBTW_HI</i> registers. This register only exists if the CM2 is configured with PDTrace.
0x0110	PDTrace Trace Word Write Pointer Register ( <i>GCR_DB_TCBWRP</i> )	R/W	TAP	Pointer into the On-Chip Trace Buffer memory for the next TraceWord write from <i>GCR_DB_TCBTW_LO</i> and <i>GCR_DB_TCBTW_HI</i> registers. This register only exists if the CM2 is configured with PDTrace.

**Table 8.51 Global Debug Block Register Map (Relative to Global Debug Block Offset)(continued)**

Register Offset	Name	Type	Reset Source	Description
0x0118	PDTrace Trace Word Start Pointer Register ( <i>GCR_DB_TCBSTP</i> )	R/W	TAP	Pointer into On-Chip Trace Buffer that is used to determine when all entries in the trace buffer have been filled. This register only exists if the CM2 is configured with PDTrace.
0x0120	CM2 Performance Counter Overflow Status Register ( <i>GCR_DB_PC_OV</i> )	R/W	CM2	Indicates which performance counters have overflowed.
0x0130	CM2 Performance Counter Event Select Register ( <i>GCR_DB_PC_EVENT</i> )	R/W	CM2	Selects event type of each performance counter.
0x0180	CM2 Performance Cycle Counter Register ( <i>GCR_DB_PC_CYCLE</i> )	R/W	CM2	Counts cycles.
0x0190	CM2 Performance Counter 0 Qualifier Register ( <i>GCR_DB_PC_QUAL0</i> )	R/W	CM2	Performance counter 0 event qualifiers.
0x0198	CM2 Performance Counter 0 Register ( <i>GCR_DB_PC_CNT0</i> )	R/W	CM2	Performance Counter 0 value.
0x01A0	CM2 Performance Counter 1 Qualifier Register ( <i>GCR_DB_PC_QUAL1</i> )	R/W	CM2	Performance counter 1 event qualifiers.
0x01A8	CM2 Performance Counter 1 Register ( <i>GCR_DB_PC_CNT1</i> )	R/W	CM2	Performance Counter 1 value.
0x0200	PDTrace Trace Word Lo Register ( <i>GCR_DB_TCBTW_LO</i> )	R/W	TAP	Access point to read TraceWords from the On-Chip Trace Buffer memory, Least Significant 32-bits.
0x0208	PDTrace Trace Word Hi Register ( <i>GCR_DB_TCBTW_HI</i> )	R/W	TAP	Access point to read TraceWords from the On-Chip Trace Buffer memory, Most Significant 32-bits.
All Others	RESERVED			


All registers are 32 bits wide and should only be accessed using 32-bit uncached load/stores. Reads from unpopulated registers in the GCR address space return 0x0 and writes to those locations should be silently dropped without generating any exceptions.

### 8.5.1.1 CM2 PDTrace TCB ControlB Register (*GCR\_DB\_TCBCONTROLB* Offset 0x0008)

The TCB includes a control register, *GCR\_DB\_TCBCONTROLB* (0x11). This register configures interfaces to the trace buffer. This register only exists if the CM2 is configured with PDTrace.

The format of the *GCR\_DB\_TCBCONTROLB* register is shown below, and the fields are described in [Table 8.52](#).

**Figure 8.45 PDTrace TCB ControlB Register Format**



**Table 8.52 PDTrace TCB ControlB Register**

Fields		Description	Read / Write	Reset State
Name	Bits			
<i>WE</i>	31	Write Enable. Only when set to 1 will the other bits of this register be written. This bit will always read 0.	R	0
<i>Reserved</i>	30:28	Reserved. Must be written as zero; returns zero on read.	R	0
TWSrcWidth	27:26	Used to indicate the number of bits used in the source field of the Trace Word. The value for the CM2 is always 2'b10, indicating a four bit source field width.	R	2'b10
Reserved	25:20	This field is used by EJTAG to access other PDTtrace registers. Although the field is R/W via core accesses, this field has no function for core accesses.	R/W	0
STCE	19	System Trace capture enable. When asserted, the System Trace port of the Funnel is enabled to capture System Trace stream data. When not asserted, System Trace stream data is not captured regardless of <i>TC_Sys_Valid[1:0]</i> input pin state.	R/W	0
TRPAD	18	Trace RAM access disable bit. When set, core reads and writes to the on-chip trace RAM using GCR accesses are inhibited.  If TRPAD is set, memory-mapped writes to the <i>GCR_DB_TCBTW_LO</i> and <i>GCR_DB_TCBTW_HI</i> registers have no effect, and memory-mapped reads from <i>GCR_DB_TCBTW_LO</i> and <i>GCR_DB_TCBTW_HI</i> do not access the Trace RAM and 0 is returned.  Also, when TRPAD is set, then memory-mapped writes to the following registers are inhibited:  <i>TCBTW</i> <i>TCBRDP</i> <i>TCBWRP</i> <i>TCBSTP</i>	R/W	0
Reserved	17	Reserved. Must be written as zero; returns zero on read.	R	0
<i>RM</i>	16	Read on-chip trace memory. When this bit is set, the read address-pointer of the on-chip memory in register <i>TCBRDP</i> is set to the value held in <i>TCBSTP</i> . Subsequent access to the <i>TCBTW</i> register (through the <i>TCBDATA</i> register), will automatically increment the read pointer in register <i>TCBRDP</i> after each read. When the write pointer is reached, this bit is automatically reset to 0, and the <i>TCBTW</i> register will read all zeros. Once set to 1, writing 1 again will have no effect. The bit is reset by setting the TR bit or by reading the last Trace word in <i>TCBTW</i> .	R/W	0

**Table 8.52 PDTrace TCB ControlB Register (continued)**

Fields		Description	Read / Write	Reset State										
Name	Bits													
<i>TR</i>	15	Trace memory reset. When written to one, the address pointers for the on-chip trace memory <i>TCBSTP</i> , <i>TCBRDP</i> and <i>TCBWRP</i> are reset to zero. Also the RM and BF bits are reset to 0. This bit is automatically reset back to 0, when the reset specified above is completed.	R/W1	0										
<i>BF</i>	14	Buffer Full indicator that the TCB uses to communicate to external software that the on-chip trace memory is full. This bit is cleared when writing a 1 to the TR bit. This bit has no function if on-chip memory is not implemented.	R	0										
<i>TM</i>	13:12	Trace Mode. This field determines how the trace memory is filled when using the simple-break control in the PDtrace™ IF to start or stop trace. <table border="1" data-bbox="646 758 1052 949"> <thead> <tr> <th>TM</th> <th>Trace Mode</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Trace-To</td> </tr> <tr> <td>01</td> <td>Trace-From</td> </tr> <tr> <td>10</td> <td>Reserved</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table> In Trace-To mode, the on-chip trace memory is filled, continuously wrapping around, overwriting older Trace Words, as long as there is trace data coming from the core. In Trace-From mode, the on-chip trace memory is filled from the point that the core starts tracing until the on-chip trace memory is full (when the write pointer address is the same as the start pointer address). If a <i>TCBTRIGx</i> trigger control register is used to start/stop tracing, then this field should be set to Trace-To mode. These bits have no function if on-chip memory is not implemented.	TM	Trace Mode	00	Trace-To	01	Trace-From	10	Reserved	11	Reserved	R/W	0
TM	Trace Mode													
00	Trace-To													
01	Trace-From													
10	Reserved													
11	Reserved													
0	11	Read as Zero. Writes ignored. Must be written with a value of 0x0.	R	0										
<i>CR</i>	10:8	Off-chip Clock Ratio. Writing this field, sets the ratio of the core clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 8.53</a> . <b>Note:</b> As the Probe interface works in double data rate (DDR) mode, a 1:2 ratio indicates one data packet sent per core clock rising edge. These bits have no function if off-chip memory is not implemented.	R/W	3'b100										



**Table 8.52 PDTrace TCB ControlB Register (continued)**


Fields		Description	Read / Write	Reset State																																																												
Name	Bits																																																															
<i>Cal</i>	7	<p>Calibrate off-chip trace interface.</p> <p>If set, the off-chip trace pins will produce the following pattern in consecutive trace clock cycles. If more than 4 data pins exist, the pattern is replicated for each set of 4 pins. The pattern repeats from top to bottom until the Cal bit is de-asserted.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th colspan="4">Calibrations pattern</th> </tr> <tr> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr><td>0</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>1</td><td>1</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>1</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>1</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>1</td></tr> <tr><td>1</td><td>1</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>1</td><td>0</td><td>1</td><td>1</td></tr> <tr><td>0</td><td>1</td><td>1</td><td>1</td></tr> </tbody> </table> <p style="text-align: center; margin-left: 20px;">This pattern is replicated for every 4 bits of <i>TR_DATA</i> pins.</p> <p><b>Note:</b> The clock source of the TCB and PIB must be running. These bits have no function if off-chip memory is not implemented.</p>	Calibrations pattern				3	2	1	0	0	0	0	0	1	1	1	1	0	0	0	0	0	1	0	1	1	0	1	0	1	0	0	0	0	1	0	0	0	0	1	0	0	0	0	1	1	1	1	0	1	1	0	1	1	0	1	1	0	1	1	1	R/W	0
Calibrations pattern																																																																
3	2	1	0																																																													
0	0	0	0																																																													
1	1	1	1																																																													
0	0	0	0																																																													
0	1	0	1																																																													
1	0	1	0																																																													
1	0	0	0																																																													
0	1	0	0																																																													
0	0	1	0																																																													
0	0	0	1																																																													
1	1	1	0																																																													
1	1	0	1																																																													
1	0	1	1																																																													
0	1	1	1																																																													
Reserved	6:2	Read as Zero. Writes ignored. Must be written with a value of 0x0.	R	0																																																												
<i>OfC</i>	1	<p>If set to 1, trace is sent to off-chip memory using <i>TR_DATA</i> pins.</p> <p>If not set, trace info is sent to on-chip memory.</p> <p>This bit is read only if one of these options exists.</p>	R/W	Preset																																																												
EN	0	<p>Funnel Trace Enable. When this bit is set, the trace funnels accepts trace information from the CM2, cores, and/or system trace and writes the information to off-chip or on-chip memory.</p> <p>When this bit is cleared, the trace funnel drops all new trace information from the those sources. The trace information already accepted by the trace funnel is sent to the off-chip or on-chip memory, but new trace information is dropped and not written out.</p>	R/W	0																																																												

**Table 8.53 Clock Ratio Encoding of the CR Field**

Encoding of CR Field	Trace Clock:Core Clock Ratio
3'b000	1:20
3'b001	1:16
3'b010	1:12
3'b011	1:10
3'b100	1:2
3'b101	1:4
3'b110	1:6
3'b111	1:8

**8.5.1.2 CM2 PDTrace TCB ControlID Register (GCR\_DB\_TCBCONTROLID Offset 0x0010)**

**Figure 8.46 PDTrace TCB ControlID Register Format**



**Table 8.54 CM2 PDTrace TCB ControlID Register Descriptions**

Name	Bits	Description	Read/Write	Reset State										
RESERVED	31:30	Reserved.	R/W	0x0										
<i>P6_Ctl</i>	29:28	Provides specific control over tracing transactions on Port 6 of the CM. (the IOCU on 6 core configurations). <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Tracing Enabled, no Address Tracing</td> </tr> <tr> <td>01</td> <td>Tracing Enabled with Address Tracing</td> </tr> <tr> <td>10</td> <td>Reserved</td> </tr> <tr> <td>11</td> <td>Tracing Disabled</td> </tr> </tbody> </table>	Encoding	Description	00	Tracing Enabled, no Address Tracing	01	Tracing Enabled with Address Tracing	10	Reserved	11	Tracing Disabled	R/W	0x0
Encoding	Description													
00	Tracing Enabled, no Address Tracing													
01	Tracing Enabled with Address Tracing													
10	Reserved													
11	Tracing Disabled													
<i>P5_Ctl</i>	27:26	Provides specific control over tracing transactions on Port 5 of the CM2 (core 5). See encoding for <i>P6_Ctl</i> .	R/W	0x0										
<i>P4_Ctl</i>	25:24	Provides specific control over tracing transactions on Port 4 of the CM2 (core 4 on 6 core configurations or the IOCU on 4 core or less configurations). See encoding for <i>P6_Ctl</i> .	R/W	0x0										
<i>P3_Ctl</i>	23:22	Provides specific control over tracing transactions on Port 3 of the CM2 (core 3). See encoding for <i>P6_Ctl</i> .	R/W	0x0										

**Table 8.54 CM2 PDTrace TCB ControlID Register Descriptions (continued)**

Name	Bits	Description	Read/Write	Reset State										
<i>P2_Ctl</i>	21:20	Provides specific control over tracing transactions on Port 2 of the CM2 (core 2). See encoding for <i>P6_Ctl</i> .	R/W	0x0										
<i>P1_Ctl</i>	19:18	Provides specific control over tracing transactions on Port 1 of the CM2 (core 1). See encoding for <i>P6_Ctl</i> .	R/W	0x0										
<i>P0_Ctl</i>	17:16	Provides specific control over tracing transactions on Port 0 of the CM2 (core 0). See encoding for <i>P6_Ctl</i> .	R/W	0x0										
RESERVED	15:12	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0										
<i>TwSrcVal</i>	11:8	The source ID inserted into the Trace Word by the CM. NOTE: When disabling trace by setting <i>Global_CM_En</i> to 0, the value in <i>TwSrcVal</i> continues to be used until all trace messages have been flushed from the CM. Therefore, when writing to this register to disabled, the correct value must still be written into the <i>TwSrcVal</i> field.	R/W	0xF										
<i>WB</i>	7	When this bit is set, Coherent Writeback requests are traced. If this bit is not set, all Coherent Writeback requests are suppressed from the CM2 PDTrace Stream.	R/W	0x0										
<i>ST_En</i>	6	System Trace Enable. Driven to the CM2 output pin <i>TC_Sys_Enable</i> . External logic can use this output to control generation of the System Trace stream.	R/W	0x0										
<i>IO</i>	5	Inhibit Overflow on the CM2 PDTrace FIFO full condition. When set to 0, the CM2 will drop a new PDTrace message if the internal PDTrace FIFOs are full. When set to 1, the CM2 will not drop PDTrace messages, but may stall transactions within the CM2 when the internal PDTrace FIFOs are full.	R/W	0x0										
<i>TLev</i>	4:3	This defines the current trace level being used by CM2 PDtrace: <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>No Timing Information</td> </tr> <tr> <td>01</td> <td>Include Stall Times, Causes</td> </tr> <tr> <td>10</td> <td>Reserved</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	00	No Timing Information	01	Include Stall Times, Causes	10	Reserved	11	Reserved	R/W	0x0
Encoding	Description													
00	No Timing Information													
01	Include Stall Times, Causes													
10	Reserved													
11	Reserved													
<i>AE</i>	2	When set to 1, address tracing is always enabled for the CM. When set to 0, address tracing may be enabled on a per-port basis through the P<x>_Ctl bits.	R/W	0x0										
<i>Global_CM_En</i>	1	Setting this bit to 1 enables tracing from the CM2 as long as the <i>CM_EN</i> bit is also enabled.	R/W	0x0										
<i>CM_EN</i>	0	This is the master trace enable for the CM. When zero, tracing from the CM2 is always disabled. When set to one, tracing is enabled from whenever the other enabling functions are also true.	R/W	0x0										

See [Section 8.5.1.2, "CM2 PDTrace TCB ControlD Register \(GCR\\_DB\\_TCBCONTROLD Offset 0x0010\),"](#) for more information about how this register is used.

This register only exists if the CM2 is configured with PDTrace.

### 8.5.1.3 CM2 PDTrace TCB ControlE Register (GCR\_DB\_TCBCONTROLE Offset 0x0020)

**Figure 8.47 PDTrace TCB ControlE Register Format**

31	R	9	TrIdle	8	WB	7	R	1	PeC	0
----	---	---	--------	---	----	---	---	---	-----	---

**Table 8.55 TCBCONTROLE Register**

Name	Bits	Description	Read / Write	Reset State
0	31:9	Reserved for future use. Must be written as zero; returns zero on read.	0	0
<i>TrIdle</i>	8	Trace Unit Idle. This bit indicates if the trace hardware is currently idle (not processing any data). This can be useful when switching control of trace from hardware to software and vice versa. The bit is read-only and updated by the trace hardware. TrIdle is set when the system traces on all cores, and the CM2, have disabled PDTrace and the trace funnel has written all outstanding trace information to the off-chip or on-chip memory.	R	1
0	7:1	Reserved for future use; Must be written as zero; returns zero on read. (Hint to architect, Reserved for future expansion of performance counter trace events).	0	0
<i>PeC</i>	0	Performance Control Tracing is not implemented.	R	0

This register only exists if the CM2 is configured with PDTrace.

See [Section 8.5.1.3, "CM2 PDTrace TCB ControlE Register \(GCR\\_DB\\_TCBCONTROLE Offset 0x0020\),"](#) for more information about how this register is used.

### 8.5.1.4 CM2 PDTrace TCB Config Register (GCR\_DB\_TCBConfig Offset 0x0028)

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

**Figure 8.48 PDTrace TCB Config Register Format**

31	30	21	20	17	16	14	13	11	10	9	8	6	5	4	3	0
CF1	R	SZ	CRMax	CRMin	PW	R	OnT	OfT	REV							

**Table 8.56 TCBCONFIG Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
<i>CF1</i>	31	This bit is set if a <i>TCBCONFIG1</i> register exists. In this revision, <i>TCBCONFIG1</i> does not exist, and this bit reads zero.	R	0

**Table 8.56 TCBCONFIG Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State										
Name	Bits													
<i>Reserved</i>	30:21	Read as Zero. Writes ignored. Must be written with a value of 0x0.	R	0										
<i>SZ</i>	20:17	On-chip trace memory size. This field holds the encoded size of the on-chip trace memory. The size in bytes is given by $2^{(SZ+8)}$ , i.e., the lowest value is 256 bytes, and the highest is 8 MB. This bit is reserved if on-chip memory is not implemented.	R	Preset										
<i>CRMax</i>	16:14	Off-chip Maximum Clock Ratio. This field indicates the maximum ratio of the core clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 8.53</a> . This bit is reserved if off-chip trace option is not implemented.	R	Preset										
<i>CRMin</i>	13:11	Off-chip Minimum Clock Ratio. This field indicates the minimum ratio of the core clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 8.53</a> . This bit is reserved if off-chip trace option is not implemented.	R	Preset										
<i>PW</i>	10:9	Probe Width: Number of bits available on the off-chip trace interface <i>TR_DATA</i> pins. The number of <i>TR_DATA</i> pins is encoded, as shown in the table. <table border="1" data-bbox="574 961 1073 1155"> <thead> <tr> <th>PW</th> <th>Number of bits used on TR_DATA</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>4 bits</td> </tr> <tr> <td>01</td> <td>8 bits</td> </tr> <tr> <td>10</td> <td>16 bits</td> </tr> <tr> <td>11</td> <td>reserved</td> </tr> </tbody> </table> This field is preset based on input signals to the TCB and the actual capability of the TCB. This bit is reserved if off-chip trace option is not implemented.	PW	Number of bits used on TR_DATA	00	4 bits	01	8 bits	10	16 bits	11	reserved	R	Preset
PW	Number of bits used on TR_DATA													
00	4 bits													
01	8 bits													
10	16 bits													
11	reserved													
Reserved	8:6	Read as Zero. Must be written with a value of 0x0.	R	0										
<i>OnT</i>	5	When set, this bit indicates that on-chip trace memory is present. This bit is preset based on the selected option when the TCB is implemented.	R	Preset										
<i>OffT</i>	4	When set, this bit indicates that off-chip trace interface is present. This bit is preset based on the selected option when the TCB is implemented, and on the existence of a PIB module ( <i>TC_PibPresent</i> asserted).	R	Preset										
<i>REV</i>	3:0	Revision of TCB.	R	0x3										

This register only exists if the CM2 is configured with PDTrace.

### 8.5.1.5 CM2 Performance Counter Control Register (GCR\_DB\_PC\_CTL Offset 0x0100)

**Figure 8.49 CM2 Performance Counter Control Register Format**

31	30	29	28	10
R	Perf-Int_En	Perf_OvF_Stop	R	

9	8	7	6	5	4	3	0
P1_Reset	P1_CountOn	P1_Reset	P1_CountOn	Cycl_Cnt_Reset	Cycl_Cnt_CountOn	Perf_Num_Cnt	

**Table 8.57 CM2 Performance Counter Control Register**

Name	Bits	Description	Read/Write	Reset State
Reserved	31	Read as Zero. Must be written with a value of 0x0.	R	0x0
<i>Perf_Int_En</i>	30	Enable Interrupt on counter overflow. If set to 1, a CM2 performance counter interrupt is generated when any enabled CM2 performance counter overflows.	R/W	0x0
<i>Perf_Ovf_Stop</i>	29	Stop Counting on overflow. If set to 1, all CM2 Performance counters stop counting when any enabled CM2 performance counter overflows i.e., the counter has reached 0xFFFF_FFFF.	R/W	0x0
Reserved	28:10	Read as Zero. Must be written with a value of 0x0.	R	0x0
<i>P1_Reset</i>	9	If set to 1, CM2 Performance Counter 1 and <i>P1_Overflow</i> bit is reset before counting is started. If set to 0 counting is resumed from previous value. This bit is automatically set to 0 when the counter is reset, so <i>P1_Reset</i> is always read as 0.	R/W	0x0
<i>P1_CountOn</i>	8	Start Counting. If this bit is set to 1 then CM2 Performance Counter 1 and the <i>P1_Overflow</i> bit starts counting the specified event. If this bit is set to 0 then CM2 Performance Counter 1 is disabled. This bit is automatically set to 0 if any counter overflows and <i>Perf_Ovf_Stop</i> is set to 1.	R/W	0x0
<i>P0_Reset</i>	7	If set to 1, CM2 Performance Counter 0 and <i>P0_Overflow</i> bit is reset before counting is started. If set to 0 counting is resumed from previous value. This bit is automatically set to 0 when the counter is reset, so <i>P0_Reset</i> is always read as 0.	R/W	0x0
<i>P0_CountOn</i>	6	Start/Stop Counting. If this bit is set to 1 then CM2 Performance Counter 0 starts counting the specified event. If this bit is set to 0 then CM2 Performance Counter 0 is disabled. This bit is automatically set to 0 if any counter overflows and <i>Perf_Ovf_Stop</i> is set to 1.	R/W	0x0
<i>Cycl_Cnt_Reset</i>	5	If set to 1, the <i>CM2 Cycle Counter Register</i> and the <i>Cycl_Cnt_Overflow</i> bit is reset before counting is started. If set to 0 counting is resumed from previous value. This bit is automatically set to 0 when the counter is reset, so <i>Cycl_Cnt_Reset</i> is always read as 0.	R/W	0x0
<i>Cycl_Cnt_CountOn</i>	4	Start/Stop the Cycle Counter. If this bit is set to 1 then CM2 Cycle Counter starts counting. If this bit is set to 0 then CM2 Cycle Counter is disabled. This bit is automatically set to 0 if any Counter Overflows and <i>Perf_Ovf_Stop</i> is set to 1.	R/W	0x0
<i>Perf_Num_Cnt</i>	3:0	The number of performance counters implemented (not including the cycle counter). The CM2 has 2 performance counters.	R	0x2


### 8.5.1.6 CM2 PDTrace TCB Trace Word Read Pointer Register (GCR\_DB\_TCBRDP Offset 0x0108)

The *TCBRDP* register is an address pointer to on-chip trace memory. It points to the TW read when reading the *TCBTW* register. When writing the *TCBCONTROLB<sub>RM</sub>* bit to 1, this pointer is reset to the current value of *TCBSTP*.

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

The format of the *TCBRDP* register is shown below and the fields are described in [Table 8.58](#). The value of n depends on the size of the on-chip trace memory. As the address points to a 64-bit TW, lower three bits are always zero.

**Figure 8.50 TCBRDP Register Format**



**Table 8.58 TCBRDP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
Data	31:(n+1)	Reserved. Must be written with zero, reads back zero.	0	0
Address	n:0	Byte address of on-chip trace memory word.	R/W	0


### 8.5.1.7 CM2 PDTrace TCB Trace Word Write Pointer Register (GCR\_DB\_TCBWRP Offset 0x0110)

The *TCBWRP* register is an address pointer to on-chip trace memory. It points to the location where the next new TW for on-chip trace will be written.

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

The format of the *TCBWRP* register is shown below and the fields are described in [Table 8.59](#). The value of n depends on the size of the on-chip trace memory. As the address points to a 64-bit TW, the lower three bits are always zero.

**Figure 8.51 TCBWRP Register Format**



**Table 8.59 TCBWRP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
Data	31:(n+1)	Reserved. Must be written zero, reads back zero.	0	0
Address	n:0	Byte address of on-chip trace memory word.	R/W	0

### 8.5.1.8 CM2 PDTrace TCB Trace Word Start Pointer Register (GCR\_DB\_TCBSTP Offset 0x0118)

The *TCBSTP* register is the start pointer register. This pointer is used to determine when all entries in the trace buffer have been filled (when *TCBWRP* has the same value as *TCBSTP*). This pointer is reset to zero when the

$TCBCONTROLB_{TR}$  bit is written to 1. If a continuous trace to on-chip memory wraps around the on-chip memory,  $TSBSTP$  will have the same value as  $TCBWRP$ .

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

The format of the  $TCBSTP$  register is shown below and the fields are described in Table 8.60. The value of n depends on the size of the on-chip trace memory. As the address points to a 64-bit TW, the lower three bits are always zero.

**Figure 8.52 TCBSTP Register Format**



**Table 8.60 TCBSTP Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Names	Bits			
Data	31:(n+1)	Reserved. Must be written zero, reads back zero.	0	0
Address	n:0	Byte address of on-chip trace memory word.	R/W	0

**8.5.1.9 CM2 PDTrace TCB System Trace User Control Register ( GCR\_DB\_TCBSYS Offset 0x0040)**

The  $TCBSYS$  register contents are driven to the  $TC\_Sys\_UserCtl[31:0]$  output signals. This register is also mapped to offset 0x0040 in the Global Debug Block of the CM GCRs. Thus, any change to this register will be reflected in these output signals. The format of the  $TCBSYS$  register is shown below, and the fields are described in Table 8.61.

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

**Figure 8.53 TCBSYS Register Format**




**Table 8.61 TCBSYS Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
STA	31	System Trace Available. Set to 1 if the System Trace Interface is present. Otherwise it is set to 0.	R	Preset
UsrCtl	30:0	User-defined Control.	R/W	0

**8.5.1.10 CM2 Performance Counter Overflow Status Register (GCR\_DB\_PC\_OV Offset 0x120)**

**Figure 8.54 Performance Counter Overflow Status Register Format**






**Table 8.62 Performance Counter Overflow Status Register**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
Reserved	31:3	Reserved. Must be written zero, reads back zero.	R	0x0
<i>P1_OF</i>	2	If this bit is set to 1, <i>CM2 Performance Counter 1</i> has overflowed i.e., the counter has reached 0xFFFF_FFFF.	R Write 1 to clear	0x0
<i>P0_OF</i>	1	If this bit is set to 1, <i>CM2 Performance Counter 0</i> has overflowed i.e., the counter has reached 0xFFFF_FFFF.	R Write 1 to clear	0x0
<i>Cycl_Cnt_OF</i>	0	If this bit is set to 1, the <i>CM2 Cycle Counter Register</i> has overflowed.	R Write 1 to clear	0x0

**8.5.1.11 CM2 Performance Counter Event Select Register (GCR\_DB\_PC\_EVENT Offset 0x130)**

**Figure 8.55 CM2 Performance Counter Event Select Register Format**




**Table 8.63 CM2 Performance Counter Event Select Register**

Name	Bits	Description	Read/Write	Reset State
Reserved	31:16	Reserved. Must be written zero, reads back zero.	R	0x0
<i>P1_Event</i>	15:8	Event Selection for CM2 Performance Counter 1. Event numbers are defined in <a href="#">Table 17.1</a> .	R/W	0x0
<i>P0_Event</i>	7:0	Event Selection for CM2 Performance Counter 0. Event numbers are defined in <a href="#">Table 17.1</a> .	R/W	0x0

**8.5.1.12 CM2 Cycle Counter Register**

The CM2 Cycle Count Register is a 32-bit register that keeps count of CM2 clock cycles. It is controlled through the *Cycl\_Cnt\_CountOn* and *Cycl\_Cnt\_Reset* bits in the CM2 Performance Counter Control Register. An overflow of the cycle counter is indicated by a 1 in the *Cycl\_Cnt\_Overflow* bit in the CM2 Performance Counter Overflow Status Register.

**Figure 8.56 CM2 Cycle Count Register Format**




**Table 8.64 CM2 Cycle Counter Register (GCR\_DB\_PC\_CYCLE Offset 0x180)**

Name	Bits	Description	Read/Write	Reset State
<i>Cycle_Cnt</i>	31:0	32-bit count of CM2 clock cycles.	R/W	0x0

**8.5.1.13 CM2 Performance Counter n Qualifier Field Register (GCR\_DB\_PC\_QUALn Offset 0x190, 0x1a0)**

**Figure 8.57 Performance Counter n Qualifier Field Register Format**




**Table 8.65 CM2 Performance Counter n Qualifier Field Register Descriptions**

Name	Bits	Description	Read/Write	Reset State
<i>Pn_Qualifier</i>	31:0	CM2 Performance Counter n Event Qualifier. The qualifier corresponds to the event configured through the <i>Performance Counter 0 Event Select Register</i> .	R/W	0x0

**8.5.1.14 CM2 Performance Counter n Register (GCR\_DB\_PC\_CNTn Offset 0x198, 0x1A8)**

**Figure 8.58 Performance Counter n Register Format**



**Table 8.66 CM2 Performance Counter n Register**

Name	Bits	Description	Read/Write	Reset State
<i>Pn_Count</i>	31:0	32-bit Performance Counter. The event counted is specified in the <i>CM2 Performance Counter Event Select Register</i> and by the corresponding <i>Qualifier Register</i> .	R/W	0x0

**8.5.1.15 CM2 PDTrace TCB Trace Word LO Register ( GCR\_DB\_TCBTW\_LO Offset 0x0200)**


Reads to this register access the contents of the On-Chip Trace Buffer entry (least significant 32-bits) which is referenced by the *GCR\_DB\_TCBRDP* register. Writes to this register modify the On-Chip Trace Buffer entry (least significant 32-bits) which is referenced by the *GCR\_DB\_TCBWRP* register.

A side effect of reading the *TCBTW\_LO* register is that the *TCBRDP* register increments to the next TW in the on-chip trace memory. If *TCBRDP* is at the max size of the on-chip trace memory, the increment wraps back to address zero. A side effect of writing the *TCBTW\_LO* register is that the *TCBWRP* register increments to the next TW in the on-chip trace memory. If *TCBWRP* is at the max size of the on-chip trace memory, the increment wraps back to

address zero. The use of load half-word or load byte instructions can lead to unpredictable results, and is not recommended.

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

**Figure 8.59 TCBTW\_LO Register Format**



**Table 8.67 TCBTW\_LO Register Field Descriptions**

Names	Bits	Description	Read / Write	Reset State
<i>Data</i>	31:0	Trace Word, least significant 32-bits.	R/W	0


**8.5.1.16 CM2 PDTrace TCB Trace Word HI Register ( *GCR\_DB\_TCBTW\_HI* Offset 0x0208)**

Reads to this register access the contents of the On-Chip Trace Buffer entry (most significant 32-bits) which is referenced by the *GCR\_DB\_TCBRDP* register. Writes to this register modify the On-Chip Trace Buffer entry (most significant 32-bits) which is referenced by the *GCR\_DB\_TCBWRP* register.

To read or write a 64-bit trace word from the Trace Buffer, the *GCR\_DB\_TCBTW\_HI* register must be accessed first before the *GCR\_DB\_TCBTW\_LO* register. The access of the *GCR\_DB\_TCBTW\_LO* register causes the appropriate pointer register to be incremented. The use of load half-word or load byte instructions can lead to unpredictable results, and is not recommended.

This register is also accessible by EJTAG via the TCBDATA instruction as described in the EJTAG Debug Support chapter.

**Figure 8.60 TCBTW\_HI Register Format**



**Table 8.68 TCBTW\_HI Register Field Descriptions**

Names	Bits	Description	Read / Write	Reset State
<i>Data</i>	31:0	Trace Word, most significant 32-bits.	R/W	0



# Global Interrupt Controller

This chapter describes the optional Global Interrupt Controller (GIC) included in the interAptiv Multiprocessing System. The GIC can control up to 256 external interrupt sources in multiples of 8. This chapter describes how software controls the configuration and use of the GIC.

The optional GIC is selected at IP Configuration time. A *Yes/No* menu selection in the Graphical User Interface (GUI) selects between the GIC described in this chapter, and no GIC. If *Yes* is selected, the MIPS GIC is instantiated and described here. If *No* is selected, the MIPS GIC is not instantiated and this chapter can be skipped.

The GIC handles the distribution of interrupts between and among the CPUs in the cluster. On a multithreaded CPU with multiple Virtual Processing Elements (VPEs), each VPE will have its own set of interrupt inputs. The GIC has the ability to route interrupts to each VPE independently.

The chapter contains the following sections:

- [Section 9.1 “GIC Terminology”](#)
- [Section 9.2 “GIC Features”](#)
- [Section 9.3 “GIC Address Map Overview”](#)
- [Section 9.4 “GIC Programming”](#)
- [Section 9.5 “Shared Register Set”](#)
- [Section 9.6 “GIC VPE-Local and VPE-Other Register Set”](#)
- [Section 9.7 “GIC User-Mode Visible Section”](#)

## 9.1 GIC Terminology

In the context of the GIC, the term ‘Processor’ will be used synonymously to refer to a single processor or a virtual processor in a Core that implements the MT ASE, such as interAptiv.

In the interAptiv core, which can contain up to two VPE’s per core, the processor numbering is as follows:

**Table 9.1 Processor Numbering with Two VPE’s per Core**

Processor Number	Core Number	VPE Number
0	0	0
1	0	1
2	1	0
3	1	1
4	2	0
5	2	1

**Table 9.1 Processor Numbering with Two VPE's per Core (continued)**

Processor Number	Core Number	VPE Number
6	3	0
7	3	1

When there is one VPE per core, the processor numbering is as follows:

**Table 9.2 Processor Numbering with One VPE per Core**

Processor Number	Core Number	VPE Number
0	0	0
1	1	0
2	2	0
3	3	0

## 9.2 GIC Features

To provide support for a multiprocessor environment, the GIC design includes the following features:

- Accepts interrupts from up to 256 external sources.
- Supports active-high, active-low, rising-edge triggered, falling-edge triggered, and dual-edge triggered interrupt signaling.
- Distributes/partitions the interrupt sources among the available cores and VPEs.
- Steers any interrupt source to any VPE interrupt input (Interrupt pin, NMI, yield qualifier).
- Allows any VPE to interrupt any other VPE.
- Backward compatible with pre-defined MIPS Technologies interrupt modes (legacy, vectored, and EIC).
- Scalable for both the number of interrupt sources as well as the number of VPE in the system.
- Able to integrate interrupt messages from peripherals such as PCI-Express.
- Supports simultaneous multithreading, as defined in the MIPS MT-ASE and implemented by the interAptiv cores, including routing interrupts to an individual Virtual Processing Element (VPE) within a CPU core and routing interrupts to the Yield Qualifier input pins of the CPU.
- Hardware assist features are configurable by software at run-time.
- Provides interval and watchdog timers.

## 9.3 GIC Address Map Overview

An interAptiv Multiprocessing System can contain up to four cores and eight VPEs. To avoid the large address space needed for VPE-specific register sets, an aliasing address scheme is used.

The GIC address space is accessed with uncached load/store commands. The physical address and the VPE number of the requester is supplied for each load/store command. The VPE number is used as an index to reference the appropriate subset of the instantiated control registers. By using the VPE number information, the hardware writes/reads the correct subset of the control registers pertaining to that VPE. Software does not need to explicitly calculate the register index for the core in question; it is done entirely by hardware.

In the interAptiv Multiprocessing System, any VPE can access the registers of any other VPE by using the *VPE-Other* address spaces. Software must write the *VPE-Other Addressing Register* before accessing these address spaces. The value of this register is used by hardware to index the appropriate subset of the control registers.

Two address “windows” are made available to the programmer:

- A window for the “Local” VPE (as specified by the VPE number information).
- A second window for an “Other” VPE that allows a VPE to access the register set belonging to another VPE. The “Other” VPE is specified by first writing the *VPE-Other Addressing Register* in the “local” VPE address space.

An additional section called the *User-Mode Visible section* is used to give quick user-mode read access to specific GIC registers. The use of this section is meant to avoid the overhead of system calls to read GIC resources, such as counter registers.

The address map of the GIC is shown in [Table 9.3](#).

**Table 9.3 GIC Address Space**

Segment	Base Offset	Addressing Method	Address Space Size	Virtual Address Space Type
Shared Section Offset	0x00000	Offset relative to <i>GCR_GIC_Base</i>	32 KB	Kernel
VPE-Local Section Offset	0x08000	Offset relative to <i>GCR_GIC_Base</i> + using VPE number as Index	16 KB	Kernel
VPE-Other Section Offset	0x0C000	Offset relative to <i>GCR_GIC_Base</i> + using <i>VPE-Other Addressing Register</i> as Index	16 KB	Kernel
User-Mode Visible Section Offset	0x10000	Offset relative to <i>GCR_GIC_Base</i>	64 KB	User

As shown in the table above, the GIC address space is divided into four types:

- A *Shared* section in which the external interrupt sources are registered, masked, and assigned to a particular VPE and interrupt pin. This section is used by all VPEs and all cores in the system.
- A *VPE-Local* section in which interrupts local to a VPE are registered, masked, and assigned to a particular interrupt pin. If External Interrupt Controller Mode (EIC) mode is used for a particular VPE, the EIC encoder is instantiated here.
- A *VPE-Other* section in which the local VPE can access the VPE-Local section of another VPE by which the interrupt can be registered, masked, and assigned to a particular interrupt pin of the other VPE. One VPE can setup the GIC for all VPEs in the system using this section.
- A *User Mode Visible* section that contains the GIC Hi/Lo counters accessible in user mode for quick user mode access. The use of this section is meant to avoid the overhead of system calls to read GIC resources, such as counter registers.

In the GIC, the *Shared*, *VPE-Local*, and *VPE-Other* sections are meant to be located in privileged system virtual address space, in which only kernel mode software can initialize and update the interrupt controller.

A separate 64KB address space is allocated so that it may be mapped to *User Mode* virtual address space. Within this address space are aliases for GIC registers that are read so often that it makes sense to make them available to user-mode programs without requiring a system call. The aliases for these registers are read-only. Currently, the only registers that are aliased into this space are the shared *GIC\_SH\_CounterLo* and *GIC\_SH\_CounterHi* registers. Refer to [Section 9.7 “GIC User-Mode Visible Section”](#) for more information.

### 9.3.1 GIC Base Address

The GIC base address is a 17-bit value that is programmed into the *GCR\_CPE\_BASE* field of the *GCR CPC Base* register located at offset address 0x0088 in the Global Control Block of the CM2 registers. Refer to the *GCR\_CPC\_BASE Register* in Chapter 8, *CM2 Global Control Registers* for more information on this register.

### 9.3.2 Block Offsets Relative to the Base Address

The block offsets for each of the three blocks listed in [Table 9.3](#) above are relative to a GIC base address as described above and can be located anywhere in physical memory. To determine the physical address of each block listed in [Table 9.4](#), the base address written to the *GCR\_GIC\_BASE Register* this value would be added to the GIC block offset ranges to derive the absolute physical address as shown in [Table 9.4](#). Note that an example base address of 0x1BDC\_0 is used for these calculations.

**Table 9.4 Example Physical Address Calculation of the GIC Register Blocks**

Example Base Address		GCR Block Offset		Absolute Physical Address	Size (bytes)	Description
0x1BDC_0	+	0x0000 - 0x7FFF	=	0x1BDC_0000 - 0x1BDC_7FFF	32 KB	GIC Shared Control Block.
0x1BDC_0	+	0x8000 - 0xBFFF	=	0x1BDC_8000 - 0x1BDC_BFFF	16 KB	GIC Core-Local Control Block.
0x1BDC_0	+	0xC000 - 0xFFFF	=	0x1BDC_C000 - 0x1BDC_FFFF	16 KB	GIC Core-Other Control Block.
0x1BDC_0	+	0x10000 - 0x1FFFF	=	0x1BDD_0000 - 0x1BDD_FFFF	64 KB	User-Mode Visible Block.

### 9.3.3 Register Offsets Relative to the Block Offsets

In addition to the block offsets, the register offsets provided in each register description of this chapter are relative to the block offsets shown in [Table 9.3](#) above. To determine the physical address of each register, the base address programmed into the *GCR\_GIC\_BASE* register is added to the corresponding GIC block offset described above, plus the actual register offset to derive the absolute physical address as shown in [Table 9.5](#). This table shows the physical address for the first few registers of the GIC Shared block. In this table an example base address of 0x1BDC\_0 is used.

**Table 9.5 Absolute Address of Individual GIC Shared Block Registers**

MIPS Default Base		Global Register Block Offset		Global Register Offset		Absolute Physical Address	Global Control Register
0x1BDC_0	+	0x0000	+	0x0000	=	0x1BDC_0000	GIC Config.
0x1BDC_0	+	0x0000	+	0x0010	=	0x1BDC_0010	GIC CounterLo.
0x1BDC_0	+	0x0000	+	0x0014	=	0x1BDC_0014	GIC CounterHi.
0x1BDC_0	+	0x0000	+	0x0020	=	0x1BDC_0020	GIC Revision.

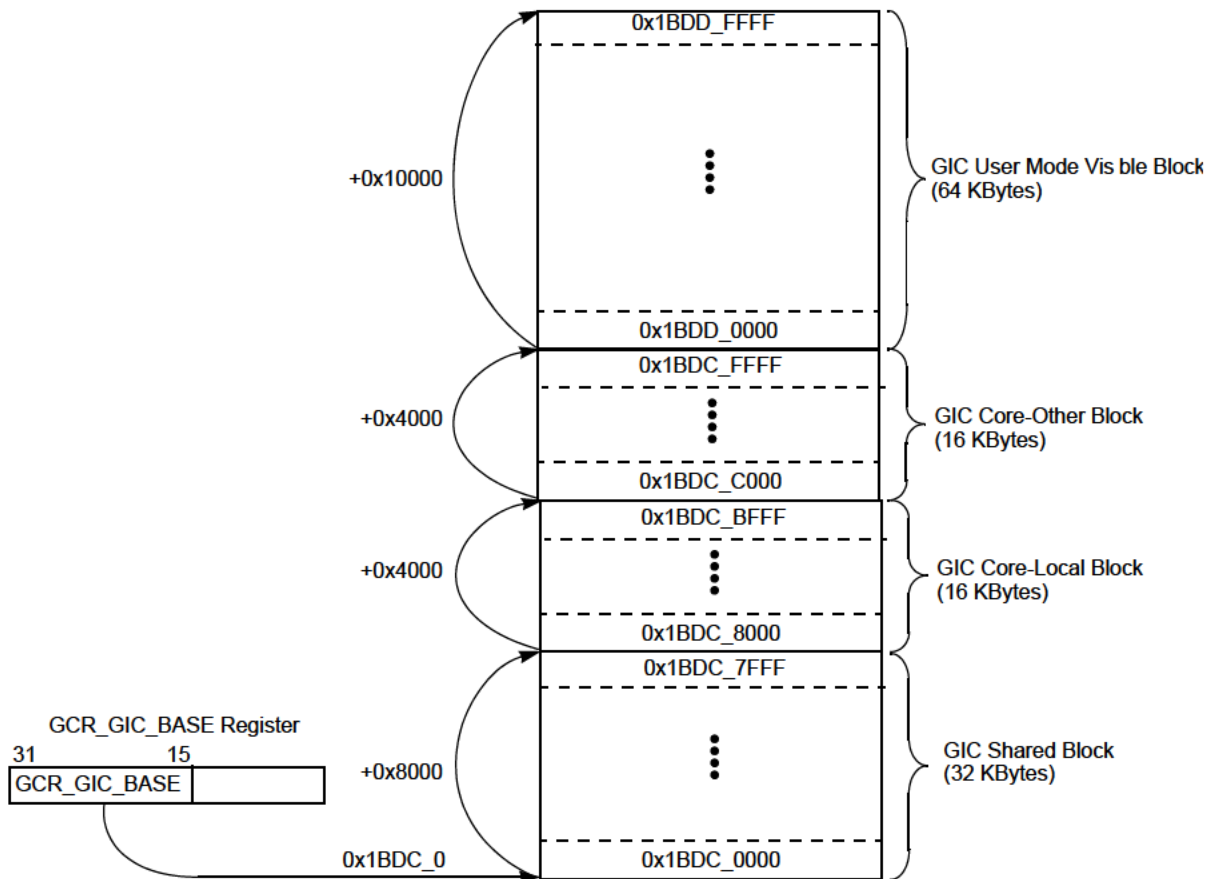


**Table 9.5 Absolute Address of Individual GIC Shared Block Registers(continued)**

MIPS Default Base		Global Register Block Offset		Global Register Offset		Absolute Physical Address	Global Control Register
0x1BDC_0	+	0x0000	+	0x0100	=	0x1BDC_0100	CPC Interrupt Polarity 0.
...	+	...	+	...	=	...	...

This concept is described in [Figure 9.1](#) below. In this figure an example base address of 0x1BDE\_0 is used.

**Figure 9.1 GIC Register Addressing Scheme Using an Example Base Address of 0x1BDC\_0**



## 9.4 GIC Programming

This section covers the programming for the following tasks.

- Setting the GIC Base Address and Enabling the GIC
- Configuration of interrupt sources:
  - External interrupt source configuration:
    - Level Sensitivity, active high or active low
    - Edge Sensitivity, dual or single edge (falling or Rising)
  - Routing of Interrupt external interrupts to specific processors.
- Enabling or Disabling interrupts
- Inter-Processor Interrupts
- Local device interrupt configuration

### 9.4.1 Setting the GIC Base Address and Enabling the GIC

As described in [Section 9.3.1 “GIC Base Address”](#), the base address for the memory mapped registers of the GIC is set using the GIC\_BASE\_ADDR field of the GCR\_GIC\_BASE Register. This field is normally programmed by the boot code executing outside of the boot process.

To enable the GIC the GIC\_EN bit must be set in this same register.

### 9.4.2 Configuring Interrupt Sources

The triggering of interrupts is configured through several registers in the GIC that are shared by all processors. All processors can access these registers but in practice these registers are usually programmed at boot time by processor 0. There are three register groups that control the interrupt triggering configuration.

- Trigger type register group
- Edge type register group
- Polarity register group

Each interrupt source is represented by one bit in each register group. Each register in a group is 32 bits so each register controls 32 interrupt sources. The first register in each group would control interrupts 0 - 31, the next 32 - 63 and so on. Since there can be 256 interrupt sources there could be 8 registers in each group. There are enough of these registers in each group to control the number of interrupt sources implemented. The number of interrupt sources is a fixed value configured at core build time. This number can be determined by reading the NUMINTERRUPTS field of the "GIC Configuration Register", GIC\_SH\_CONFIG. Refer to [Section 9.5.3.1 “Global Config Register”](#) for more information.

Each of the interrupt sources can be of either positive (asserted high) or negative (asserted low) polarity. Similarly, any of these sources can be either level-sensitive, single-edge-sensitive, or dual-edge-sensitive. Through the polarity control registers (*GIC\_SH\_POLx\_y*), the trigger type control registers (*GIC\_SH\_TRIGx\_y*) and dual edge control registers (*GIC\_SH\_DUALx\_y*), all of the sources are normalized to positive, level-sensitive signals. This is the interrupt type supported by the CPU interrupt inputs.

For single-edged signaling, the *Polarity* register denotes which edge is used for setting the interrupt register and which edge is ignored. For double-edged signaling, both the rising and falling edges are used to set the interrupt register. These three registers work in conjunction with one another to define the characteristics of each specific interrupt in the system. Each bit of each register corresponds to an interrupt. So for a given bit, the corresponding interrupt characteristics would be defined as shown in [Table 9.6](#). The ‘n’ in the table entries denotes that it can be any bit of a given register, but must be the same bit of each register.

**Table 9.6 Selecting Interrupt Polarity, Edge Sensitivity, and Triggering**

<b>Polarity (GIC_SH_POL[n])</b>	<b>Trigger (GIC_SH_TRIG[n])</b>	<b>Single/Dual Edge (GIC_SH_DUAL[n])</b>	<b>Description</b>
0	0	x	Interrupt is level sensitive and active low. In this case the contents of the GIC_SH_DUAL have no meaning because level triggering is enabled.
1	0	x	Interrupt is level sensitive and active high. In this case the contents of the GIC_SH_DUAL have no meaning because level triggering is enabled.
0	1	0	Interrupt is single edge triggered on the falling edge of the signal.
1	1	0	Interrupt is single edge triggered on the rising edge of the signal.
x	1	1	Interrupt is dual edge triggered. In this case the contents of the GIC_SH_POL have no meaning because interrupts occur on both the rising and falling edges of the signal.

#### 9.4.2.1 Trigger Type Register Group

The trigger type register group is made up of shared "Global Interrupt Trigger Type Registers", GIC\_SH\_TRIG. The trigger type can be set to level or edge sensitive. Setting the source bit configures the source to be edge sensitive and clearing it configures it to be level sensitive. For example to set the interrupt source 32 to edge sensitive bit 0 of the second GIC\_SH\_TRIG Register should be set. Refer to [Section 9.5.3.6 “Global Interrupt Trigger Type Registers”](#), for more information on how to assign this parameter.

#### 9.4.2.2 Edge Type Register Group

The edge type register group is made up of shared "Global Dual Edge Registers", GIC\_SH\_DUAL. This register group is used if the Trigger type described in the last section is set to edge sensitive and has no effect if the trigger type is level sensitive. The edge type can be either single or dual edge. Setting the source bit configures the source to be dual edge and clearing it configures it to be single edge. For example, to set interrupt source 32 to dual edge sensitive bit 0 of the second Global Dual Edge Registers should be set.

Refer to [Section 9.5.3.7 “Global Interrupt Dual Edge Registers”](#) for more information on how to assign this parameters.

#### 9.4.2.3 Polarity Type Register Group

The polarity register group is made up of shared "Global Interrupt Polarity Registers", GIC\_SH\_POL. This register group is used to determine the polarity sensitivity of the source.

If the interrupt source type is level sensitive then setting the source bit configures the source to be active High, and clearing it configures it to be active low.

If the interrupt is single edge sensitive then setting the source bit configures the source to rising edge toggle and setting clearing it configure it to be falling edge toggle.

This register group has no effect if the edge type was set to dual edge sensitive.

Refer to [Section 9.5.3.5 “Global Interrupt Polarity Registers”](#) for more information on how to assign this parameter.

### 9.4.3 Interrupt Routing

The routing of interrupts to a specific input on a specific processor is controlled by the setting of 2 registers.

- Global Interrupt Map to Processor register, GIC\_SH\_MAP\_VPE — maps the interrupt to a processor.
- Global Interrupt Map to Pin Register, GIC\_SH\_MAP\_PIN — maps interrupt to a specific signal on a processor.

There is one of each of these 32 bit registers for each external interrupt source. The mapping of external interrupt pins and the registers that control them is listed in [Table 9.7](#).

**Table 9.7 Mapping of External Interrupts**

External Interrupt	Offset	Register Name	External Interrupt	Offset	Register Name
0	0x2000	GIC_SH_MAP0_VPE	248	0x3F00	GIC_SH_MAP248_VPE
	0x0500	GIC_SH_MAP0_PIN		0x08E0	GIC_SH_MAP248_PIN
1	0x2020	GIC_SH_MAP1_VPE	249	0x3F20	GIC_SH_MAP249_VPE
	0x0504	GIC_SH_MAP1_PIN		0x08E4	GIC_SH_MAP249_PIN
2	0x2040	GIC_SH_MAP2_VPE	250	0x3F40	GIC_SH_MAP250_VPE
	0x0508	GIC_SH_MAP2_PIN		0x08E8	GIC_SH_MAP250_PIN
3	0x2060	GIC_SH_MAP3_VPE	251	0x3F60	GIC_SH_MAP251_VPE
	0x050C	GIC_SH_MAP3_PIN		0x08EC	GIC_SH_MAP251_PIN
4	0x2080	GIC_SH_MAP4_VPE	252	0x3F80	GIC_SH_MAP252_VPE
	0x0510	GIC_SH_MAP4_PIN		0x08F0	GIC_SH_MAP252_PIN
5	0x20A0	GIC_SH_MAP5_VPE	253	0x3FA0	GIC_SH_MAP253_VPE
	0x0514	GIC_SH_MAP5_PIN		0x08F4	GIC_SH_MAP253_PIN
6	0x20C0	GIC_SH_MAP6_VPE	254	0x3FC0	GIC_SH_MAP254_VPE
	0x0518	GIC_SH_MAP6_PIN		0x08F8	GIC_SH_MAP254_PIN
7	0x20E0	GIC_SH_MAP7_VPE	255	0x3FE0	GIC_SH_MAP255_VPE
	0x051C	GIC_SH_MAP7_PIN		0x08FC	GIC_SH_MAP255_PIN
8 - 247	0x2100 - 0x3EE0	GIC_SH_MAP8_VPE - GIC_SH_MAP247_VPE			
	0x0520 - 0x08DC	GIC_SH_MAP8_PIN - GIC_SH_MAP247_PIN			

### 9.4.3.1 Mapping an Interrupt Source to a Processor

There is one shared "Global Interrupt Map to VPE Register", GIC\_SH\_MAP\_VPE for each interrupt source that maps that source to a processor. Bit 0 would map the interrupt source to processor 0; bit 1 would map the interrupt to processor 1 and so on. Refer to [Section 9.5.3.14 "Global Interrupt Map to VPE Registers"](#) for more information.

### 9.4.3.2 Mapping and Interrupt Source to a Specific Processor Pin

There is one shared "Global Interrupt Map to Pin Register", GIC\_SH\_MAP\_PIN for each external interrupt source that further maps that source to a specific signal on the processor. There are two bits that control the type of signals that can be assigned to the interrupt source. Refer to [Section 9.5.3.13 "Global Interrupt Map to Pin Registers"](#) for more information.

- If set, the MAP\_TO\_PIN bit will map the external interrupt source to Interrupt Pending bits in the CP0 Cause register of the local processor. The actual Interrupt Pending value is set in the MAP field of this register.
  - Note that in EIC mode, the MAP Field of this register contains the encoded value of the number (0 -63). For example, a value of 0x20 asserts Interrupt 32 (decimal). For vectored interrupt mode, only values of 0x0 through 0x5 should be used.
- If set, the MAP\_TO\_NMI bit will map the external interrupt source to the NMI bit in the CP0 Status register. This in essence will cause the processor to soft boot using the boot exception vector as the start of the interrupt routine.
- MAP\_TO\_YQ bit is only present for a multi threaded core. If set it determines that the source is a Yield Qualifier. The actual Yield Qualifier setting is set in the MAP field of the register.

### 9.4.3.3 Mapping an Interrupt Source to a Register Set

Each processor has one register per interrupt source used when the processor is in "EIC mode" to map the interrupt source to a register set. This is the "EIC Shadow Set Register", GIC\_VPEi\_EICSS, located in the GIC local and other sections. Refer to [Section 9.6.3.6 "Local EIC Shadow Set Registers"](#) for more information.

The first register corresponds to interrupt source 0; the second to interrupt source 1 and so on. The EIC\_SS field is set to the register set number.

## 9.4.4 Enabling, Disabling, and Polling Interrupts

The Enabling, Disabling and Polling of interrupts is configured through several registers in the GIC that are shared by all processors.

There are 4 shared registers groups for Enabling, Disabling and Polling of interrupts.

- Enabling an interrupt using the "GIC Set Mask Registers", GIC\_SH\_SMASK
- Disabling an interrupt using the "GIC Reset Mask Registers", GIC\_SH\_RMASK
- Determining the Enable/Disable state of an interrupt state using "GIC Mask Register", GIC\_SH\_MASK
- Polling the interrupt active state using the "GIC Pending Register", GIC\_PEND\_MASK

Like the trigger registers, each interrupt source is represented by one bit in each register group. Each register in a group is 32 bits so each controls 32 interrupt sources. The first register in each group would control interrupts sources 0 - 31, the next 32 - 63 and so on. Since there can be 256 interrupt sources there could be 8 registers in each group. There are enough of these registers in each group to control the number of interrupt sources implemented. The number of interrupt sources is a fixed value configured at core build time. This number can be determined by reading the

NUMINTERRUPTS field of the "GIC Configuration Register", GIC\_SH\_CONFIG. Refer to [Section 9.5.3.1 “Global Config Register”](#) for more information.

#### 9.4.4.1 Enabling External Interrupts

The GIC Set Mask register group is used to enable external interrupts. It is made up of "GIC Set Mask Registers", GIC\_SH\_SMASK. For synchronization purposes this is a write only register. Setting the source bit enables the interrupt. Refer to [Section 9.5.3.10 “Global Interrupt Set Mask Registers”](#) for more information.

#### 9.4.4.2 Disabling External Interrupts

The GIC Reset Mask register group is used to disable external interrupts. It is made up of "GIC reset Mask Registers", GIC\_SH\_RMASK. For synchronization purposes; this is a write only register. Setting the source bit disables the interrupt. Refer to [Section 9.5.3.9 “Global Interrupt Reset Mask Registers”](#) for more information.

#### 9.4.4.3 Determining the Enabled or Disabled Interrupt State

The GIC Mask register group is used to determine if an external interrupt is enabled. It is made up of GIC Mask Registers, GIC\_SH\_MASK. For synchronization purposes; this is a read only register. If a bit is set the corresponding interrupt source is enable. If it is clear the corresponding interrupt is disabled. Refer to [Section 9.5.3.11 “Global Interrupt Mask Registers”](#) for more information.

#### 9.4.4.4 Polling for an Active Interrupt

The GIC Pending register group is used to determine if a external interrupt is active. It is made up of GIC Pending Registers, GIC\_PENDING\_MASK. This is a read only register. If a bit is set the corresponding interrupt source is active. If it is clear the corresponding interrupt is inactive. Refer to [Section 9.5.3.12 “Global Interrupt Pending Registers”](#) for more information.

#### 9.4.4.5 Programming Example

Incoming interrupts are registered in the *Global Interrupt Pending* registers (*GIC\_SH\_PENDING<sub>x,y</sub>*). This is the register that software needs to probe to discern the source of the interrupt. The *Global Interrupt Mask* registers (*GIC\_SH\_MASK<sub>x,y</sub>*) allow software to temporarily disable any particular interrupt source.

There are separate set (*GIC\_SH\_SMASK<sub>x,y</sub>*) and reset (*GIC\_SH\_RMASK<sub>x,y</sub>*) mask registers to set/clear individual interrupts to avoid any read-modify-write hazards within the system (multiple VPEs reading/writing the mask register simultaneously). This mechanism is shown in [Figure 9.2](#) for interrupts 31:0. For interrupts 64:32, a different set of registers is used. Similar for interrupts 95:64, and so on through interrupts 255:224.

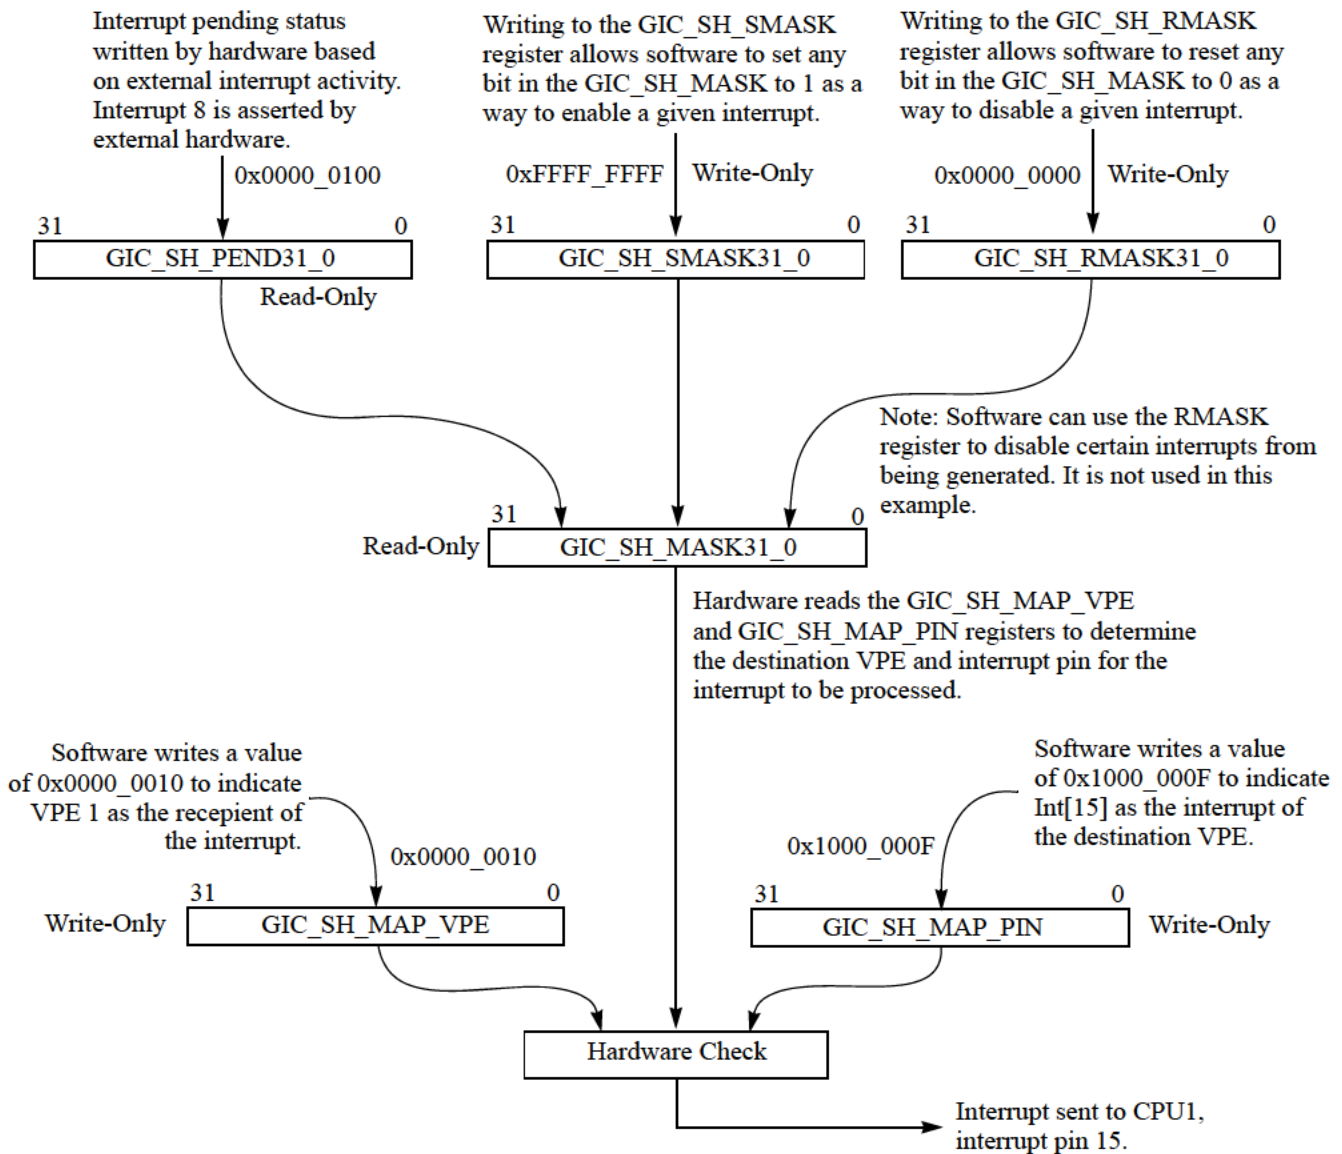
When an interrupt occurs, the corresponding bit in the *GIC\_SH\_PENDING* register is set by hardware. If the corresponding interrupt enable bit in the *GIC\_SH\_MASK* bit is set, the GIC delivers the interrupt to the appropriate VPE. The hardware does this by using the *GIC\_SH\_MAP\_VPE* register to send the interrupt to the appropriate VPE and the *GIC\_SH\_MAP\_PIN* register to set the interrupt pins for that VPE.

In the following example:

- External interrupt 8 is asserted
- All bits of the *GIC\_SH\_SMASK* register are set, enabling all 32 interrupts.
- The receiving VPE is #1, and the receiving interrupt is #15.

This example is shown in [Figure 9.2](#) below.

**Figure 9.2 Masking and Mapping of Interrupts in the GIC**



### 9.4.5 Inter-processor Interrupts

Each processor in the system can interrupt any other processor. Each inter-processor interrupt is configured just like an external interrupt using sources not being used by external devices. The interrupt source must be configured to be edge sensitive.

The "Global Interrupt Write Edge Register", GIC\_SH\_WEDGE is a shared register used to deliver an interrupt to another processor (only one per system). It is also used to clear an interrupt. There are two fields in the GIC\_SH\_WEDGE register used to do this.



- The RW bit determines if the interrupt is being set (delivered) or cleared. Setting this bit delivers an interrupt and clearing the bit clears the interrupt.
- The Interrupt field should be set to the interrupt number to be set or cleared.

### 9.4.5.1 WEDGE Register Programming Example

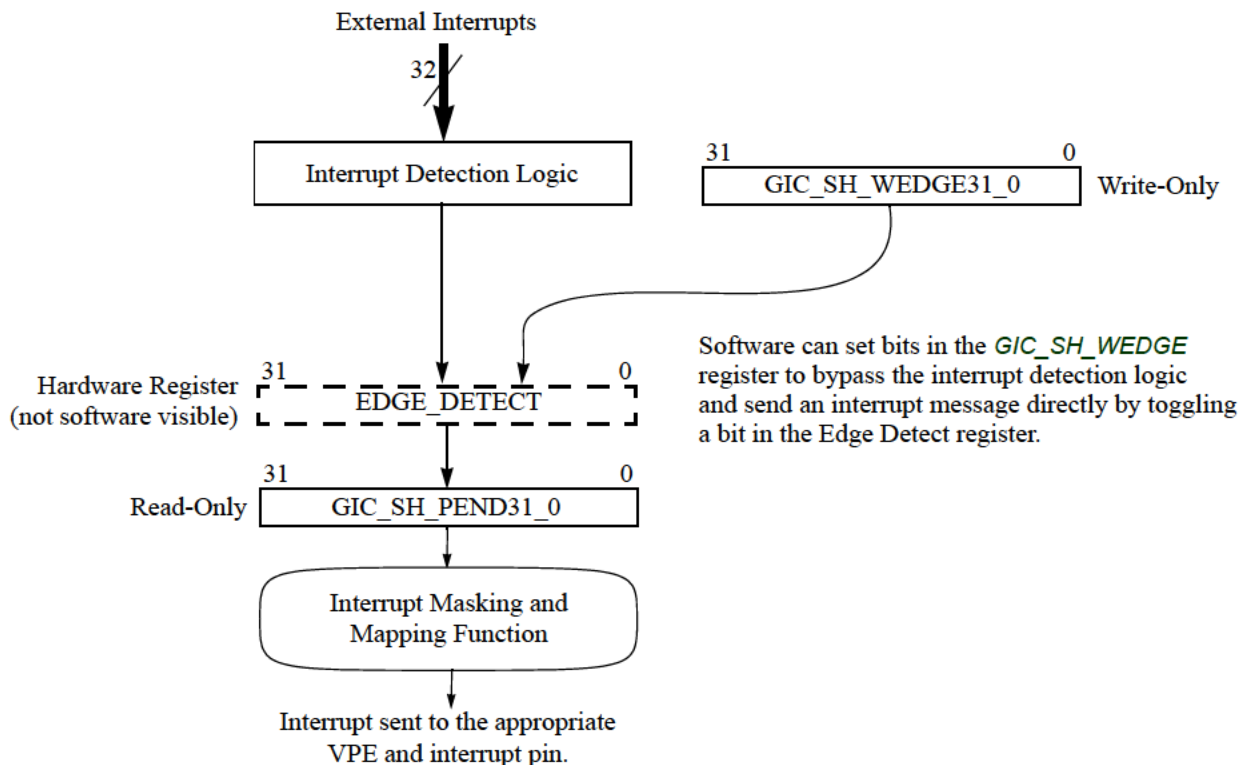
Setting a bit in the *Write Edge* register is treated equivalently to having the edge detection logic see an active edge. Because the programming of the Write Edge register has a direct effect on the state of the internal Edge Detect register, the *Write Edge* register can be used to bypass the edge detection logic. Thus, it does not matter whether the corresponding interrupt is configured to be rising, falling, or dual edge sensitive.

When VPE 0 wants to interrupt VPE 1, the number of the interrupt to be used is programmed into the GIC\_SH\_WEDGE31\_0 register. The selected interrupt must be mapped to the target VPE (VPE1 in this example) using the GIC\_SH\_MAPi\_VPE register).

For example, assume VPE 0 wants to toggle interrupt 40. In this case, software writes a value of 0x28 into the GIC\_SH\_WEDGE31\_0 register. Hardware then writes the value in the WEDGE register into the Edge Detect hardware register, effectively bypassing the edge detection logic. Hardware determines that interrupt being toggled belongs to VPE 1, not VPE 0. The GIC routing logic then routes interrupt 40 onto the appropriate VPE 1 interrupt pins.

Figure 9.3 shows how the *Write Edge* register can be used to bypass the interrupt detection logic and assert interrupt directly. Setting a bit in the *Write Edge* register in turn sets the corresponding bit in the internal Edge Detect register, forcing an interrupt to be generated and allowing for inter-processor interrupts within the GIC.

**Figure 9.3 Sending Inter-Processor Interrupts in the GIC**





### 9.4.5.2 Inter-Processor Interrupt Code Example

Here is an example on how to set up interrupt sources 32 through 39 for inter-processor interrupts. First here is a table of what the #defines are set to.

**Table 9.8 Setting Interrupt Sources 32 Through 39**

#define	Value	Description
GIC_BASE_ADDR	0xbbdc0000	Virtual Base memory address of the GIC memory mapped registers
GIC_P_BASE_ADDR	0x1bdc0000	Physical Base address of the GIC memory mapped registers
GIC_SH_RMASK63_32	0x0304	Offset into the GIC registers for the GIC Reset Mask Register
GIC_SH_POL63_32	0x0104	Offset into the GIC registers for the GIC Reset Polarity Register
GIC_SH_TRIG63_32	0x0184	Offset into the GIC registers for the GIC Trigger Register
GIC_SH_SMASK63_32	0x0384	Offset into the GIC registers for the GIC Set Mask Register
GCR_CONFIG_ADDR	0xbfbf8000	Base address of the Global Configuration Register
GCR_GIC_BASE	0x0080	Offset int the GCR of the GIC base Address
GIC_SH_MAP0_VPE31_0	0x2000	Offset into the GIC for first map register
GIC_SH_MAP_SPACER	0x20	Spacing between map registers

```

// First load GIC base address into the GCR and enable the GIC
li    a1, GCR_CONFIG_ADDR + GCR_GIC_BASE // load the address of the GIC Base Address register
li    a0, (GIC_P_BASE_ADDR | 1)          // Physical address + enable
sw    a0, 0(a1)                          // Store the Physical address of the GIC and the enable
                                           // bit to the GCR

// Configure the source pins for inter-processor interrupts
li    a1, GIC_BASE_ADDR                  // load GIC base address
li    a0, 0xff                           // load bits for interrupts 32..39 lower 8 bits of 2nd group)
sw    a0, GIC_SH_RMASK63_32(a1)         // (disable interrupts 32..39)
sw    a0, GIC_SH_TRIG63_32(a1)         // (set source to be edge sensitive for interrupts 32..39)
sw    a0, GIC_SH_POL63_32(a1)          // (set Polarity to rising edge for interrupts32..39)
sw    a0, GIC_SH_SMASK63_32(a1) // (enable interrupts 32..39)

// Map interrupts to a processor

// The register offset into the GIC for the MAP TO VPE register is obtained by multiplying the
// interrupt number by the spacing size (GIC_SH_MAP_SPACER) and adding the offset for the Global
// Interrupt Map to VPE Registers (GIC_SH_MAP0_VPE31_0).

li    a0, 1 // set bit 0 processor 0
// Map Source 32 processor 0

```

```

sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 32)(a1)
sll a0, a0, 1 // set bit 1 for processor 1
// Source 33 to processor 1
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 33)(a1)
sll a0, a0, 1 // set bit 2 for processor 2
// Source 34 to processor 2
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 34)(a1)
sll a0, a0, 1 // set bit 3 for processor 3 or for MT vpe3
// Source 35 to processor 3
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 35)(a1)
sll a0, a0, 1 // set bit 4 for processor 4
// Source 36 to processor 4
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 36)(a1)
sll a0, a0, 1 // set bit 5 for processor 5
// Source 37 to processor 5
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 37)(a1)
sll a0, a0, 1 // set bit 6 for processor 6
// Source 38 to processor 6
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 38)(a1)
sll a0, a0, 1 // set bit 7 for processor 7
// Source 39 to processor 7
sw a0,GIC_SH_MAP0_VPE31_0+(GIC_SH_MAP_SPACER * 39)(a1)

```

At this point the Map-to-Pin Registers could be used to map each interrupt source to Interrupt Pending bits in the CP0 Cause register of a processor. The default values for the "Map to Pin" registers are the MAP\_TO\_PIN bit is set and the MAP field is cleared. This example does not change the default values therefore the interrupts are mapped to IP2, Hardware Interrupt 0.

#### 9.4.5.3 Example of Sending an Inter-Processor Interrupt

The following is a C coding example of sending an inter-processor interrupt. First the #defines

**Table 9.9**

#define	Value	Description
GIC_SH_WEDGE	*((volatile unsigned int*) (0xbbdc0280))	Address of the GIC_WEDGE_REGISTER.
FIRST_IPI	32	Source number for the first IPI.

```

void set_ipi(int cpu_num) {
// Add the enable bit, the first IPI number and the cpu number
// and write it to the GIC_SH_WEDGE register
    GIC_SH_WEDGE = 0x80000000 + FIRST_IPI + cpu_num ;

```

#### 9.4.5.4 Example of Clearing an Inter-Processor Interrupt

Once received, the interrupt routine should do whatever action is intended for the interrupt and clear the interrupt by writing the interrupt number to the GIC\_SH\_WEDGE register before executing the eret instruction. NOTE: only the interrupt number is set before the write so the R/W bit will be cleared indicating that the interrupt is to be cleared.

```

li    k0, (GIC_SH_WEDGE | GIC_BASE_ADDR)
mfc0  k1, C0_EBASE           // Get cp0 EBase
ext   k1, k1, 0, 10         // Extract CPUNum
addiu k1, 0x20              // Offset to base of IPI interrupts.
sw    k1, 0(k0)            // Clear this IPI.

```

### 9.4.6 Local Device Interrupt Configuration

The GIC also controls how devices within the processor and the GIC are configured and mapped locally to the processor.

There are 2 devices that are added as part of the GIC described in this section:

- GIC Interval Timer - a 64 bit timer that compares a local compare registers, GIC\_VPE\_CompareLo/Hi of a processor with a global counter, GIC\_SH\_CounterLo/Hi in the GIC and activates an interrupt when they match.
- GIC Watchdog Timer - a 32 bit decrementing counter, GIC\_VPE\_WD\_COUNT that can be used as liveness signal for a processor.

#### 9.4.6.1 GIC Interval Timer

The interval timer is similar to the CP0 Count/Compare timer within each processor. The difference is the GIC CounterLo/Hi register is global to the CPS so all processors will have the same time reference.

Both the interval count and interval compare values are 8 bytes wide and are made up of 2 (Lo/Hi) registers. For each Lo register overflow the Hi register is incremented. If the Hi register overflows, both registers rollover to 0.

#### **Counter Registers**

The counter registers, GIC\_SH\_CounterLo/Hi are in the shared section of the GIC memory map. The counter must be stopped before it is set. This is done by setting the COUNTSTOP bit of the GIC\_SH\_CONFIG register (link to register reference of GIC\_SH\_CONFIG). In practical use the counter is usually set by an OS at boot time by one processor. These counter registers are also available (read only) in user mode located at offset 0 of the User Mode Visible Section of the GIC.

The COUNTBITS field of the GIC\_SH\_CONFIG register in [Section 9.5.3.1, "Global Config Register"](#) is used to set up the width of the GIC\_SH\_CounterHi register. In the GIC design, this field is fixed at a value of 0x8, indicating a total counter size of 64-bits.

The shared counter registers are defined as follows:

- *GIC\_SH\_CounterLo* register in [Section 9.5.3.2, "GIC CounterLo"](#). Used in conjunction with the *GIC\_SH\_CounterHi* register. Sets the lower 32-bits of the starting count value.
- *GIC\_SH\_CounterHi* register in [Section 9.5.3.3, "GIC CounterHi"](#). Used in conjunction with the *GIC\_SH\_CounterLo* register. Sets the upper 32-bits of the starting count value.

### **Compare Registers**

The compare registers, *GIC\_VPE\_CompareLo/Hi* are located in the local section of the GIC memory map making the count specific to each processor. These registers can be written at any time. When the count value equals the compare value an Interval Timer interrupt is asserted. The interrupt is cleared (de-asserted) by writing to either *GIC\_VPE\_CompareLo/Hi* register. The compare registers are defined as follows:

- *GIC\_VPEi\_CompareLo* register in [Section 9.6.3.4, "CompareLo Register"](#). Used in conjunction with the *GIC\_VPEi\_CompareHi* register to set the count value at which an internal interrupt is generated.
- *GIC\_VPEi\_CompareHi* register in [Section 9.6.3.5, "VPE-Local CompareHi Register"](#). Used in conjunction with the *GIC\_VPEi\_CompareLo* register to set the count value at which an internal interrupt is generated.

### **Determining the Counter Width**

The counter used for GIC internal interrupt generation has a minimum width of 32 bits, meaning that all of the *GIC\_SH\_CounterLo* register is used. In the GIC design, the width of the *GIC\_SH\_CounterHi* register is also fixed at 32 bits as indicated by a value of 0x8 in the 4-bit COUNTBITS field in the *GIC\_SH\_CONFIG* register. To derive the total width of the counter, the following formula is used:

$$32 + \text{COUNTBITS} \times 4$$

Where:

'32' is the width of the *GIC\_SH\_CounterLo* register and 'COUNTBITS' is the value in the COUNTBITS field of the *GIC\_SH\_CONFIG* register.

Since the COUNTBITS field contains a fixed value of 0x8, the overall width of the counter would be:

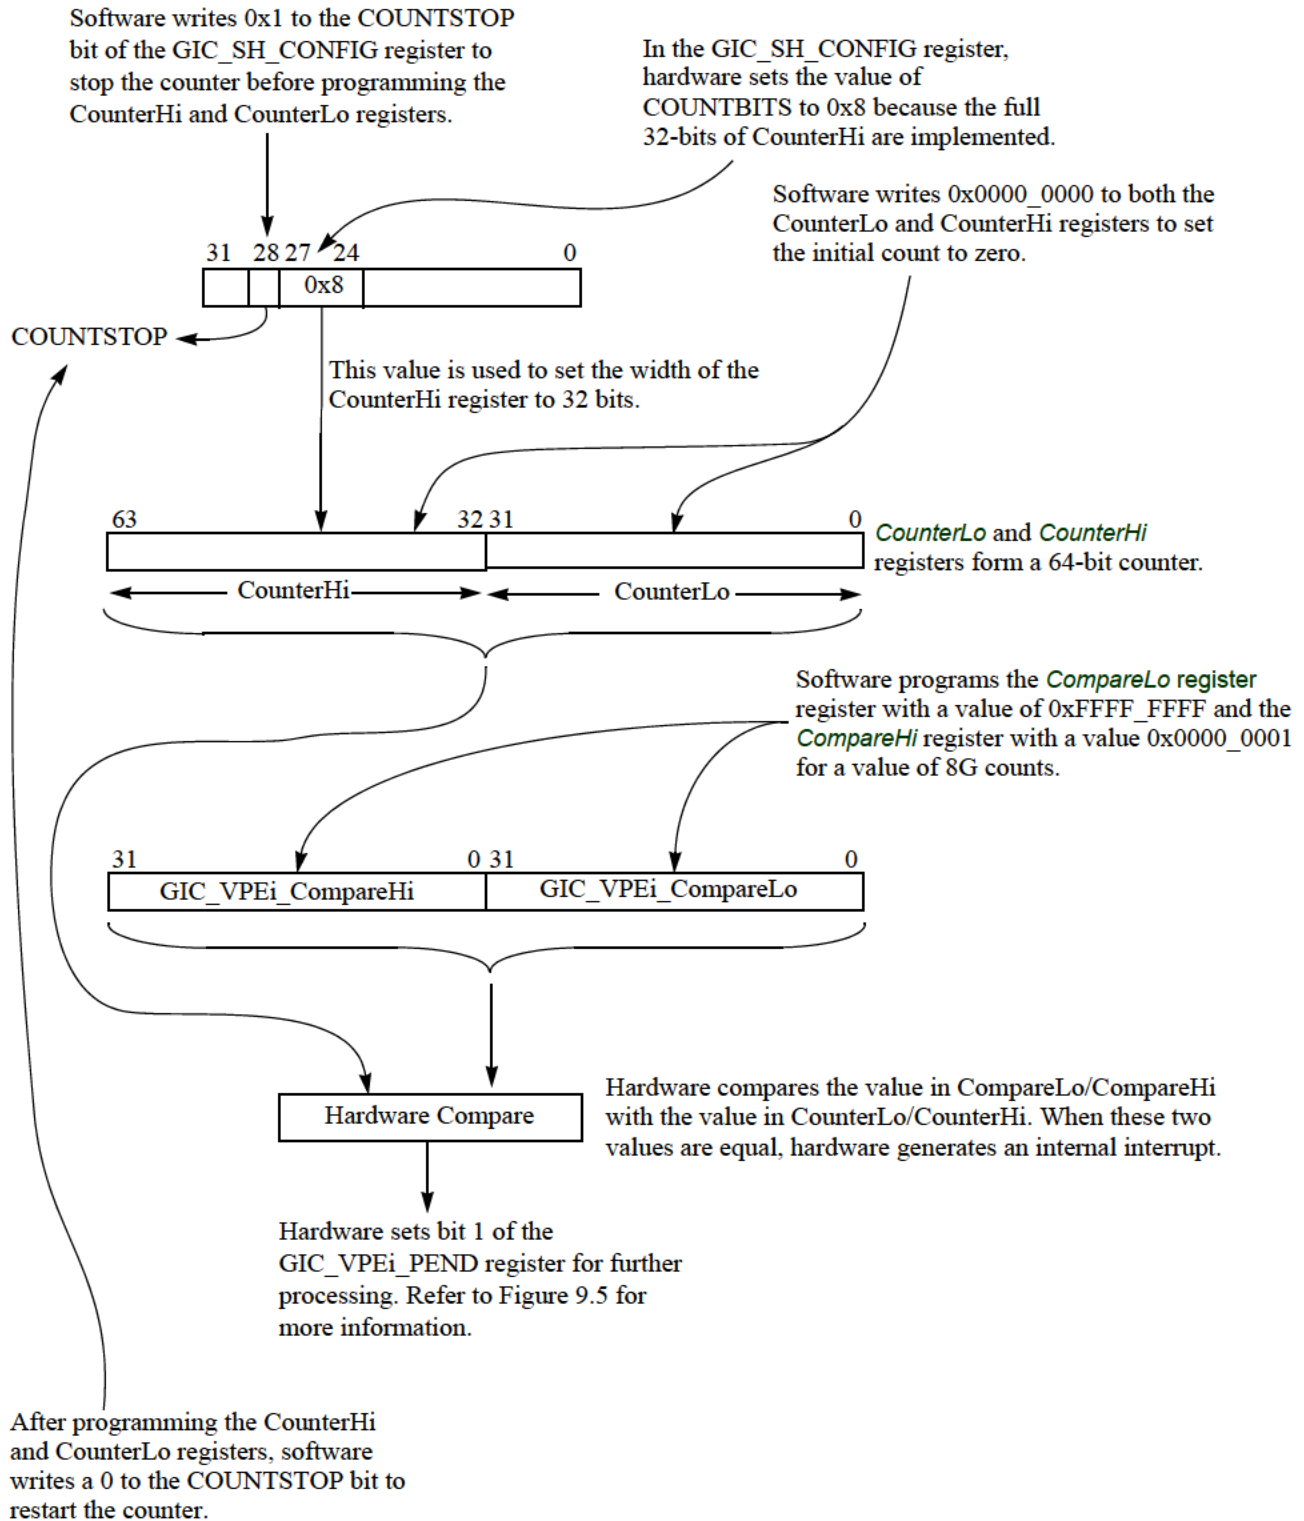
$$32 + 8 \times 4 = 64 \text{ bits}$$

In the GIC design, the COUNTBITS field is fixed at a value of 0x8, indicating a total counter size of 64-bits.

### **Counter Based Interrupt Example**

In the example shown in [Figure 9.4](#), the width of the counter is 64-bits, and the CompareLo/Hi value is 0x1\_FFFF\_FFFF which corresponds to 8G clock cycles. When this count is reached, hardware generates an internal interrupt.

**Figure 9.4 Example of GIC Internal Counter-Based Interrupt Generation**



### 9.4.6.2 GIC Watchdog Timer

Each core supports a Watchdog timer that is controlled by the following three registers.

- The "GIC Watchdog Timer Configuration Register", GIC\_COREi\_WD\_CONFIG is local to each processor and reports state information and configures the characteristics of the timer.
- The "Watchdog Timer Initial Count Register", GIC\_COREi\_WD\_INITIAL is local to each processor and is used to set the timer interval.
- The "Watchdog Timer Count Register", GIC\_VPEi\_WD\_COUNT is a read only register local to each processor that contains the current value of the countdown.

#### **GIC Watchdog Timer Configuration Register**

The GIC Watchdog Timer Configuration register contains bits that control the function of the timer.

- Clearing the WAIT bit of GIC\_COREi\_WD\_CONFIG register (default value) will cause the counter stop counting when the processor is executing a wait instruction or is in a low power state controlled by the Cluster Power Controller. Setting this bit to 1 will cause it to continue counting down in these states. Usually this bit is left unset.
- Clearing the Debug bit (default value) will cause the counter to stop the count when the processor enters debug mode. When set the count will continue counting down. Usually this bit is left unset.
- The TYPE field in bits 3:1 of this register determines what happens when the timer reaches 0.

**Table 9.10 GIC Watchdog Timer Modes**

Encoding	Mode	Behavior
0x2	One Trip	An interrupt is asserted and the timer stops.
0x1	Second Countdown	An interrupt is asserted and the timer reloads. If the timer expires for the second time before being reloaded again all processors in the CPS will be reset. This mode provides a way to distinguish between a Software hang and a Hardware Hang. Usually the Watchdog Timer Interrupt is routed to NMI. This will cause the processor to soft reboot. In this mode that is what happens when the timer expires the first time so if this was a software hang during the reboot the software should reload the Watchdog Timer thus avoiding the second expiration. If the processor itself does not respond to the interrupt then it is assumed to be a hardware issue so when the count expires the second time a reset signal will be sent to all processors in the system.
0x3	Programmable Interval Timer	An interrupt is asserted, the initial count is reloaded and the time starts counting down again interrupting each time the counter reaches 0. This mode provides a per processor interval timer. This is one mode where the interrupt should not be routed to NMI. It should instead be routed to a normal interrupt where for example the interrupt could be used in a time slicing OS.

Clearing the WDEN bit disables the timer and when it is set it enables the timer. Writing WDEN with a 1 triggers a reload of the GIC\_VPE\_WD\_COUNT register with the value in the GIC\_COREi\_WD\_INITIAL register. Refer to [Section 9.6.3.1, "Watchdog Timer Config Register"](#) for more information.

## Watchdog Timer Initial Count Register

The "Watchdog Timer Initial Count Register", `GIC_COREi_WD_INITIAL` is local to each processor and is used to set the timer interval. To start the counter for the first time the counter should be disabled by clearing the `WDEN` bit in the `GIC_COREi_WD_CONFIG` register and the countdown value loaded into this register and then the counter enabled by setting the `WDEN` bit. Refer to Section 9.6.3.3, "Watchdog Timer Initial Count Register" for more information.

## Watchdog Timer Count Register

The "Watchdog Timer Count Register", `GIC_VPE_WD_COUNT` is a read only register local to each processor that contains the current value of the countdown. This register is reloaded with the value in the `GIC_COREi_WD_INITIAL` register each time the `WDEN` bit in the `GIC_COREi_WD_CONFIG` register is set. Refer to Section 9.6.3.2, "Watchdog Timer Count Register" for more information.

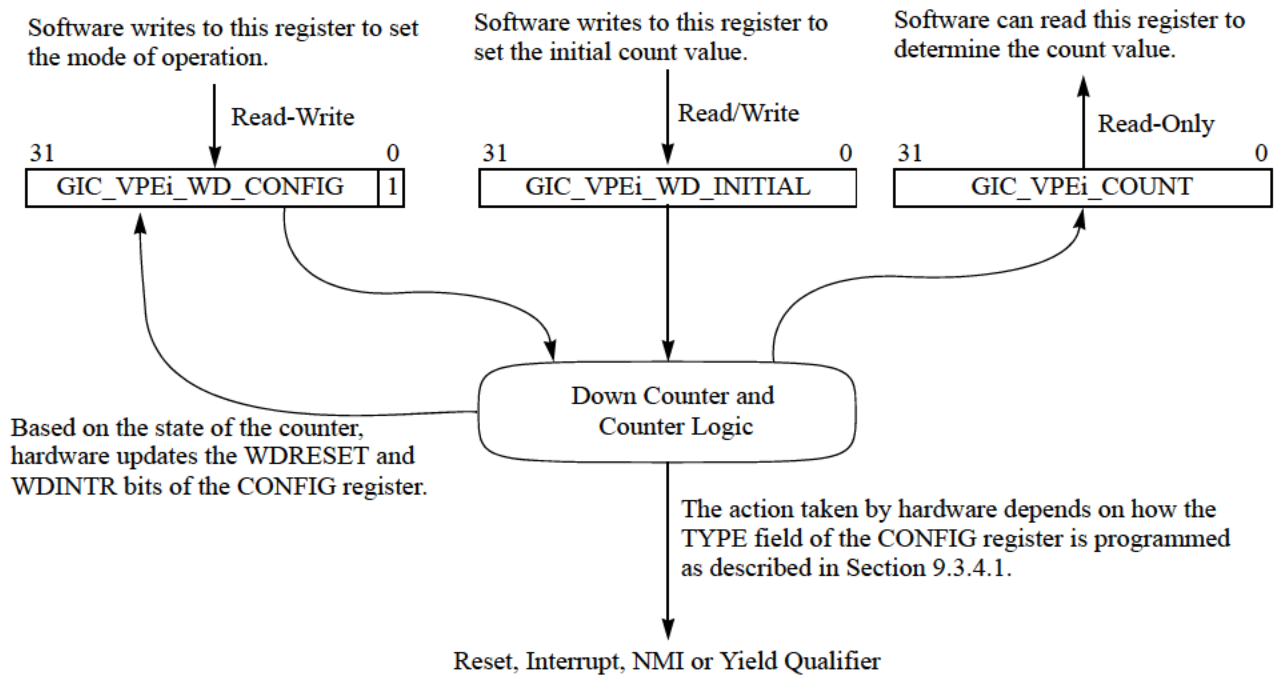
## Configuring the Watchdog Timer

Software can configure the WatchDog timer with a starting count value by programming the *WatchDog Timer Initial Count* register (`GIC_VPEi_WD_INITIAL`) located at offset address 0x0098. Refer to Section 9.6.3.3 "Watchdog Timer Initial Count Register" for more information.

Software can read the state of the count at any time by reading the the *WatchDog Timer Count* register (`GIC_VPEi_WD_COUNT`) located at offset address 0x0094. Refer to Section 9.6.3.2 "Watchdog Timer Count Register" for more information.

Figure 9.7 shows the timer counter configuration process.

Figure 9.5 Local Watchdog Timer Interrupt Count Configuration



## Watchdog Timer Masking and Mapping

Figure 9.5 above shows the process used to configure the Watchdog timer. Once a Watchdog timer interrupt is generated (output of Figure 9.5), hardware sets bit 0 of the *Local Interrupt Pending* register (*GIC\_VPEi\_PEND*) at offset address 0x0004. Hardware then reads the state of bit 0 in the *Local Interrupt Mask* register (*GIC\_VPEi\_MASK*) at offset address 0x0008 to determine whether the Watchdog timer interrupt has been masked. The *GIC\_VPEi\_MASK* register is a read-only register.

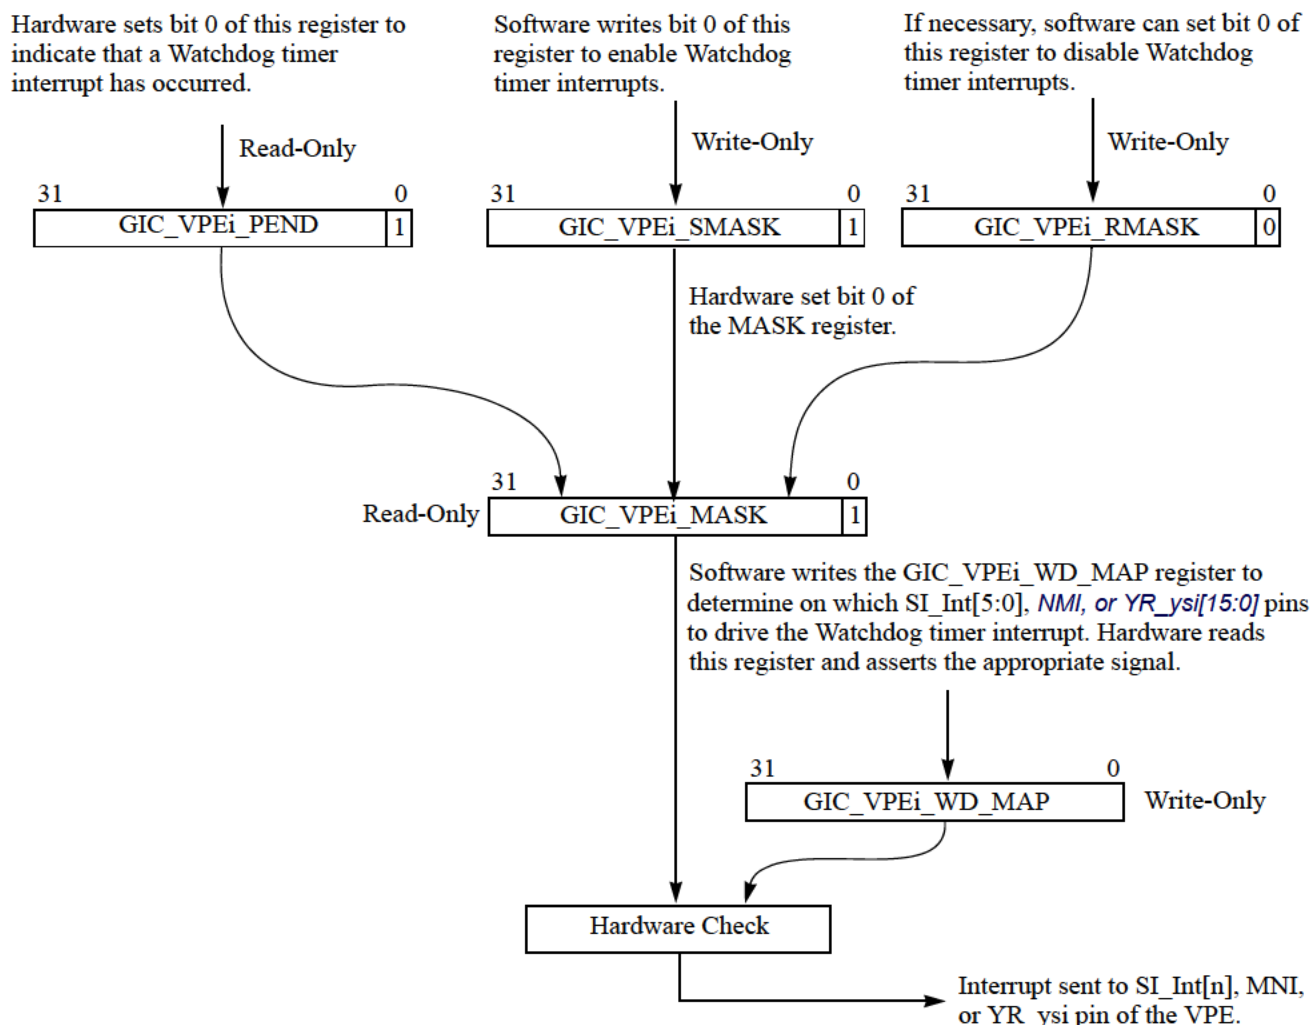
Software can affect the state of this register using the write-only *Local Interrupt Set Mask* register (*GIC\_VPEi\_SMASK*) at offset address 0x0010 and the *Local Interrupt Reset Mask* register (*GIC\_VPEi\_RMASK*) at offset address 0x000C. Software sets bit 0 of the *SMASK* register to enable the Watchdog timer interrupt, or it can set bit 0 of the *RMASK* register to disable Watchdog timer interrupts. Note that when the WatchDog timer is programmed to generate a hardware reset, the reset cannot be masked by the *Local Interrupt Mask* register

Once hardware has determine the masking characteristics of the interrupt, it uses the *Watchdog Timer Map-to-Pin* register at offset address 0x0040 to determine which *SI\_Int[5:0]*, *NMI*, or *YR\_ys[15:0]* pins the interrupt will be driven onto. In non-EIC mode, bits 5:0 of this register are used to select one of 6 VPE interrupts. For example, if software programs this field with a value of 0x2, then the Watchdog timer interrupt will be driven into *SI\_Int[2]*. In non-EIC mode, only encodings 0 - 5 are valid.

In EIC mode, the VPE encodes this field to support up to 64 interrupts. For example, if software programs this field with a value of 0x20, then the Watchdog timer interrupt corresponds to interrupt 33. This encoded value is then driven onto *SI\_Int[5:0]*.



**Figure 9.6 Watchdog Timer Interrupt Masking and Mapping in the GIC**



### Watchdog Timer and Debug Mode

Under certain conditions, software may want to suspend Watchdog timer operation while the interAptiv Multiprocessing System is in debug mode. This can be accomplished by clearing the `DEBUGMODE_CTRL` bit of the *Watchdog Timer Config* register located at offset address 0x0090. When this bit is cleared, counting is stopped. Note that the `DM` bit of the *CP0 Debug* register (`DEBUGDM`) must be set to place the device in debug mode.

If this bit is set by software, entering debug mode has no effect on the Watchdog timer counting process.

### Watchdog Timer and Low Power Mode

Under certain conditions, software may want to suspend Watchdog timer operation while the interAptiv Multiprocessing System is in low power mode. This can be accomplished by clearing the `WAITMODE_CTRL` bit of the *Watchdog Timer Config* register located at offset address 0x0090. When this bit is cleared, counting is stopped (including when low power mode is entered via the `WAIT` instruction).

If this bit is set by software, entering low power mode has no effect on the Watchdog timer counting process.

## 9.4.7 Local Interrupt Routing

### 9.4.7.1 Routability of Local Interrupts

Local interrupts (except for the Watchdog timer, GIC Interval Timer and software interrupts) can be hardwired to local pins when the CPS is configured or can be more flexible and left to software to route the local interrupts to local pins on the processor. The "Local Interrupt Control Register", `GIC_COREi_CTL` (link to register reference of `GIC_COREi_CTL`) reports the routable state of the local interrupts. If the bit for the particular interrupt is set then the interrupt is routable within the GIC. The following table describes the behavior if not set.

Bits 4:1 of the `GIC_VPEi_CTL` register determines the routing of the following interrupts. In the interAptiv GIC design, these bits are hard-wired to 1. Note that Software Interrupts from the VPE are routed internally by the CPU in vectored interrupt mode, and are only routed through the GIC when the GIC is in EIC mode, regardless of the `GIC_VPEi_CTL` register.

**Table 9.11 GIC\_COREi\_CTL Register Fields**

Bit Field Name	Behavior if cleared
FDC_ROUTABLE	The CPU Fast Debug Channel Interrupt is hardwired to one of the SI_Int pins as described by the CPU's COP0 IntCtlPFDCI register field.
SWINT_ROUTABLE	The CPU SW Interrupts are routed back to the CPU directly.
PERFCOUNT_ROUTABLE	The CPU Performance Counter Interrupt is hardwired to one of SI_Int pins as described by the CPU's COP0 IntCtlPPCI register field.
TIMER_ROUTABLE	The CPU Timer Interrupt is hardwired to one of the SI_Int pins, as described by the CPU's COP0 IntCtlIPTI register field

### 9.4.7.2 Routing Local Interrupts

If a local interrupt is routable it can be routed to a local signal of the local processor, much the same as an external interrupt.

There is a Local Interrupt Map to Pin Register (link to register reference of Local WatchDog Timer/Compare/CPU Timer/PerfCount/SWInt0-1 Map to Pin Registers) for each local interrupt source that further maps the local interrupt to a specific input on the processor. There are two bits, `MAP_TO_PIN` and `MAP_TO_NMI` that control the type of input that is assigned to the interrupt source. Only one of these bits can be set at any one time.

- If set the `MAP_TO_PIN` bit will map the local interrupt source to Interrupt Pending bits in the CP0 Cause register of the processor. The actual Interrupt Pending bit is set in the MAP field of this register. The MAP Field of this register contains the encoded value of the number (0 -63). For example, a value of 0x20 asserts Interrupt 32 (decimal). For vectored interrupt mode, only use values of 0x0 to 0x5.
- If set the bit will map the local interrupt source to the NMI bit in the CP0 Status register. This in essence will cause the processor to soft boot using the boot exception vector as the start of the interrupt routine.

Each of these interrupt types is described in the following subsections. [Table 9.12](#) lists the registers and associated bits that would be programmed to facilitate each type of interrupt listed above.

**Table 9.12 Local Interrupt Masking and Mapping Register Usage Per Interrupt Type**

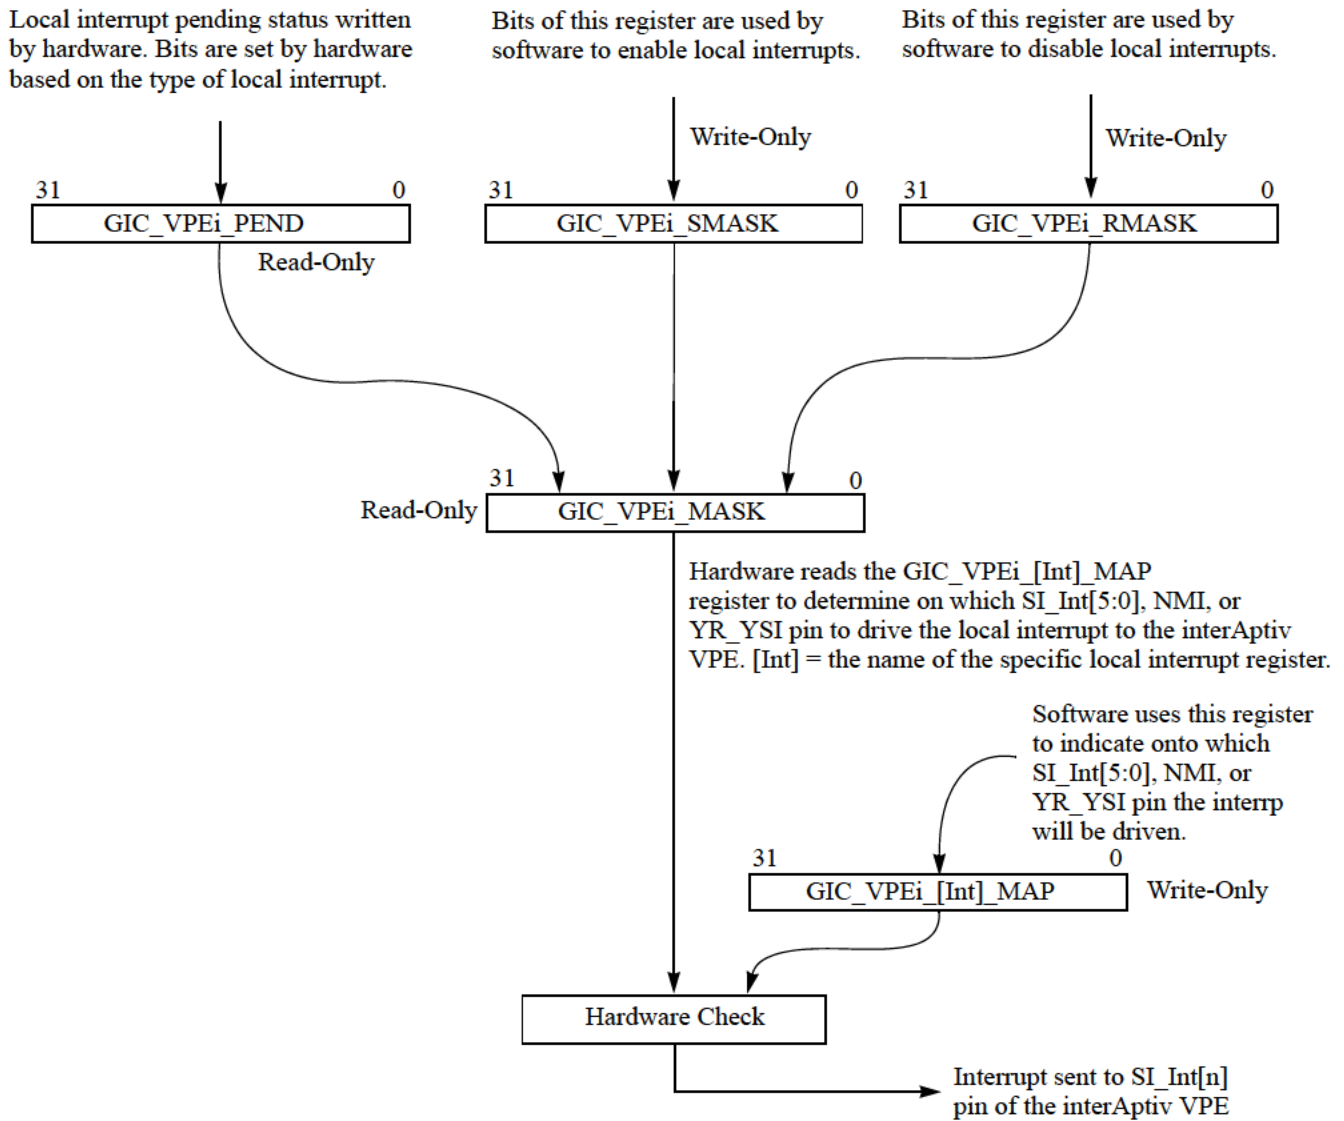
Interrupt	Register Name	Offset	Bits Used	Function
WatchDog	GIC_VPEi_PEND	0x0004	0	Set by hardware on a local WatchDog timer interrupt.
	GIC_VPEi_MASK	0x0008	0	Set by hardware based on the state of bit 0 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	0	Used by software to disable WatchDog timer interrupts.
	GIC_VPEi_SMASK	0x0010	0	Used by software to enable WatchDog timer interrupts.
	GIC_VPEi_WD_MAP	0x0040	31, 5:0	Used by software to map the WatchDog timer interrupt to one of the SI_Int[5:0] pins of the interAptiv VPE.
Count and Compare	GIC_VPEi_PEND	0x0004	1	Set by hardware on a local Count/Compare interrupt.
	GIC_VPEi_MASK	0x0008	1	Set by hardware based on the state of bit 1 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	1	Used by software to disable Count/Compare interrupts.
	GIC_VPEi_SMASK	0x0010	1	Used by software to enable Count/Compare interrupts.
	GIC_VPEi_COMPARE_MAP	0x0044	31, 5:0	Used by software to map the Count/Compare interrupt to one of the SI_Int[5:0] pins of the interAptiv VPE.
Timer	GIC_VPEi_PEND	0x0004	2	Set by hardware on a local timer interrupt.
	GIC_VPEi_MASK	0x0008	2	Set by hardware based on the state of bit 2 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	2	Used by software to disable timer interrupts.
	GIC_VPEi_SMASK	0x0010	2	Used by software to enable timer interrupts.
	GIC_VPEi_TIMER_MAP	0x0048	31, 5:0	Used by software to map the timer interrupt to one of the SI_Int[5:0] pins of the interAptiv VPE.
Performance Counter	GIC_VPEi_PEND	0x0004	3	Set by hardware on a performance counter interrupt.
	GIC_VPEi_MASK	0x0008	3	Set by hardware based on the state of bit 3 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	3	Used by software to disable performance counter interrupts.
	GIC_VPEi_SMASK	0x0010	3	Used by software to enable performance counter interrupts.
	GIC_VPEi_PERFCTR_MAP	0x0050	31, 5:0	Used by software to map the performance counter interrupt to one of the SI_Int[5:0] pins of the interAptiv VPE.
Software Interrupt 0	GIC_VPEi_PEND	0x0004	4	Set by hardware on a software interrupt 0 occurrence.
	GIC_VPEi_MASK	0x0008	4	Set by hardware based on the state of bit 4 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	4	Used by software to disable software interrupt 0 interrupts.
	GIC_VPEi_SMASK	0x0010	4	Used by software to enable software interrupt 0 interrupts.
	GIC_VPEi_SWInt0_MAP	0x0054	31, 5:0	Used by software to map software interrupt 0 to one of the SI_Int[5:0] pins of the interAptiv VPE.

**Table 9.12 Local Interrupt Masking and Mapping Register Usage Per Interrupt Type (continued)**

<b>Interrupt</b>	<b>Register Name</b>	<b>Offset</b>	<b>Bits Used</b>	<b>Function</b>
Software Interrupt 1	GIC_VPEi_PEND	0x0004	5	Set by hardware on a software interrupt 1 occurrence.
	GIC_VPEi_MASK	0x0008	5	Set by hardware based on the state of bit 5 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	5	Used by software to disable software interrupt 1 interrupts.
	GIC_VPEi_SMASK	0x0010	5	Used by software to enable software interrupt 1 interrupts.
	GIC_VPEi_SWInt1_MAP	0x0058	31, 5:0	Used by software to map software interrupt 1 to one of the SI_Int[5:0] pins of the interAptiv VPE.
Fast Debug Channel	GIC_VPEi_PEND	0x0004	6	Set by hardware on a Fast Debug Channel (FDC) interrupt.
	GIC_VPEi_MASK	0x0008	6	Set by hardware based on the state of bit 6 of the SMASK and RMASK registers. Used to determine whether the interrupt will be processed or ignored.
	GIC_VPEi_RMASK	0x000C	6	Used by software to disable FDC interrupts.
	GIC_VPEi_SMASK	0x0010	6	Used by software to enable FDC interrupts.
	GIC_VPEi_FDC_MAP	0x004C	31, 5:0	Used by software to map the FDC interrupt to one of the SI_Int[5:0] pins of the interAptivVPE.

The general overview of the local interrupt pending, masking, and mapping process is shown in [Figure 9.7](#).

**Figure 9.7 Local Interrupt Masking and Mapping in the GIC**



Each of the registers listed in [Figure 9.7](#) above can be found in the following sections:

- [Section 9.6.2.2 “Local Interrupt Pending Register”](#)
- [Section 9.6.2.5 “Local Interrupt Set Mask Register”](#)
- [Section 9.6.2.4 “Local Interrupt Reset Mask Register”](#)
- [Section 9.6.2.3 “Local Interrupt Mask Register”](#)
- [Section 9.6.2.6 “Local Map to Pin Registers”](#)

### 9.4.7.3 Watchdog Timer Interrupts

For more information, refer to [Section 9.4.6.2, “GIC Watchdog Timer”](#).

#### 9.4.7.4 Count and Compare Interrupts

A count and compare interrupt occurs when the contents of the of *GIC\_VPEi\_CompareLo* and *GIC\_VPEi\_CompareHi* registers match the contents of *GIC\_SH\_CounterLo* and *GIC\_SH\_CounterHi*, the Count/Compare interrupt is triggered. Refer to [Section “Counter Based Interrupt Example”](#) for more information.

When a count and compare interrupt is generated, hardware sets bit 1 of the *Local Interrupt Pending* register (*GIC\_VPEi\_PEND*) at offset address 0x0004. Hardware then reads the state of bit 1 in the *Local Interrupt Mask* register (*GIC\_VPEi\_MASK*) at offset address 0x0008 to determine whether the count and compare interrupt has been masked. The *GIC\_VPEi\_MASK* register is a read-only register.

Software can affect the state of this register using the write-only *Local Interrupt Set Mask* register (*GIC\_VPEi\_SMASK*) at offset address 0x0010 and the *Local Interrupt Reset Mask* register (*GIC\_VPEi\_RMASK*) at offset address 0x000C. Software sets bit 1 of the *SMASK* register to enable the count and compare interrupt, or it can set bit 1 of the *RMASK* register to disable count and compare interrupts.

Once hardware has determined the masking characteristics of the interrupt, it uses the *Count/Compare Map-to-Pin* register at offset address 0x0044 to determine which *SI\_Int[5:0]*, *NMI*, or *YR\_ysif[15:0]* pins the interrupt will be driven onto. In vectored interrupt mode, bits 5:0 of this register are used to select one of 6 VPE interrupts. In this mode, only encodings 0 - 5 are valid. In EIC mode, the VPE encodes this field to support up to 63 interrupts. For example, if software programs this field with a value of 0x20, then the WatchDog timer interrupt corresponds to interrupt level 32. This encoded value is then driven onto *SI\_Int[5:0]*.

#### 9.4.7.5 Timer Interrupts

When a timer interrupt is generated, hardware sets bit 2 of the *Local Interrupt Pending* register (*GIC\_VPEi\_PEND*) at offset address 0x0004. Hardware then reads the state of bit 2 in the *Local Interrupt Mask* register (*GIC\_VPEi\_MASK*) at offset address 0x0008 to determine whether the timer interrupt has been masked. The *GIC\_VPEi\_MASK* register is a read-only register.

Software can affect the state of this register using the write-only *Local Interrupt Set Mask* register (*GIC\_VPEi\_SMASK*) at offset address 0x0010 and the *Local Interrupt Reset Mask* register (*GIC\_VPEi\_RMASK*) at offset address 0x000C. Software sets bit 2 of the *SMASK* register to enable the timer interrupt, or it can set bit 2 of the *RMASK* register to disable timer interrupts.

Once hardware has determine the masking characteristics of the interrupt, it uses the *Timer Map-to-Pin* register at offset address 0x0048 to determine which *SI\_Int[5:0]*, *NMI*, or *YR\_ysif[15:0]* pins the interrupt will be driven onto. In non-EIC mode, bits 5:0 of this register are used to select one of 6 VPE interrupts. In non-EIC mode, only encodings 0 - 5 are valid. In EIC mode, the VPE encodes this field to support up to 63 interrupts.

#### 9.4.7.6 Performance Counter Interrupts

When a timer interrupt is generated, hardware sets bit 3 of the *Local Interrupt Pending* register (*GIC\_VPEi\_PEND*) at offset address 0x0004. Hardware then reads the state of bit 3 in the *Local Interrupt Mask* register (*GIC\_VPEi\_MASK*) at offset address 0x0008 to determine whether the performance counter interrupt has been masked. The *GIC\_VPEi\_MASK* register is a read-only register.

Software can affect the state of this register using the write-only *Local Interrupt Set Mask* register (*GIC\_VPEi\_SMASK*) at offset address 0x0010 and the *Local Interrupt Reset Mask* register (*GIC\_VPEi\_RMASK*) at offset address 0x000C. Software sets bit 3 of the *SMASK* register to enable the performance counter interrupt, or it can set bit 3 of the *RMASK* register to disable timer interrupts.

Once hardware has determined the masking characteristics of the interrupt, it uses the *Performance Counter Map-to-Pin* register at offset address 0x0050 to determine which *SI\_Int[5:0]*, *NMI*, or *YR\_ysif[15:0]* pins the interrupt will be driven onto. In non-EIC mode, bits 5:0 of this register are used to select one of 6 VPE interrupts. In non-EIC mode, only encodings 0 - 5 are valid. In EIC mode, the VPE encodes this field to support up to 63 interrupts.

#### 9.4.7.7 Software Interrupts

Each VPE provides two software interrupts; 0 and 1. Software interrupts originate from the CPU and are only used in EIC mode. In non-EIC mode they are routed internally.

When software interrupt 0 is generated, hardware sets bit 4 of the *Local Interrupt Pending* register (*GIC\_VPEi\_PEND*) at offset address 0x0004. Hardware then reads the state of bit 4 in the *Local Interrupt Mask* register (*GIC\_VPEi\_MASK*) at offset address 0x0008 to determine whether the performance counter interrupt has been masked. The *GIC\_VPEi\_MASK* register is a read-only register.

Software can affect the state of this register using the write-only *Local Interrupt Set Mask* register (*GIC\_VPEi\_SMASK*) at offset address 0x0010 and the *Local Interrupt Reset Mask* register (*GIC\_VPEi\_RMASK*) at offset address 0x000C. Software sets bit 4 of the *SMASK* register to enable the software interrupt 0, or it can set bit 4 of the *RMASK* register to disable software interrupt 0.

Once hardware has determined the masking characteristics of the interrupt, it uses the *Software Interrupt 0 Map-to-Pin* register at offset address 0x0054 to determine which *SI\_Int[5:0]*, *NMI*, or *YR\_ysif[15:0]* pins the interrupt will be driven onto. In non-EIC mode, bits 5:0 of this register are used to select one of 6 VPE interrupts. In non-EIC mode, only encodings 0 - 5 are valid. In EIC mode, the VPE encodes this field to support up to 63 interrupts.

The sequence is the same for software interrupt 1, except that bit 5 of each register noted above is set instead of bit 4. In addition, software uses the *Software Interrupt 1 Map-to-Pin* register at offset address 0x0058 to determine which *SI\_Int[5:0]* pin the interrupt will be driven onto.

#### 9.4.7.8 Fast Debug Channel Interrupts

When a Fast Debug Channel (FDC) interrupt is generated, hardware sets bit 6 of the *Local Interrupt Pending* register (*GIC\_VPEi\_PEND*) at offset address 0x0004. Hardware then reads the state of bit 6 in the *Local Interrupt Mask* register (*GIC\_VPEi\_MASK*) at offset address 0x0008 to determine whether the fast debug channel interrupt has been masked. The *GIC\_VPEi\_MASK* register is a read-only register.

Software can affect the state of this register using the write-only *Local Interrupt Set Mask* register (*GIC\_VPEi\_SMASK*) at offset address 0x0010 and the *Local Interrupt Reset Mask* register (*GIC\_VPEi\_RMASK*) at offset address 0x000C. Software sets bit 6 of the *SMASK* register to enable the fast debug channel interrupt, or it can set bit 6 of the *RMASK* register to disable fast debug channel interrupts.

Once hardware has determined the masking characteristics of the interrupt, it uses the *Fast Debug Channel Map-to-Pin* register at offset address 0x004C to determine which *SI\_Int[5:0]*, *NMI*, or *YR\_ysif[15:0]* pins the interrupt will be driven onto. In non-EIC mode, bits 5:0 of this register are used to select one of 6 VPE interrupts. In non-EIC mode, only encodings 0 - 5 are valid. In EIC mode, the interAptiv VPE encodes this field to support up to 63 interrupts.



## 9.4.8 EIC Mode Setting

EIC mode is controlled through software by setting the EIC\_MODE bit in the Local interrupt Control Register, GIC\_VPE\_CTL. Setting this bit enables EIC mode. This bit defaults to 0, vectored interrupt mode. Refer to [Section 9.6.2.1 “Local Interrupt Control Register”](#) for more information.

## 9.4.9 Enabling, Disabling, and Polling Local Interrupts

The Enabling, Disabling and Polling of local interrupts is configured through several registers in the GIC that are local to each processor.

There are 4 registers for Enabling, Disabling and Polling of local interrupts.

- Enabling an interrupt using the "GIC Local Set Mask Registers", GIC\_VPE\_SMASK
- Disabling an interrupt using the "GIC Local Reset Mask Registers", GIC\_VPE\_RMASK
- Determining the Enable/Disable state of an interrupt state using "GIC Local Interrupt Mask Register", GIC\_VPE\_MASK
- Polling the interrupt active state using the "GIC Local Interrupt Pending Register", GIC\_VPE\_PEND

### 9.4.9.1 Enabling External Interrupts

The "GIC Local Set Mask Register", GIC\_VPE\_SMASK is used to enable individual local interrupts. For synchronization purposes this is a write only register. Setting the bit enables the interrupt. The following table shows which field to set for each local interrupt. Refer to [Section 9.6.2.5 “Local Interrupt Set Mask Register”](#) for more information.

**Table 9.13 Enabling External Interrupts**

Field Name	Interrupt Controlled
FDC_MASK_SET	Fast Debug Channel
SWINT1_MASK_SET	Software interrupt 1
SWINT2_MASK_SET	Software interrupt 2
PERFCOUNT_MASK_SET	Local Performance Counter
TIMER_MASK_SET	CP0 Local Count/Compare Timer
COMPARE_MASK_SET	GIC Local Count/Compare Timer
WD_MASK_SET	Watchdog

### 9.4.9.2 Disabling External Interrupts

The "GIC Local Reset Mask Register", GIC\_VPE\_RMASK is used to disable individual local interrupts. For CPS synchronization purposes this is a write only register. Setting the bit disables the interrupt. The following table shows which field to set for each local interrupt. Refer to [Section 9.6.2.4 “Local Interrupt Reset Mask Register”](#) for more information.



**Table 9.14 Disabling External Interrupts**

Field Name	Interrupt Controlled
FDC_RESET_MASK	Fast Debug Channel
SWINT1_RESET_MASK	Software interrupt 1
SWINT2_RESET_MASK	Software interrupt 2
PERFCOUNT_RESET_MASK	Local Performance Counter
TIMER_RESET_MASK	CP0 Local Count/Compare Timer
COMPARE_RESET_MASK	GIC Local Count/Compare Timer
WD_RESET_MASK	Watchdog

**9.4.9.3 Determining the Enabled or Disabled Interrupt state**

The "GIC Local Mask Register", GIC\_VPE\_MASK is used to determine if a local interrupt is enabled. For CPS synchronization purposes this is a read only register. If a bit is set the corresponding interrupt source is enabled. If it is clear the corresponding interrupt is disabled. The following table shows which field corresponds to each local interrupt. Refer to [Section 9.6.2.3 "Local Interrupt Mask Register"](#) for more information

**Table 9.15 Determining the Enabled of Disabled Interrupt State**

Field Name	Interrupt Controlled
FDC_MASK	Fast Debug Channel
SWINT1_MASK	Software interrupt 1
SWINT2_MASK	Software interrupt 2
PERFCOUNT_MASK	Local Performance Counter
TIMER_MASK	CP0 Local Count/Compare Timer
COMPARE_MASK	GIC Local Count/Compare Timer
WD_MASK	Watchdog

**9.4.9.4 Polling for an Active Interrupt**

The "GIC Pending Register", GIC\_VPE\_PEND is used to determine if a external interrupt is active. This is a read only register. If a bit is set the corresponding local interrupt is active. If it is clear the corresponding interrupt is inactive. The following table shows which field corresponds to each local interrupt. Refer to [Section 9.6.2.2 "Local Interrupt Pending Register"](#) for more information

**Table 9.16 Polling for an Active Interrupt**

Field Name	Interrupt Controlled
FDC_PEND	Fast Debug Channel
SWINT1_PEND	Software interrupt 1
SWINT2_PEND	Software interrupt 2
PERFCOUNT_PEND	Local Performance Counter

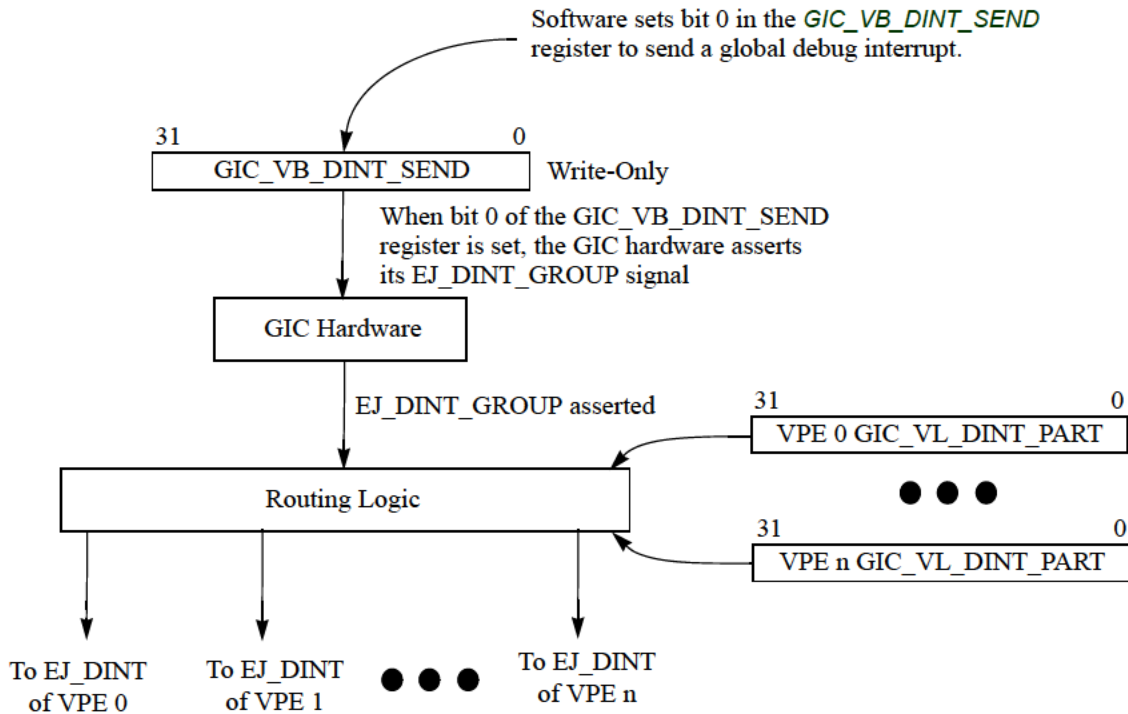
**Table 9.16 Polling for an Active Interrupt**

Field Name	Interrupt Controlled
TIMER_PEND	CP0 Local Count/Compare Timer
COMPARE_PEND	GIC Local Count/Compare Timer
WD_PEND	Watchdog

### 9.4.10 Debug Interrupt Generation

The GIC of the interAptiv Multiprocessing System allows software to globally assert a debug interrupt to all VPEs in the system. When the *Send\_DINT* bit of the *DINT Send to Group* register (*GIC\_VB\_DINT\_SEND*) in [Section 9.5.3.15, "DINT Send to Group Register"](#) is set, the *EJ\_DINT\_GROUP* signal of the GIC is asserted. Based on the state of this signal and the VPE-Local *GIC\_VL\_DINT\_PART* registers, hardware asserts the *EJ\_DINT* signal of each VPE in the system. This concept is shown in [Figure 9.8](#).

**Figure 9.8 Global EJTAG Debug Interrupt Generation in the GIC**



## 9.5 Shared Register Set

This section describes the various registers in the Shared register set.

### 9.5.1 GIC Register Field Types

For each register described below, field descriptions include the read/write properties of the field, and the reset state of the field. For single bit fields, the name is truncated to a single character which is then shown outside brackets in the Fields|Name column. For the read/write properties of the field, the following notation is used:

**Table 9.17 CP0 Register Field Types**

Notation	Hardware Interpretation	Software Interpretation
R/W	A field in which all bits are readable and writable by software and, potentially, by hardware. Hardware updates of this field are visible by software reads. Software updates of this field are visible by hardware reads. If the reset state of this field is “Undefined,” either software or hardware must initialize the value before the first read will return a predictable value. This should not be confused with the formal definition of <b>UNDEFINED</b> behavior.	
R	A field that is either static or is updated only by hardware. If the Reset State of this field is either “0” or “Preset”, hardware initializes this field to zero or to the appropriate state, respectively, on power up. If the Reset State of this field is “Undefined”, hardware updates this field only under those conditions specified in the description of the field.	A field to which the value written by software is ignored by hardware. Software may write any value to this field without affecting hardware behavior. Software reads of this field return the last value updated by hardware. If the Reset State of this field is “Undefined,” software reads of this field result in an <b>UNPREDICTABLE</b> value except after a hardware update done under the conditions specified in the description of the field.
W	A field that can be written by software but which can not be read by software. Software reads of this field will return an <b>UNDEFINED</b> value.	
0	A field that hardware does not update, and for which hardware can assume a zero value.	A field to which the value written by software must be zero. Software writes of non-zero values to this field may result in <b>UNDEFINED</b> behavior of the hardware. Software reads of this field return zero as long as all previous software writes are zero. If the Reset State of this field is “Undefined,” software must write this field with zero before it is guaranteed to read as zero.

## 9.5.2 Shared Section Register Map

The register map of the shared section is shown in [Table 9.18](#). These registers are accessible by any VPE. For the base address of this block, see [Table 9.3](#).

All registers are 32 bits wide and should only be accessed using 32-bit uncached load/stores. Reads from unpopulated registers in the GCMP address space should return 0x0, and writes to those locations should be silently dropped without generating any exceptions.

The addresses for the registers within the Shared Section of the GIC are calculated as follows:

$$\text{SharedSection\_Register\_Physical\_Address} = \text{GIC\_baseaddress} + \text{SharedSection\_baseoffset} + \text{Register\_Offset}$$

**Table 9.18 Shared Section Register Map**

Register Offset	Name	Type	Description
0x0000	GIC Config Register (GIC_SH_CONFIG)	R	Indicates the number of interrupts, number of VPEs, etc.
0x0010	GIC CounterLo (GIC_SH_CounterLo)	R/W	Shared Global Counter.
0x0014	GIC CounterHi (GIC_SH_CounterHi)	R/W	
0x0020	GIC Revision Register (GIC_RevisionID)	R	RevisionID of the GIC hardware.
0x0100	Global Interrupt Polarity Register0 (GIC_SH_POL31_0)	R/W	Polarity of the interrupt. For Level Type: 0x0 - Active Low 0x1 - Active High For Single Edge Type: 0x0 - Falling Edge used to set edge register 0x1 - Rising Edge used to set edge register At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.
0x0104	Global Interrupt Polarity Register1 (GIC_SH_POL63_32)	R/W	
0x0108	Global Interrupt Polarity Register2 (GIC_SH_POL95_64)	R/W	
0x010c	Global Interrupt Polarity Register3 (GIC_SH_POL127_96)	R/W	
0x0110	Global Interrupt Polarity Register4 (GIC_SH_POL159_128)	R/W	
0x0114	Global Interrupt Polarity Register5 (GIC_SH_POL191_160)	R/W	
0x0118	Global Interrupt Polarity Register6 (GIC_SH_POL223_192)	R/W	
0x011c	Global Interrupt Polarity Register7 (GIC_SH_POL255_224)	R/W	

**Table 9.18 Shared Section Register Map (continued)**

Register Offset	Name	Type	Description
0x0180	Global Interrupt Trigger Type Register0 (GIC_SH_TRIG31_0)	R/W	Edge or Level triggered 0x0 - Level 0x1 - Edge At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.
0x0184	Global Interrupt Trigger Type Register1 (GIC_SH_TRIG63_32)	R/W	
0x0188	Global Interrupt Trigger Type Register2 (GIC_SH_TRIG95_64)	R/W	
0x018c	Global Interrupt Trigger Type Register3 (GIC_SH_TRIG127_96)	R/W	
0x0190	Global Interrupt Trigger Type Register4 (GIC_SH_TRIG159_128)	R/W	
0x0194	Global Interrupt Trigger Type Register5 (GIC_SH_TRIG191_160)	R/W	
0x0198	Global Interrupt Trigger Type Register6 (GIC_SH_TRIG223_192)	R/W	
0x019c	Global Interrupt Trigger Type Register7 (GIC_SH_TRIG255_224)	R/W	
0x0200	Global Interrupt Dual Edge Register (GIC_SH_DUAL31_0)	R/W	Writing a 0x1 to any bit location sets the appropriate external interrupt source to be type dual-edged. At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.
0x0204	Global Interrupt Dual Edge Register (GIC_SH_DUAL63_32)	R/W	
0x0208	Global Interrupt Dual Edge Register (GIC_SH_DUAL95_64)	R/W	
0x020c	Global Interrupt Dual Edge Register (GIC_SH_DUAL127_96)	R/W	
0x0210	Global Interrupt Dual Edge Register (GIC_SH_DUAL159_128)	R/W	
0x0214	Global Interrupt Dual Edge Register (GIC_SH_DUAL191_160)	R/W	
0x0218	Global Interrupt Dual Edge Register (GIC_SH_DUAL223_192)	R/W	
0x021c	Global Interrupt Dual Edge Register (GIC_SH_DUAL255_224)	R/W	
0x0280	Global Interrupt Write Edge Register (GIC_SH_WEDGE)	W	Used for Interrupt Messages. Writes to this register atomically set or clear a specified bit in the <i>Edge Detect Register</i> .

**Table 9.18 Shared Section Register Map (continued)**

Register Offset	Name	Type	Description	
0x0300	Global Interrupt Reset Mask Register (GIC_SH_RMASK31_0)	W	Writing a 0x1 to any bit location masks off (disables) that interrupt. At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.	
0x0304	Global Interrupt Reset Mask Register (GIC_SH_RMASK63_32)	W		
0x0308	Global Interrupt Reset Mask Register (GIC_SH_RMASK95_64)	W		
0x030c	Global Interrupt Reset Mask Register (GIC_SH_RMASK127_96)	W		
0x0310	Global Interrupt Reset Mask Register (GIC_SH_RMASK159_128)	W		
0x0314	Global Interrupt Reset Mask Register (GIC_SH_RMASK191_160)	W		
0x0318	Global Interrupt Reset Mask Register (GIC_SH_RMASK223_192)	W		
0x031c	Global Interrupt Reset Mask Register (GIC_SH_RMASK255_224)	W		
0x0380	Global Interrupt Set Mask Register (GIC_SH_SMASK31_00)	W		Writing a 0x1 to any bit location sets the mask (enables) for that interrupt. At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.
0x0384	Global Interrupt Set Mask Register (GIC_SH_SMASK63_32)	W		
0x0388	Global Interrupt Set Mask Register (GIC_SH_SMASK95_64)	W		
0x038c	Global Interrupt Set Mask Register (GIC_SH_SMASK127_96)	W		
0x0390	Global Interrupt Set Mask Register (GIC_SH_SMASK159_128)	W		
0x0394	Global Interrupt Set Mask Register (GIC_SH_SMASK191_160)	W		
0x0398	Global Interrupt Set Mask Register (GIC_SH_SMASK223_192)	W		
0x039c	Global Interrupt Set Mask Register (GIC_SH_SMASK255_224)	W		

**Table 9.18 Shared Section Register Map (continued)**

Register Offset	Name	Type	Description
0x0400	Global Interrupt Mask Register (GIC_SH_MASK31_00)	R	Shows the enabled global interrupts. If bit N is set, global interrupt N is enabled. At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.
0x0404	Global Interrupt Mask Register (GIC_SH_MASK63_32)	R	
0x0408	Global Interrupt Mask Register (GIC_SH_MASK95_64)	R	
0x040c	Global Interrupt Mask Register (GIC_SH_MASK127_96)	R	
0x0410	Global Interrupt Mask Register (GIC_SH_MASK159_128)	R	
0x0414	Global Interrupt Mask Register (GIC_SH_MASK191_160)	R	
0x0418	Global Interrupt Mask Register (GIC_SH_MASK223_192)	R	
0x041c	Global Interrupt Mask Register (GIC_SH_MASK255_224)	R	
0x0480	Global Interrupt Pending Register (GIC_SH_PEND31_00)	R	
0x0484	Global Interrupt Pending Register (GIC_SH_PEND63_32)	R	
0x0488	Global Interrupt Pending Register (GIC_SH_PEND95_64)	R	
0x048c	Global Interrupt Pending Register (GIC_SH_PEND127_96)	R	
0x0490	Global Interrupt Pending Register (GIC_SH_PEND159_128)	R	
0x0494	Global Interrupt Pending Register (GIC_SH_PEND191_160)	R	
0x0498	Global Interrupt Pending Register (GIC_SH_PEND223_192)	R	
0x049c	Global Interrupt Pending Register (GIC_SH_PEND255_224)	R	
0x0500	Global Interrupt Map Src0 to Pin Register (GIC_SH_MAP0_PIN)	R/W	Maps this interrupt source to a particular pin - within <i>Int[5:0]</i> or <i>NMI</i> . At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources.
0x0504	Global Interrupt Map Src1 to Pin Register (GIC_SH_MAP1_PIN)	R/W	
0x0508	Global Interrupt Map Src2 to Pin Register (GIC_SH_MAP2_PIN)	R/W	
...	...	R/W	
0x08fc	Global Interrupt Map Src255 to Pin Register (GIC_SH_MAP255_PIN)	R/W	

**Table 9.18 Shared Section Register Map (continued)**

Register Offset	Name	Type	Description
0x2000	Global Interrupt Map Src0 to VPE Register (GIC_SH_MAP0_VPE31_0)	R/W	Assigns this interrupt source to a particular VPE.  At IP configuration time, the appropriate number of these registers are instantiated to support the number of External Interrupt Sources and the number of VPEs.
0x2020	Global Interrupt Map Src1 to VPE Register (GIC_SH_MAP1_VPE31_0)	R/W	
0x2040	Global Interrupt Map Src2 to VPE Register (GIC_SH_MAP2_VPE31_0)	R/W	
.....	.....	R/W	
0x3fe0	Global Interrupt Map Src255 to VPE Register (GIC_SH_MAP255_VPE31_0)	R/W	
0x6000	DINT Send to Group Register (GIC_VB_DINT_SEND)	R/W	Sends the DebugInterrupt to the specified VPE.
All other offsets	Reserved for future extensions		Reserved for future extensions.


### 9.5.3 Shared Section Register Descriptions

The physical address for the Shared Section registers is calculated as follows:

$$\text{GIC\_BaseAddress} + \text{SharedSection\_BaseAddress} + \text{RegisterOffset}$$

#### 9.5.3.1 Global Config Register

**Figure 9.9 Global Config Register Format**



**Table 9.19 GIC Config Register (GIC\_SH\_CONFIG — Offset 0x0000)**

Register Fields		Description	Read/ Write	Reset State
Name	Bits			
R	31:29	Reserved. Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
<i>COUNTSTOP</i>	28	Setting this bit will stop <i>GIC_CounterHi</i> and <i>GIC_CounterLo</i> . Used to freeze the shared counters when VPEs go into power-down or debug modes.	R/W	0
<i>COUNTBITS</i>	27:24	Number of Implemented Bits in <i>GIC_CounterHi</i> . Total Number of Counter Bits = 32 + COUNTBITS*4, E.g.: 0x0 = 32bits, <i>GIC_CounterHi</i> not implemented 0x1 = 36bits, <i>GIC_CounterHi</i> width = 4 bits 0x2 = 40bits, <i>GIC_CounterHi</i> width = 8 bits ... 0x7 = 60bits, <i>GIC_CounterHi</i> width = 28 bits 0x8 = 64bits, <i>GIC_CounterHi</i> width = 32 bits 9-15 Reserved	R	0x8



**Table 9.19 GIC Config Register (GIC\_SH\_CONFIG — Offset 0x0000) (continued)**

Register Fields		Description	Read/Write	Reset State																				
Name	Bits																							
<i>NUMINTERRUPTS</i>	23:16	Number of External Interrupt Sources  <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>8</td> </tr> <tr> <td>0x1</td> <td>16</td> </tr> <tr> <td>0x2</td> <td>24</td> </tr> <tr> <td>0x3</td> <td>32</td> </tr> <tr> <td>0x4</td> <td>40</td> </tr> <tr> <td>...</td> <td></td> </tr> <tr> <td>0x1E</td> <td>248</td> </tr> <tr> <td>0x1F</td> <td>256</td> </tr> <tr> <td>All others</td> <td>Reserved</td> </tr> </tbody> </table> Value is fixed by customer at IP configuration time.	Encoding	Meaning	0x0	8	0x1	16	0x2	24	0x3	32	0x4	40	...		0x1E	248	0x1F	256	All others	Reserved	R	IP Configuration Value
Encoding	Meaning																							
0x0	8																							
0x1	16																							
0x2	24																							
0x3	32																							
0x4	40																							
...																								
0x1E	248																							
0x1F	256																							
All others	Reserved																							
R	15:9	Reserved. Read as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0																				
<i>PVPES</i>	8:0	Total number of VPEs in the system.  0: 1 VPE 1: 2 VPE's 2: 3 VPE's 3: 4 VPE's 4: 5 VPE's 5: 6 VPE's 6: 7 VPE's 7: 8 VPE's 8: 9 VPE's	R	IP Configuration Value																				

### 9.5.3.2 GIC CounterLo

Figure 9.10 GIC CounterLo Register Format



Table 9.20 GIC CounterLo (GIC\_SH\_CounterLo — Offset 0x0010)

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>GIC_SH_CounterLo</i>	31:0	<p>Lower Half of an up-counter.</p> <p>When the counter reaches its maximum value, the counter rolls over to a value of 0x0.</p> <p>The counter is running at an implementation-specific frequency which is fixed, that is, not changing dynamically due to power management. It is recommended that this frequency be as close as possible to the highest clock frequency of the CPU subsystem.</p> <p>This counter is disabled by writing the <i>COUNTSTOP</i> bit in the <i>GIC_SH_CONFIG</i> register.</p> <p>This counter should only be written when <i>GIC_SH_CONFIG</i><sub>COUNTSTOP</sub> = 1; otherwise, the registers results after the write are unpredictable.</p>	R/W	0

### 9.5.3.3 GIC CounterHi

Figure 9.11 GIC CounterHi Register Format




Table 9.21 GIC CounterHi (GIC\_SH\_CounterHi — Offset 0x0014)

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>GIC_SH_CounterHi</i>	31:0	<p>Upper Half of an up-counter.</p> <p>When the counter reaches its maximum value, the counter rolls over to a value of 0x0.</p> <p>The counter is running at an implementation-specific frequency which is fixed, that is, not changing dynamically due to power management. It is recommended that this frequency be as close as possible to the highest clock frequency of the CPU subsystem.</p> <p>This counter is disabled by writing the <i>COUNTSTOP</i> bit in the <i>GIC_SH_CONFIG</i> register.</p> <p>This counter should only be written when <i>GIC_SH_CONFIG</i><sub>COUNTSTOP</sub> = 1; otherwise, the register results after the write are unpredictable.</p> <p>Unimplemented bits ignore writes and return 0 when read.</p>	R/W	0

### 9.5.3.4 GIC Revision Register

**Figure 9.12 GIC Revision Register Format**



**Table 9.22 GIC Revision Register (GIC\_RevisionID — Offset 0x0020)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
0	31:16	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0
<i>MAJOR_REV</i>	15:8	This field reflects the major revision of the GIC block. A major revision might reflect the changes from one product generation to another.	R	Preset
<i>MINOR_REV</i>	7:0	This field reflects the minor revision of the GIC block. A minor revision might reflect the changes from one release to another.	R	Preset

### 9.5.3.5 Global Interrupt Polarity Registers

There are eight Global Interrupt Polarity registers to cover all 256 possible system interrupts. These registers work in conjunction with the eight *Global Interrupt Trigger Type (GIC\_SH\_TRIGN)* and *Global Interrupt Dual Edge (GIC\_SH\_DUALn)* registers to select the polarity, active high/low trigger, and single/dual edge for each of the 256 interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

They are located at the following eight offsets.

**Table 9.23 Global Interrupt Polarity Register Mapping**

Offset	Acronym	Register Name
0x0100	GIC_SH_POL31_0	Polarity selection for interrupt pins 31:0
0x0104	GIC_SH_POL63_32	Polarity selection for interrupt pins 63:32
0x0108	GIC_SH_POL95_64	Polarity selection for interrupt pins 95:64
0x010C	GIC_SH_POL127_96	Polarity selection for interrupt pins 127:96
0x0110	GIC_SH_POL159_128	Polarity selection for interrupt pins 159:128
0x0114	GIC_SH_POL191_160	Polarity selection for interrupt pins 191:160
0x0118	GIC_SH_POL223_192	Polarity selection for interrupt pins 223:191
0x011C	GIC_SH_POL255_224	Polarity selection for interrupt pins 255:192

In the register below, the *x\_y* nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_POL63\_32 indicates that this register handles the polarity for interrupts 63:32.

**Figure 9.13 GIC Interrupt Polarity Register Format**

31

0

GIC_SH_POLx_y
---------------

**Table 9.24 Global Interrupt Polarity Registers (GIC\_SH\_POLx\_y — See Table 9.23 for Mapping)**

Register Fields		Description	Read/ Write	Reset State												
Name	Bits															
<i>GIC_SH_POLx_y</i>	31:0	<p>Each bit in this register represents an interrupt source. The state of the bit indicates the polarity of the interrupt. If the interrupt type (as denoted by <i>Global Interrupt Trigger Type</i> and <i>Global Interrupt Dual Edge</i> registers) is Level:</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th style="text-align: center;">Encoding</th> <th style="text-align: center;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">Active Low</td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">Active High</td> </tr> </tbody> </table> <p>If the interrupt type is Single-edge:</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th style="text-align: center;">Encoding</th> <th style="text-align: center;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">Falling Edge denotes interrupt source has toggled.</td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">Rising Edge denotes interrupt source has toggled.</td> </tr> </tbody> </table> <p>If the interrupt type is Dual-edge, this register is not used</p>	Encoding	Meaning	0	Active Low	1	Active High	Encoding	Meaning	0	Falling Edge denotes interrupt source has toggled.	1	Rising Edge denotes interrupt source has toggled.	R/W	0
Encoding	Meaning															
0	Active Low															
1	Active High															
Encoding	Meaning															
0	Falling Edge denotes interrupt source has toggled.															
1	Rising Edge denotes interrupt source has toggled.															

### 9.5.3.6 Global Interrupt Trigger Type Registers

There are eight Global Interrupt Trigger Type registers to cover all 256 possible system interrupts. These registers work in conjunction with the eight *Global Interrupt Polarity (GIC\_SH\_POLn)* and *Global Interrupt Dual Edge (GIC\_SH\_DUALn)* registers to select the polarity, active high/low trigger, and single/dual edge for each of the 256 interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

They are located at the following eight offsets.

**Table 9.25 Global Interrupt Trigger Type Register Mapping**

Offset	Acronym	Register Name
0x0180	GIC_SH_TRIG31_0	Interrupt trigger selection for interrupt pins 31:0
0x0184	GIC_SH_TRIG63_32	Interrupt trigger selection for interrupt pins 63:32
0x0188	GIC_SH_TRIG95_64	Interrupt trigger selection for interrupt pins 95:64
0x018C	GIC_SH_TRIG127_96	Interrupt trigger selection for interrupt pins 127:96
0x0190	GIC_SH_TRIG159_128	Interrupt trigger selection for interrupt pins 159:128
0x0194	GIC_SH_TRIG191_160	Interrupt trigger selection for interrupt pins 191:160
0x0198	GIC_SH_TRIG223_192	Interrupt trigger selection for interrupt pins 223:191
0x019C	GIC_SH_TRIG255_224	Interrupt trigger selection for interrupt pins 255:192

In the register below, the x\_y nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_TRIG63\_32 indicates that this register handles the trigger level for interrupts 63:32.

**Figure 9.14 GIC Interrupt Trigger Type Register Format**

31

0

GIC_SH_TRIGx_y
----------------

**Table 9.26 Global Interrupt Trigger Type Registers (GIC\_SH\_TRIGx\_y — See Table 9.25 for Mapping)**

Register Fields		Description	Read/ Write	Reset State						
Name	Bits									
<i>GIC_SH_TRIGx_y</i>	31:0	Each bit in this register represents an interrupt source. The state of the bit indicates the nature of the interrupt signaling.	R/W	0						
		<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">Encoding</th> <th style="text-align: center;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center; padding: 2px;">0</td> <td style="padding: 2px;">Level</td> </tr> <tr> <td style="text-align: center; padding: 2px;">1</td> <td style="padding: 2px;">Edge Single edge or dual-edge signaling denoted by <i>Global Interrupt Dual Edge Register</i>.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Level	1	Edge Single edge or dual-edge signaling denoted by <i>Global Interrupt Dual Edge Register</i> .		
Encoding	Meaning									
0	Level									
1	Edge Single edge or dual-edge signaling denoted by <i>Global Interrupt Dual Edge Register</i> .									

### 9.5.3.7 Global Interrupt Dual Edge Registers

There are eight Global Interrupt Dual Edge registers to cover all 256 possible system interrupts. These registers work in conjunction with the eight *Global Interrupt Polarity (GIC\_SH\_POLn)* and *Global Interrupt Trigger Type (GIC\_SH\_TRIGn)* registers to select the polarity, active high/low trigger, and single/dual edge for each of the 256 interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

They are located at the following eight offsets.

**Table 9.27 Global Interrupt Dual Edge Register Mapping**

Offset	Acronym	Register Name
0x0200	GIC_SH_DUAL31_0	Interrupt single/dual edge selection for interrupt pins 31:0
0x0204	GIC_SH_DUAL63_32	Interrupt single/dual edge selection for interrupt pins 63:32
0x0208	GIC_SH_DUAL95_64	Interrupt single/dual edge selection for interrupt pins 95:64
0x020C	GIC_SH_DUAL127_96	Interrupt single/dual edge selection for interrupt pins 127:96
0x0210	GIC_SH_DUAL159_128	Interrupt single/dual edge selection for interrupt pins 159:128
0x0214	GIC_SH_DUAL191_160	Interrupt single/dual edge selection for interrupt pins 191:160
0x0218	GIC_SH_DUAL223_192	Interrupt single/dual edge selection for interrupt pins 223:191
0x021C	GIC_SH_DUAL255_224	Interrupt single/dual edge selection for interrupt pins 255:192

In the register below, the x\_y nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_DUAL63\_32 indicates that this register handles the edge triggering for interrupts 63:32.

**Figure 9.15 GIC Interrupt Dual Edge Register Format**

31

0

GIC_SH_DUALx_y
----------------

**Table 9.28 Global Dual Edge Registers (GIC\_SH\_DUALx\_y — See Table 9.27 for Mapping)**

Register Fields		Description	Read/Write	Reset State						
Name	Bits									
<i>GIC_SH_DUALx_y</i>	31:0	<p>Each bit in this register represents an interrupt source. This register is only meaningful if the equivalent bit in the <i>Global Interrupt Trigger Type</i> register is set to 0x1 = Edge signaling. Indicates single or dual-edged signaling.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>Single edge</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Dual Edge</td> </tr> </tbody> </table>	Encoding	Meaning	0	Single edge	1	Dual Edge	R/W	0
Encoding	Meaning									
0	Single edge									
1	Dual Edge									

### 9.5.3.8 Global Interrupt Write Edge Register

This register is used to support interrupt messages. A write to this register will atomically set or clear one bit in the *Edge Detect Register*. Setting a bit in this register will be treated equivalently to having the edge detection logic see an active edge. This bypasses the edge detection logic and thus it does not matter whether the corresponding interrupt is configured to be rising, falling, or dual edge sensitive. However, the behavior is undefined unless the equivalent bit in the *Global Interrupt Trigger Type* register is set to 0x1 indicating edge signaling.

**Figure 9.16 GIC Interrupt Write Edge Register Format**

31 30

0

RW	INTERRUPT
----	-----------

**Table 9.29 Global Interrupt Write Edge Registers (GIC\_SH\_WEDGE Offset 0x0280)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>RW</i>	31	<p>Controls whether this write is setting or clearing a bit in the <i>Edge Detect Register</i>. If this bit is set, the selected bit in the register is set. If this bit is cleared, the selected bit in the register is cleared.</p>	W	Undefined
<i>Interrupt</i>	30:0	<p>This field is the encoded value of the interrupt that is being cleared or set. For example, a value of 0xB means interrupt 11 (decimal).</p>	W	Undefined

### 9.5.3.9 Global Interrupt Reset Mask Registers

There are eight Global Interrupt Reset Mask registers to cover all 256 possible system interrupts. These registers work in conjunction with the eight *Global Interrupt Set Mask (GIC\_SH\_SMASKn)* registers to enable and disable individual interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

These registers are located at the following eight offsets.

**Table 9.30 Global Interrupt Reset Mask Register Mapping**

Offset	Acronym	Register Name
0x0300	GIC_SH_RMASK31_0	Interrupt reset mask for interrupt pins 31:0
0x0304	GIC_SH_RMASK63_32	Interrupt reset mask for interrupt pins 63:32
0x0308	GIC_SH_RMASK95_64	Interrupt reset mask for interrupt pins 95:64
0x030C	GIC_SH_RMASK127_96	Interrupt reset mask for interrupt pins 127:96
0x0310	GIC_SH_RMASK159_128	Interrupt reset mask for interrupt pins 159:128
0x0314	GIC_SH_RMASK191_160	Interrupt reset mask for interrupt pins 191:160
0x0318	GIC_SH_RMASK223_192	Interrupt reset mask for interrupt pins 223:192
0x031C	GIC_SH_RMASK255_224	Interrupt reset mask for interrupt pins 255:192

In the register below, the x\_y nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_RMASK63\_32 indicates that this register handles the reset mask for interrupts 63:32.

**Figure 9.17 GIC Interrupt Reset Mask Register Format**



**Table 9.31 Global Interrupt Reset Mask Registers (GIC\_SH\_RMASKx\_y — See Table 9.30 for Mapping)**

Register Fields		Description	Read/ Write	Reset State
Name	Bits			
<i>GIC_SH_RMASKx_y</i>	31:0	Each bit in this register represents an interrupt source. Writing this register with a 0x1 in any bit position(s) will cause only the corresponding bit/interrupt(s) in the <i>Global Interrupt Mask Register</i> to be reset (value->0). This is used by software to temporarily disable interrupts.	W	Undefined

### 9.5.3.10 Global Interrupt Set Mask Registers

There are eight Global Interrupt Set Mask registers to cover all 256 possible system interrupts. These registers work in conjunction with the eight *Global Interrupt Reset Mask (GIC\_SH\_RMASKn)* registers to enable and disable individual interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

These registers are located at the following eight offsets.

**Table 9.32 Global Interrupt Set Mask Register Mapping**

Offset	Acronym	Register Name
0x0380	GIC_SH_SMASK31_0	Interrupt set mask for interrupt pins 31:0
0x0384	GIC_SH_SMASK63_32	Interrupt set mask for interrupt pins 63:32
0x0388	GIC_SH_SMASK95_64	Interrupt set mask for interrupt pins 95:64
0x038C	GIC_SH_SMASK127_96	Interrupt set mask for interrupt pins 127:96

**Table 9.32 Global Interrupt Set Mask Register Mapping (continued)**


Offset	Acronym	Register Name
0x0390	GIC_SH_SMASK159_128	Interrupt set mask for interrupt pins 159:128
0x0394	GIC_SH_SMASK191_160	Interrupt set mask for interrupt pins 191:160
0x0398	GIC_SH_SMASK223_192	Interrupt set mask for interrupt pins 223:192
0x039C	GIC_SH_SMASK255_224	Interrupt set mask for interrupt pins 255:192

In the register below, the x\_y nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_SMASK63\_32 indicates that this register handles the set mask for interrupts 63:32.

**Figure 9.18 GIC Interrupt Set Mask Register Format**

31

0



**Table 9.33 Global Set Mask Registers (GIC\_SH\_SMASKx\_y — See Table 9.32 for Mapping)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
GIC_SH_SMASKx_y	31:0	Each bit in this register represents an interrupt source. Writing this register with a 0x1 in any bit position(s) will cause only the corresponding bit/interrupt(s) in the <i>Global Interrupt Mask Register</i> to be set (value->0x1). This is used by software to enable interrupts.	W	Undefined

### 9.5.3.11 Global Interrupt Mask Registers

There are eight Global Interrupt Reset Mask registers to cover all 256 possible system interrupts. These read-only registers are used to indicate when an external interrupt occurs. An individual interrupt bit is set when an interrupt occurs and the corresponding Global Interrupt Set Mask bit is set, thereby enabling the interrupt. Refer to [Section 9.5.3.10, "Global Interrupt Set Mask Registers"](#) for more information.

These registers work in conjunction with the eight *Global Interrupt Set Mask (GIC\_SH\_SMASKn)* and *Global Interrupt Reset Mask (GIC\_SH\_RMASKn)* registers to manage and process interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

These registers are located at the following eight offsets.

**Table 9.34 Global Interrupt Mask Register Mapping**

Offset	Acronym	Register Name
0x0400	GIC_SH_MASK31_0	Interrupt status for interrupt pins 31:0
0x0404	GIC_SH_MASK63_32	Interrupt status for interrupt pins 63:32
0x0408	GIC_SH_MASK95_64	Interrupt status for interrupt pins 95:64
0x040C	GIC_SH_MASK127_96	Interrupt status for interrupt pins 127:96
0x0410	GIC_SH_MASK159_128	Interrupt status for interrupt pins 159:128
0x0414	GIC_SH_MASK191_160	Interrupt status for interrupt pins 191:160



**Table 9.34 Global Interrupt Mask Register Mapping (continued)**

Offset	Acronym	Register Name
0x0418	GIC_SH_MASK223_192	Interrupt status for interrupt pins 223:191
0x041C	GIC_SH_MASK255_224	Interrupt status for interrupt pins 255:192

In the register below, the x\_y nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_MASK63\_32 indicates that this register handles the masking for interrupts 63:32.

**Figure 9.19 GIC Interrupt Mask Register Format**



**Table 9.35 Global Interrupt Mask Registers (GIC\_SH\_MASKx\_y — See Table 9.34 for Mapping)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>GIC_SH_MASKx_y</i>	31:0	Each bit in this register represents an interrupt source. Reports which of the external interrupt sources are enabled. Used by software to determine which interrupt sources are currently enabled.	R	0x00000000

### 9.5.3.12 Global Interrupt Pending Registers

There are eight Global Interrupt Pending registers to cover the pending status of all 256 possible system interrupts. These read-only registers are set by hardware when an external interrupt is pending.

These registers work in conjunction with the eight *Global Interrupt Set Mask (GIC\_SH\_SMASKn)*, *Global Interrupt Reset Mask (GIC\_SH\_RMASKn)*, and *Global Interrupt Mask (GIC\_SH\_MASKn)* registers to manage and process interrupts. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

These registers are located at the following eight offsets.

**Table 9.36 Global Interrupt Pending Register Mapping**

Offset	Acronym	Register Name
0x0480	GIC_SH_PEND31_0	Interrupt pending status for interrupt pins 31:0
0x0484	GIC_SH_PEND63_32	Interrupt pending status for interrupt pins 63:32
0x0488	GIC_SH_PEND95_64	Interrupt pending status for interrupt pins 95:64
0x048C	GIC_SH_PEND127_96	Interrupt pending status for interrupt pins 127:96
0x0490	GIC_SH_PEND159_128	Interrupt pending status for interrupt pins 159:128
0x0494	GIC_SH_PEND191_160	Interrupt pending status for interrupt pins 191:160
0x0498	GIC_SH_PEND223_192	Interrupt pending status for interrupt pins 223:191
0x049C	GIC_SH_PEND255_224	Interrupt pending status for interrupt pins 255:192

In the register below, the x\_y nomenclature indicates the bit range covered by each register shown above. For example, GIC\_SH\_PEND63\_32 indicates that this register handles the interrupt pending status for interrupts 63:32.

**Figure 9.20 GIC Interrupt Pending Register Format**



**Table 9.37 Global Interrupt Pending Registers (GIC\_SH\_PENDx\_y — See Table 9.36 for Mapping)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>GIC_SH_PENDx_y</i>	31:0	There are eight Interrupt Pending register that are used to indicate the pending status of all 256 possible interrupts in the system Each bit indicates which of the external interrupt sources are asserted/pending before masking.  Used by software to find the external source that caused the CPU interrupt.	R	Undefined


**9.5.3.13 Global Interrupt Map to Pin Registers**

There are up to 256 Global Interrupt Map-to-Pin registers in the GIC to cover the mapping of all 256 possible system interrupts. This corresponds to one register per external interrupt signal. The number of registers instantiated at build time depends on the number of external system interrupts. These are write-only registers. Software is not expected to change these registers frequently. Software is expected to keep a shadow copy of these registers in memory so that Read-Modify-Write hazards are avoided.

Each interrupt pin can be mapped to one of three signal types: *SI\_Int[5:0]*, *SI\_NMI*, or *SI\_YSI[15:0]*. Bits 31:29 of this register are used to indicate to which signal type the interrupt will be mapped. Only one of these bits can be set at any given time. Bits 5:0 indicate the actual mapping for each external interrupt pin. For example, if bit 31 of this register is set, the external interrupt is routed to the *SI\_Int[5:0]* pins of the appropriate VPE.

For the register offset addresses corresponding to each register, refer [Table 9.7, "Mapping of External Interrupts"](#)

**Figure 9.21 GIC Interrupt Map to Pin Register Format**



**Table 9.38 Global Interrupt Map to Pin Registers (GIC\_SH\_MAPx\_y)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>MAP_TO_PIN</i>	31	If this bit is set, this interrupt source is mapped to a VPE interrupt pin (specified by the <i>MAP</i> field below). Only one of the <i>MAP_TO_PIN</i> or <i>MAP_TO_NMI</i> bits can be set at any one time.	RW	0x1
<i>MAP_TO_NMI</i>	30	If this bit is set, this interrupt source is mapped to NMI. Only one of the <i>MAP_TO_PIN</i> or <i>MAP_TO_NMI</i> , or <i>MAP_TO_YQ</i> bits can be set at any one time.	RW	0

**Table 9.38 Global Interrupt Map to Pin Registers (GIC\_SH\_MAPx\_y) (continued)**

Register Fields		Description	Read/ Write	Reset State																
Name	Bits																			
<i>MAP_TO_YQ</i>	29	If this bit is set, this interrupt source is mapped to an MT Yield Qualifier pin. Only one of the <i>MAP_TO_PIN</i> , <i>MAP_TO_NMI</i> , or <i>MAP_TO_YQ</i> bits can be set at any one time.	RW	0																
Reserved	28:6	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	-	0																
<i>MAP</i>	5:0	<p>When the <i>MAP_TO_PIN</i> bit is set, this field contains the encoded value of the VPE interrupts signals <i>Int[62:0]</i>.</p> <p>In EIC mode, this represents one less than the EIC interrupt level (e.g. a value of 0x20 represents interrupt level 21).</p> <p>For non-EIC mode, the value represents the CPU interrupt to be asserted (e.g. a value of 0x03 represents interrupt 3), and only values of 0 to 5 are legal.</p> <p>When <i>MAP_TO_YP</i> is set, this field contains the encoded signal selection of the Yield Qualifier.</p> <table border="1" data-bbox="634 871 1027 1186"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td><i>YR_ysi[0]</i></td> </tr> <tr> <td>0x1</td> <td><i>YR_ysi[1]</i></td> </tr> <tr> <td>...</td> <td></td> </tr> <tr> <td>0xF</td> <td><i>YR_ysi[15]</i></td> </tr> <tr> <td>0x10</td> <td>NULL</td> </tr> <tr> <td>...</td> <td></td> </tr> <tr> <td>0x1F</td> <td>NULL</td> </tr> </tbody> </table> <p>Since YQ is per-CORE rather than per-VPE, software needs to apply proper protection across VPEs by using, for example, the <i>cop0 YQMask</i> register.</p>	Encoding	Meaning	0x0	<i>YR_ysi[0]</i>	0x1	<i>YR_ysi[1]</i>	...		0xF	<i>YR_ysi[15]</i>	0x10	NULL	...		0x1F	NULL	RW	0x5 (Timer, FDC, PerfCnt) 0x0 (All Others)
Encoding	Meaning																			
0x0	<i>YR_ysi[0]</i>																			
0x1	<i>YR_ysi[1]</i>																			
...																				
0xF	<i>YR_ysi[15]</i>																			
0x10	NULL																			
...																				
0x1F	NULL																			

### 9.5.3.14 Global Interrupt Map to VPE Registers

There are up to 512 Global Interrupt Map-to-VPE registers in the GIC to cover the mapping of all 256 possible system interrupts. This corresponds to two registers per external interrupt signal. However, the high-order register is not used in the interAptiv core as described in [Section 9.5.3.14, "Global Interrupt Map to VPE Registers"](#).


The number of registers instantiated at build time depends on the number of external system interrupts. These are write-only registers. Software is not expected to change these registers frequently. Software is expected to keep a shadow copy of these registers in memory so that Read-Modify-Write hazards are avoided.

For the register offset addresses corresponding to each register, refer [Table 9.7, "Mapping of External Interrupts"](#)

**Figure 9.22 GIC Interrupt Map to VPE Register Format**

31

0




**Table 9.39 Global Interrupt Map to VPE Registers (GIC\_SH\_MAP\_VPEn — See [Table 9.7](#) for Mapping)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>GIC_SH_MAPi_VPEn</i>	31:0	Setting any bit in this register causes the interrupt source to be routed to the corresponding VPE. If the GIC_SH_MAPi_PIN[MAP_TO_YQ] is set, each bit in this register represents a CPU core rather than a VPE, and the interrupt source will be routed to the corresponding CPU core.  For all GIC_SH_MAPi_VPE registers, only one bit may be set at a time. That is, an interrupt source will be routed to one and only one VPE or CPU core.	W	0

### 9.5.3.15 DINT Send to Group Register

This register allows software to assert the `EJ_DINT_GROUP` signal directly. Refer to [Section 9.4.10 “Debug Interrupt Generation”](#) for more information.

**Figure 9.23 DINT Send to Group Register Format**



**Table 9.40 DINT Send to Group Register (GIC\_VB\_DINT\_SEND Offset 0x6000)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
R	[31:1]	Read as Zero. Writes ignored.	-	0x0
<i>SEND_DINT</i>	[0]	If this register field is written with a value of 0x1, the <i>EJ_DINT_GROUP</i> signal is asserted in a one-shot manner.	W	0x0

See [Chapter 17, “Multi-CPU Debug” on page 745](#) for more information about how this register is used.

## 9.6 GIC VPE-Local and VPE-Other Register Set

### 9.6.1 VPE-Local and VPE-Other Register Maps

The VPE-Local and VPE-Other interrupt register maps are described in [Table 9.41](#) below. For the base addresses of these blocks, see [Table 9.3](#). Each VPE in the interAptiv core contains a set of these registers.

The physical address for the registers within the VPE-Local section are calculated as follows:

$$\text{VPE-Local\_Register\_Physical\_Address} = \text{GIC\_BaseAddress} + \text{VPE-Local\_BaseOffset} + \text{Register Offset}$$

Similarly, for the VPE-Other section:

$$\text{VPE-Other\_Register\_Physical\_Address} = \text{GIC\_BaseAddress} + \text{VPE-Other\_BaseOffset} + \text{Register Offset}$$

All registers are 32 bits wide and should only be accessed using 32-bit uncached load/stores. Reads from unpopulated registers in the GCMP address space will return 0x0, and writes to those locations will be silently dropped without generating any exceptions.

**Table 9.41 VPE-Local and VPE-Other Register Maps**

Register Offset	Name	Type	Description
0x0000	Local Interrupt Control Register (GIC_VPEi_CTL)	R/W	Enable EIC Mode.
0x0004	Local Interrupt Pending Register (GIC_VPEi_PEND)	R	Status of the local interrupts before masking.
0x0008	Local Mask Register (GIC_VPEi_MASK)	R	Mask bits, if set will enable the corresponding interrupts in the interrupt vector.
0x000c	Local Reset Mask Register (GIC_VPEi_RMASK)	W	Setting a bit in this register causes the corresponding bits in the <i>GIC_VPEi_MASK</i> register to be cleared atomically with respect to other bits.
0x0010	Local Set Mask Register (GIC_VPEi_SMASK)	W	Setting a bit in this register causes the corresponding bits in the <i>GIC_VPEi_MASK</i> register to be set atomically with respect to other bits.
0x0040	Local WatchDog Map-to-Pin Register (GIC_VPEi_WD_MAP)	R/W	This register is used to route the local WatchDog interrupt to the desired VPE pin.
0x0044	Local GIC Counter/Compare Map-to-Pin Register (GIC_VPEi_COMPARE_MAP)	R/W	This register is used to route the local GIC Compare/Count Interrupt to the desired VPE pin. This is an optional register instantiated at IP configuration time.
0x0048	Local CPU Timer Map-to-Pin Register (GIC_VPEi_TIMER_MAP)	R/W	This register is used to route the local CPU Timer interrupt to the desired VPE pin.

**Table 9.41 VPE-Local and VPE-Other Register Maps (continued)**

Register Offset	Name	Type	Description
0x004c	Local CPU Fast Debug Channel Map-to-Pin Register (GIC_VPEi_FDC_MAP)	R/W	This register is used to route the local CPU Fast Debug Channel interrupt to the desired VPE pin. This is an optional register instantiated at IP configuration time.
0x0050	Local Perf Counter Map-to-Pin Register (GIC_VPEi_PERFCTR_MAP)	R/W	This register is used to route the local Performance Counter interrupt to the desired VPE pin. This is an optional register instantiated at IP configuration time.
0x0054	Local SWInt0 Map-to-Pin Register (GIC_VPEi_SWInt0_MAP)	R/W	This register is used to route the local SWInt0 interrupt to the desired VPE pin. This is an optional register instantiated at IP configuration time.
0x0058	Local SWInt1 Map-to-Pin Register (GIC_VPEi_SWInt1_MAP)	R/W	This register is used to route the local SWInt1 interrupt to the desired VPE pin. This is an optional register instantiated at IP configuration time.
0x0080	VPE-Other Addressing Register (GIC_VPEi_OTHER_ADDR)	R/W	Sets the <i>VPENum</i> of the register that will be accessed through the VPE-Other address space.
0x0088	VPE-Local Identification Register (GIC_VPEi_IDENT)	R	Indicates the VPE number of the local VPE.
0x0090	Programmable/Watchdog Timer0 Config Register (GIC_VPEi_WD_CONFIG0)	R/W	Local Programmable or Watchdog Timer0 related registers. See register description for more details.
0x0094	Programmable/Watchdog Timer0 Count Register (GIC_VPEi_WD_COUNT0)	R	
0x0098	Programmable/Watchdog Timer0 Initial Count Register (GIC_VPEi_WD_INITIAL0)	R/W	
0x00A0	CompareLo Register (GIC_VPEi_CompareLo)	R/W	Compare Register. See register description for more details.
0x00A4	CompareHi Register (GIC_VPEi_CompareHi)	R	
0x0100	EIC Shadow Set for Interrupt Src0 (GIC_VPEi_EICSS0)	R/W	EIC Shadow Set for Interrupt Source0.
0x0104	EIC Shadow Set for Interrupt Src1 (GIC_VPEi_EICSS1)	R/W	EIC Shadow Set for Interrupt Source1.
0x0108 - 0x01F8	EIC Shadow Set for Interrupt Src2 through InterrVPEi_EICSS62)	R/W	EIC Shadow Set for Interrupt Source2 through Source62.
0x01FC	EIC Shadow Set for Interrupt Src63 (GIC_VPEi_EICSS63)	R/W	EIC Shadow Set for Interrupt Source63.
0x3000	VPE-Local DINT Group Participate Register (GIC_VL_DINT_PART GIC_VO_DINT_PART)	R/W	Controls whether this VPE pays attention to the <i>DebugInt_GroupRequest</i> register.

**Table 9.41 VPE-Local and VPE-Other Register Maps (continued)**

Register Offset	Name	Type	Description
0x3080	VPE-Local DebugBreak Group Register (GIC_VL_BRK_GROUP GIC_VO_BRK_GROUP)	R/W	Allows multiple VPE to simultaneously enter Debug Mode.
All Other Offsets	RESERVED		Reserved for Future Extensions.

## 9.6.2 VPE-Local and VPE-Other Section Register Description

The following subsections describes the registers of the VPE-Local and VPE-Other sections.

### 9.6.2.1 Local Interrupt Control Register

**Figure 9.24 Local Interrupt Control Register Format**

31	5	4	3	2	1	0
R	FDC_ ROUTABLE	SWINT_ ROUTABLE	PERFCOUNT_ ROUTABLE	TIMER_ ROUTABLE	EIC_MODE	

**Table 9.42 Local Interrupt Control Register (GIC\_VPEi\_CTL — Offset 0x0000)**

Register Fields		Description	Read/ Write	Reset State
Name	Bits			
RESERVED	31:5	Read as 0x0. Writes ignored. Must be written with a value of 0x0.		0
<i>FDC_ROUTABLE</i>	4	If this bit is set, the CPU Fast Debug Channel Interrupt is routable within the GIC. If this bit is clear, the CPU Fast Debug Channel Interrupt is hardwired to one of the <i>SI_Int</i> pins as described by the CPU's COP0 <i>IntCtlIPFDCI</i> register field.	R	IP Configuration Value
<i>SWINT_ROUTABLE</i>	3	If this bit is set, the CPU SW Interrupts are routable within the GIC. If this bit is clear, then the CPU SW Interrupts are routed back to the CPU directly.	R	IP Configuration Value
<i>PERFCOUNT_ROUTABLE</i>	2	If this bit is set, the CPU Performance Counter Interrupt is routable within the GIC. If this bit is clear, the CPU Performance Counter Interrupt is hardwired to one of <i>SI_Int</i> pins as described by the CPU's COP0 <i>IntCtlPPCI</i> register field.	R	IP Configuration Value
<i>TIMER_ROUTABLE</i>	1	If this bit is set, the CPU Timer Interrupt is route-able within the GIC. If this bit is clear, the CPU Timer Interrupt is hardwired to one of the <i>SI_Int</i> pins, as described by the CPU's COP0 <i>IntCtlPTI</i> register field.	R	IP Configuration Value
<i>EIC_MODE</i>	0	Writing a 1 to this bit will set this local interrupt controller to EIC (External Interrupt Controller) mode.	R/W	0



### 9.6.2.2 Local Interrupt Pending Register

**Figure 9.25 Local Interrupt Pending Register Format**

31	7	6	5	4	3	2	1	0
R	FDCPEND	SWINT1_ PEND	SWINT0_ PEND	PERFCOUNT_ PEND	TIMER_ PEND	COMPARE_ PEND	WQ_PEND	

**Table 9.43 Local Interrupt Pending Register (GIC\_VPEi\_PEND — Offset 0x0004)**

Register Fields		Description	Read/ Write	Reset State
Name	Bits			
R	31:7	Read as 0x0. Writes ignored. Must be written with a value of 0x0.		0
<i>FDC_PEND</i>	6	Indicates the status of the local Fast Debug Channel interrupt prior to masking	R	Undefined
<i>SWINT1_PEND</i>	5	Indicates the status of the local SW interrupt1 prior to masking.	R	Undefined
<i>SWINT0_PEND</i>	4	Indicates the status of the local SW interrupt0 prior to masking.	R	Undefined
<i>PERFCOUNT_PEND</i>	3	Indicates the status of the local Performance Counter interrupt prior to masking.	R	Undefined
<i>TIMER_PEND</i>	2	Indicates the status of the local CPU Timer interrupt prior to masking.	R	Undefined
<i>COMPARE_PEND</i>	1	Indicates the status of the local GIC Count/Compare interrupt prior to masking.	R	Undefined
<i>WD_PEND</i>	0	Indicates the status of the local WatchDog interrupt prior to masking.	R	Undefined

### 9.6.2.3 Local Interrupt Mask Register

This is a read-only register. Refer to [Section 9.4.2, "Configuring Interrupt Sources"](#) for more information.

**Figure 9.26 Local Interrupt Mask Register Format**

31	7	6	5	4	3	2	1	0
R	FDC- MASK	SWINT1_ MASK	SWINT0_ MASK	PERFCOUNT_ MASK	TIMER_ MASK	COMPARE_ MASK	WQ_MASK	

**Table 9.44 Local Interrupt Mask Register (GIC\_VPEi\_MASK — Offset 0x0008)**


Register Fields		Description	Read/ Write	Reset State
Name	Bits			
RESERVED	31:7	Read as 0x0	R	0x0
<i>FDC_MASK</i>	6	If this bit is set, the local Fast Debug Channel interrupt is enabled	R	0x1
<i>SWINT1_MASK</i>	5	If this bit is set, the local SWInt1 Interrupt is enabled.	R	0x1
<i>SWINT0_MASK</i>	4	If this bit is set, the local SWInt0 Interrupt is enabled.	R	0x1

**Table 9.44 Local Interrupt Mask Register (GIC\_VPEi\_MASK — Offset 0x0008)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>PERFCNT_MASK</i>	3	If this bit is set, the local Performance Counter Interrupt is enabled.	R	0x1
<i>TIMER_MASK</i>	2	If this bit is set, the local CPU Timer Interrupt is enabled.	R	0x1
<i>COMPARE_MASK</i>	1	If this bit is set, the local GIC Count/Compare Interrupt is enabled.	R	0x1
<i>WD_MASK</i>	0	If this bit is set, the local WatchDog Interrupt is enabled.	R	0x1

**9.6.2.4 Local Interrupt Reset Mask Register**

**Figure 9.27 Local Interrupt Reset Mask Register Format**



**Table 9.45 Local Interrupt Reset Mask Register (GIC\_VPEi\_RMASK — Offset 0x000C)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:7	Writes ignored. Must be written with a value of 0x0.		Undefined
<i>FDC_RMASK</i>	6	Writing a 0x1 to this bit disables the local Fast Debug Channel interrupt	W	Undefined
<i>SWINT1_RMASK</i>	5	Writing a 0x1 to this bit disables the local SWInt1 Interrupt.	W	Undefined
<i>SWINT0_RMASK</i>	4	Writing a 0x1 to this bit disables the local SWInt0 Interrupt.	W	Undefined
<i>PERFCNT_RMASK</i>	3	Writing a 0x1 to this bit disables the local Performance Counter Interrupt.	W	Undefined
<i>TIMER_RMASK</i>	2	Writing a 0x1 to this bit disables the local Timer Interrupt.	W	Undefined
<i>COMPARE_RMASK</i>	1	Writing a 0x1 to this bit disables the local GIC Count/Compare Interrupt.	W	Undefined
<i>WD_RMASK</i>	0	Writing a 0x1 to this bit disables the local WatchDog Timer Interrupt.	W	Undefined

### 9.6.2.5 Local Interrupt Set Mask Register

This is a write-only register. For more information, refer to [Section 9.4.2, "Configuring Interrupt Sources"](#).

**Figure 9.28 Local Interrupt Set Mask Register Format**

31	7	6	5	4	3	2	1	0
R	FDC SMASK	SWINT1_ SMASK	SWINT0_ SMASK	PERFCOUNT_ SMASK	TIMER_ SMASK	COMPARE_ SMASK	WQ_SMASK	

**Table 9.46 Local Interrupt Set Mask Register (GIC\_VPEi\_SMASK — Offset 0x0010)**

Register Fields		Description	Read/ Write	Reset State
Name	Bits			
RESERVED	31:7	Writes ignored. Must be written with a value of 0x0.		Undefined
<i>FDC_SMASK</i>	6	Writing a 0x1 to this bit enables the local Fast Debug Channel Interrupt	W	Undefined
<i>SWINT1_SMASK</i>	5	Writing a 0x1 to this bit enables the local SWInt1 Interrupt.	W	Undefined
<i>SWINT0_SMASK</i>	4	Writing a 0x1 to this bit enables the local SWInt0 Interrupt.	W	Undefined
<i>PERFCNT_SMASK</i>	3	Writing a 0x1 to this bit enables the l	W	Undefined
<i>TIMER_SMASK</i>	2	Writing a 0x1 to this bit enables the local Timer Interrupt.	W	Undefined
<i>COMPARE_SMASK</i>	1	Writing a 0x1 to this bit enables the local GC Count/Compare Interrupt.	W	Undefined
<i>WD_SMASK</i>	0	Writing a 0x1 to this bit enables the local WatchDog Timer Interrupt.	W	Undefined


### 9.6.2.6 Local Map to Pin Registers

This section includes the local map to pin registers for the interrupt types described in [Table 9.47](#). The bit assignments for each of these registers is identical. There is one register per instantiated VPE. The ‘i’ indicates a number between 1 and 8 depending on the number of VPEs in the system.

**Table 9.47 Local Map-to-Pin Register Mapping**

Offset	Acronym	Register Name
0x0040	GIC_VPEi_WD_MAP	Local Watchdog Map-to-Pin register.
0x0044	GIC_VPEi_COMPARE_MAP	Local Counter/Compare Map-to-Pin register.
0x0048	GIC_VPEi_TIMER_MAP	Local Timer Map-to-Pin register.
0x004C	GIC_VPEi_FDC_MAP	Local Fast Debug Channel Map-to-Pin register.
0x0050	GIC_VPEi_PERFCTR_MAP	Local Performance Counter Map-to-Pin register.
0x0054	GIC_VPEi_SWInt0_MAP	Local Software Interrupt 0 Map-to-Pin register.
0x0058	GIC_VPEi_SWInt1_MAP	Local Software Interrupt 1 Map-to-Pin register.

**Figure 9.29 GIC Interrupt Map to Pin Register Format**




**Table 9.48 Local Map to Pin Registers (Offset 0x0040 - 0x0058 — See above for mapping)**

Register Fields		Description	Read/Write	Reset State																
Name	Bits																			
<i>MAP_TO_PIN</i>	31	If this bit is set, this interrupt source is mapped to a VPE interrupt pin (specified by the <i>MAP</i> field below). Only one of the <i>MAP_TO_PIN</i> , <i>MAP_TO_NMI</i> , or <i>MAP_TO_YQ</i> bits can be set at any one time.	-	0x1 for Timer, PerfCount and SWIntx; 0x0 for WatchDog																
<i>MAP_TO_NMI</i>	30	If this bit is set, this interrupt source is mapped to a VPE NMI interrupt pin. Only one of the <i>MAP_TO_PIN</i> , <i>MAP_TO_NMI</i> , or <i>MAP_TO_YQ</i> bits can be set at any one time.	R/W	0x1 for WatchDog; 0x0 for Others																
<i>MAP_TO_YQ</i>	29	If this bit is set, this interrupt source is mapped to an MT Yield Qualifier pin (specified by the <i>MAP</i> field below). Only one of the <i>MAP_TO_PIN</i> , <i>MAP_TO_NMI</i> , or <i>MAP_TO_YQ</i> bits can be set at any one time.	R/W	0																
R	28:6	Read as 0x0. Writes ignored. Must be written with a value of 0x0.		0																
<i>MAP</i>	5:0	<p>When the <i>MAP_TO_PIN</i> bit is set, this field contains the encoded value of interrupts signals <i>Int[62:0]</i>.</p> <p>In EIC mode, this represents one less than the EIC interrupt level (e.g. a value of 0x20 represents interrupt level 21).</p> <p>For non-EIC mode, the value represents the CPU interrupt to be asserted (e.g. a value of 0x03 represents interrupt 3), and only values of 0 to 5 are legal.</p> <p>When <i>MAP_TO_YQ</i> is set, this field contains the encoded signal selection of the Yield Qualifier.</p> <table border="1" style="margin: 10px auto; border-collapse: collapse; text-align: center;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td><i>YR_ysi[0]</i></td> </tr> <tr> <td>0x1</td> <td><i>YR_ysi[1]</i></td> </tr> <tr> <td>...</td> <td></td> </tr> <tr> <td>0xF</td> <td><i>YR_ysi[15]</i></td> </tr> <tr> <td>0x10</td> <td>NULL</td> </tr> <tr> <td>...</td> <td></td> </tr> <tr> <td>0x1F</td> <td>NULL</td> </tr> </tbody> </table> <p>Since YQ is per-CORE rather than per-VPE, software needs to apply proper protection across VPEs by using, for example, the cop0 <i>YQMask</i> register.</p>	Encoding	Meaning	0x0	<i>YR_ysi[0]</i>	0x1	<i>YR_ysi[1]</i>	...		0xF	<i>YR_ysi[15]</i>	0x10	NULL	...		0x1F	NULL	W	0
Encoding	Meaning																			
0x0	<i>YR_ysi[0]</i>																			
0x1	<i>YR_ysi[1]</i>																			
...																				
0xF	<i>YR_ysi[15]</i>																			
0x10	NULL																			
...																				
0x1F	NULL																			

**9.6.2.7 VPE-Other Addressing Register**

This register must be written with the correct value before accessing the VPE-Other address section.

**Figure 9.30 VPE-Other Addressing Register Format**



**Table 9.49 VPE-Other Addressing Register (GIC\_VPEi\_OTHER\_ADDRESS — Offset 0x0080)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
R	31:16	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0
<i>VPENUM</i>	15:0	Number of the register set to be accessed in the VPE-Other address space.	R/W	0

### 9.6.2.8 VPE-Local Identification Register

The aliased memory scheme is normally invisible to software when accessing GIC registers within the VPE-Local Control Block. What actually happens is that an offset is used to make a subset of the GIC registers appear in the VPE-Local addressing Window.

This register reports the VPE number that is used as the addressing offset for the VPE-Local Control Block.

**Figure 9.31 VPE-Local Addressing Register Format**



**Table 9.50 VPE-Local Identification Register (GIC\_VPEi\_IDENT — Offset 0x0088)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>VPENUM</i>	31:0	This number is used as an index to the registers within the GIC when accessing the VPE-local control block for this VPE.	R	-

## 9.6.3 Local Timer Register Descriptions

### 9.6.3.1 Watchdog Timer Config Register

For more information on the usage of this register, refer to [Section 9.4.6.2, "GIC Watchdog Timer"](#).

**Figure 9.32 Watchdog Timer Config Register Format**

31	8	7	6	5	4	3	1	0
R	WDRESET	WDINTR	WAIT	DEBUG	TYPE	WDSTART		

**Table 9.51 Watchdog Timer Config Register (GIC\_VPEi\_WD\_CONFIG — Offset 0x0090)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
R	31:8	Read as 0x0. Writes ignored. Must be written with a value of 0x0.		0
<i>WDRESET</i>	7	Status bit which indicates that a Watchdog was responsible for resetting the interAptiv MPS. A write of 0x1 to this bit of this register automatically clears this bit. This bit needs to survive a watchdog triggered reset.	R/WC	0
<i>WDINTR</i>	6	Status bit which indicates that a Watchdog was responsible for generating this interrupt. A write of 0x1 to this bit automatically clears the bit. Typically this interrupt is routed to the <i>NMI</i> interrupt input of the VPE, but could be routed to another interrupt as well.	R/WC	Undefined
<i>WAIT</i>	5	Stop countdown if the VPE is in an implementation-defined low power mode (including the mode which is entered on a <i>WAIT</i> instruction). 0x0 - Stop countdown if VPE is in low power mode. 0x1 - Low power mode has no effect on countdown.	R/W	0
<i>DEBUG</i>	4	Stop countdown if the VPE is in debug mode. 0x0 - Stop countdown if VPE is in Debug Mode (CP0 <i>DEBUG<sub>DM</sub></i> bit is set). 0x1 - Debug Mode has no effect on countdown.	R/W	0

**Table 9.51 Watchdog Timer Config Register (GIC\_VPEi\_WD\_CONFIG — Offset 0x0090)(continued)**

Register Fields		Description	Read/Write	Reset State										
Name	Bits													
<i>TYPE</i>	3:1	<p>There are <i>three</i> ways to setup the watchdog timer:</p> <ol style="list-style-type: none"> <li>1. It can be setup such that if it decrements to 0x0, it causes an interrupt and then stops.</li> <li>2. It can be setup such that after the first countdown, the initial value is reloaded and the countdown continues. If on the second trip, the counter reaches 0x0, a reset is broadcast to all VPEs in the system.</li> <li>3. The counter can be used as a Programmable Interval Timer (PIT).</li> </ol> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>WD One Trip Mode. This asserts an interrupt, typically an NMI, and stops.</td> </tr> <tr> <td>0x1</td> <td>WD Second Countdown Mode. This asserts <i>SL_Reset</i> on all VPEs.</td> </tr> <tr> <td>0x2</td> <td>PIT Mode. This asserts an interrupt and reloads and keeps going.</td> </tr> <tr> <td>0x3..0x7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	0	WD One Trip Mode. This asserts an interrupt, typically an NMI, and stops.	0x1	WD Second Countdown Mode. This asserts <i>SL_Reset</i> on all VPEs.	0x2	PIT Mode. This asserts an interrupt and reloads and keeps going.	0x3..0x7	Reserved	R/W	0
Encoding	Meaning													
0	WD One Trip Mode. This asserts an interrupt, typically an NMI, and stops.													
0x1	WD Second Countdown Mode. This asserts <i>SL_Reset</i> on all VPEs.													
0x2	PIT Mode. This asserts an interrupt and reloads and keeps going.													
0x3..0x7	Reserved													
WD_START	0	<p>Watchdog timer start/stop. Setting this bit starts the Watchdog timer, while clearing the bit stops the timer.</p> <p>0 - Stop the Watchdog timer 1 - Reload the initial count and start the Watchdog timer.</p>	R/W	0										

### 9.6.3.2 Watchdog Timer Count Register

For more information on the usage of this register, refer to [Section 9.4.7.3, "Watchdog Timer Interrupts"](#).

**Figure 9.33 Watchdog Timer Count Register Format**




**Table 9.52 Watchdog Timer Count Register (GIC\_VPEi\_WD\_COUNT — Offset 0x0094)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>COUNT</i>	31:0	This read-only register indicates the state of the decrementing counter. The width of the counter is 32 bits.	R	Undefined

### 9.6.3.3 Watchdog Timer Initial Count Register

For more information on the usage of this register, refer to [Section 9.4.7.3, "Watchdog Timer Interrupts"](#).

**Figure 9.34 Watchdog Timer Initial Count Register Format**



**Table 9.53 Watchdog Timer Initial Count Register (GIC\_VPEi\_WD\_INITIAL — Offset 0x0098)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>INIT</i>	31:0	Initial value to be loaded into the Watchdog counter. Needs to be done with the counter disabled; otherwise, the results are UNPREDICTABLE.	R/W	Undefined

### 9.6.3.4 CompareLo Register

For more information on the usage of this register, refer to [Section 9.4.7.3, "Watchdog Timer Interrupts"](#).

**Figure 9.35 CompareLo Register Format**



**Table 9.54 CompareLo Register (GIC\_VPEi\_CompareLo — Offset 0x00A0)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>COMPARELO</i>	31:0	When the contents of <i>GIC_VPEi_CompareLo</i> and <i>GIC_VPEi_CompareHi</i> registers match the contents of <i>GIC_SH_CounterLo</i> and <i>GIC_SH_CounterHi</i> , the <i>VPEi_Compare</i> interrupt is triggered. This registered interrupt can only be deasserted by writing either the <i>GIC_VPEi_CompareLo</i> or <i>GIC_VPEi_CompareHi</i> registers.	R/W	0xFFFF_FFFF

### 9.6.3.5 VPE-Local CompareHi Register

For more information on the usage of this register, refer to [Section 9.4.7.3, "Watchdog Timer Interrupts"](#).

**Figure 9.36 CompareHi Register Format**





**Table 9.55 VPE-Local CompareHi Register (GIC\_VPEi\_CompareHi — Offset 0x00A4)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>COMPAREHI</i>	31:0	See description for <i>GIC_VPEi_CompareLo</i> . The width of this register matches the width of <i>GIC_SH_COUNTER</i> .	R/W	All instantiated bits = 0x1

### 9.6.3.6 Local EIC Shadow Set Registers

These registers are only instantiated if the GIC is configured to include EIC mode. There are 64 EIC Shadow registers located at offset addresses 0x0100 - 0x01FC. These registers are mapped as follows:

**Table 9.56 Local Map-to-Pin Register Mapping**

Offset	Interrupt Source	Register Acronym	Offset	Interrupt Source	Acronym
0x0100	0	GIC_VPEi_EICSRCSS0	0x0180	32	GIC_VPEi_EICSRCSS32
0x0104	1	GIC_VPEi_EICSRCSS1	0x0184	33	GIC_VPEi_EICSRCSS33
0x0108	2	GIC_VPEi_EICSRCSS2	0x0188	34	GIC_VPEi_EICSRCSS34
0x010C	3	GIC_VPEi_EICSRCSS3	0x018C	35	GIC_VPEi_EICSRCSS35
0x0110	4	GIC_VPEi_EICSRCSS4	0x0190	36	GIC_VPEi_EICSRCSS36
0x0114	5	GIC_VPEi_EICSRCSS5	0x0194	37	GIC_VPEi_EICSRCSS37
0x0118	6	GIC_VPEi_EICSRCSS6	0x0198	38	GIC_VPEi_EICSRCSS38
0x011C	7	GIC_VPEi_EICSRCSS7	0x019C	39	GIC_VPEi_EICSRCSS39
0x0120	8	GIC_VPEi_EICSRCSS8	0x01A0	40	GIC_VPEi_EICSRCSS40
0x0124	9	GIC_VPEi_EICSRCSS9	0x01A4	41	GIC_VPEi_EICSRCSS41
0x0128	10	GIC_VPEi_EICSRCSS10	0x01A8	42	GIC_VPEi_EICSRCSS42
0x012C	11	GIC_VPEi_EICSRCSS11	0x01AC	43	GIC_VPEi_EICSRCSS43
0x0130	12	GIC_VPEi_EICSRCSS12	0x01B0	44	GIC_VPEi_EICSRCSS44
0x0134	13	GIC_VPEi_EICSRCSS13	0x01B4	45	GIC_VPEi_EICSRCSS45
0x0138	14	GIC_VPEi_EICSRCSS14	0x01B8	46	GIC_VPEi_EICSRCSS46
0x013C	15	GIC_VPEi_EICSRCSS15	0x01BC	47	GIC_VPEi_EICSRCSS47
0x0140	16	GIC_VPEi_EICSRCSS16	0x01C0	48	GIC_VPEi_EICSRCSS48
0x0144	17	GIC_VPEi_EICSRCSS17	0x01C4	49	GIC_VPEi_EICSRCSS49
0x0148	18	GIC_VPEi_EICSRCSS18	0x01C8	50	GIC_VPEi_EICSRCSS50
0x014C	19	GIC_VPEi_EICSRCSS19	0x01CC	51	GIC_VPEi_EICSRCSS51
0x0150	20	GIC_VPEi_EICSRCSS20	0x01D0	52	GIC_VPEi_EICSRCSS52
0x0154	21	GIC_VPEi_EICSRCSS21	0x01D4	53	GIC_VPEi_EICSRCSS53
0x0158	22	GIC_VPEi_EICSRCSS22	0x01D8	54	GIC_VPEi_EICSRCSS54
0x015C	23	GIC_VPEi_EICSRCSS23	0x01DC	55	GIC_VPEi_EICSRCSS55
0x0160	24	GIC_VPEi_EICSRCSS24	0x01E0	56	GIC_VPEi_EICSRCSS56
0x0164	25	GIC_VPEi_EICSRCSS25	0x01E4	57	GIC_VPEi_EICSRCSS57
0x0168	26	GIC_VPEi_EICSRCSS26	0x01E8	58	GIC_VPEi_EICSRCSS58

**Table 9.56 Local Map-to-Pin Register Mapping**

Offset	Interrupt Source	Register Acronym	Offset	Interrupt Source	Acronym
0x016C	27	GIC_VPEi_EICSRCSS27	0x01EC	59	GIC_VPEi_EICSRCSS59
0x0170	28	GIC_VPEi_EICSRCSS28	0x01F0	60	GIC_VPEi_EICSRCSS60
0x0174	29	GIC_VPEi_EICSRCSS29	0x01F4	61	GIC_VPEi_EICSRCSS61
0x0178	30	GIC_VPEi_EICSRCSS30	0x01F8	62	GIC_VPEi_EICSRCSS62
0x017C	31	GIC_VPEi_EICSRCSS31	0x01FC	63	GIC_VPEi_EICSRCSS63

**Figure 9.37 Local EIC Shadow Set Register Format**




**Table 9.57 Local EIC Shadow Set Registers (GIC\_VPEi\_EICSSi — Offset 0x0100 - 0x01FC)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
R	31:4	Read as 0x0. Writes ignored. Must be written with a value of 0x0.	-	0
<i>EIC_SSn</i>	3:0	Encoded value that indicates the Shadow Set number to use for this particular interrupt.	R/W	Undefined

### 9.6.3.7 VPE-Local DINT Group Participate Register

When bit 0 of this register is set, the local VPE monitors the state of the DINT\_Send\_to\_Group register in the Shared register set, as well as the EJ\_DINT\_IN pin for debug activity. Refer to [Section 9.4.10, "Debug Interrupt Generation"](#) for more information.

**Figure 9.38 VPE-Local EIC DINT Group Participate Register Format**



**Table 9.58 VPE-Local DINT Group Participate Register (GIC\_Vx\_DINT\_PART — Offset 0x3000)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
RESERVED	31:1	Reads as 0x0. Writes ignored. Must be written with a value of 0x0.	R	0x0

**Table 9.58 VPE-Local DINT Group Participate Register (GIC\_Vx\_DINT\_PART — Offset 0x3000)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>DINT_GP</i>	0	<p>If this bit is set, the local VPE pays attention to the <i>DINT_Send_to_Group</i> register as well as the external <i>EJ_DINT_IN</i> signal pin.</p> <p>For this case, when the <i>Send_DINT</i> bit within the <i>DINT_Send_to_Group</i> register is asserted (or the external <i>EJ_DINT_IN</i> signal is asserted), the local VPE will have its <i>EJ_DINT</i> or <i>EJ_DINT_1</i> signal asserted.</p> <p>If this bit is clear, the local VPE is not affected by the <i>DINT_Send_to_Group</i> register nor the external <i>EJ_DINT_IN</i> pin signal.</p>	R/W	0x1

See [Chapter 17, “Multi-CPU Debug” on page 745](#) for more information about how this register is used.

### 9.6.3.8 VPE-Local DebugBreak Group Register

When the local VPE enters Debug Mode (denoted by the local *EJTAG\_TAP.DebugM* bit being asserted), this register defines which other VPEs in the system will subsequently also receive a Debug Interrupt. This allows multiple VPEs to be synchronized to a single software debugger by entering debug mode somewhat simultaneously.

**Figure 9.39 VPE-Local EIC DINT Group Participate Register Format**



**Table 9.59 VPE-Local DebugBreak Group Register (GIC\_Cx\_BRK\_GROUP — Offset 0x3080)**

Register Fields		Description	Read/Write	Reset State
Name	Bits			
<i>JOIN_DB</i>	31:0	<p>Each bit in this register represents a VPE in the system.</p> <p>If the bit is set, the corresponding VPE will have its <i>EJ_DINT</i> or <i>EJ_DINT_1</i> signal asserted when the local VPE enters Debug Mode.</p> <p>If the bit is clear, the corresponding VPE is not affected when the VPE enters Debug Mode.</p> <p>The bit which represents the local VPE cannot be used to disable Debug Mode for the local VPE. For example, if the local VPE is represented by bit <i>i</i>, clearing bit <i>i</i> will NOT disable Debug Mode for the local VPE.</p>	R/W	All zeros

See [Chapter 17, “Multi-CPU Debug” on page 745](#) for more information about how this register is used.

## 9.7 GIC User-Mode Visible Section

The Shared, VPE-local, and VPE-other sections are meant to be located in privileged system virtual address space, in which only kernel mode software can initialize and update the interrupt controller.

A separate 64KB address space is allocated so that it may be mapped to user-mode virtual address space. Within this address space are aliases for GIC registers that are read so often that it makes sense to make them available to user-mode programs without requiring a system call. The aliases for these registers are read-only. Currently, the only registers that are aliased into this space are the shared Counter registers.

The addresses for the registers within the User-Mode Visible Section of the GIC are calculated as follows:

$$\text{SharedSection\_Register\_Physical\_Address} = \text{GIC\_baseaddress} + \text{UMVisible\_Section\_baseoffset} + \text{Register\_Offset}$$

**Table 9.60 User-Mode Visible Section Register Map**

Register Offset	Name	Type	Description
0x0000	GIC CounterLo (GIC_SH_CounterLo)	R	Read-only alias for GIC Shared CounterLo.
0x0004	GIC CounterHi (GIC_SH_CounterHi)	R	Read-only alias for GIC Shared CounterHi.
Any Other Offsets	Reserved		Reserved for future extensions.

## 9.8 GIC Timing Diagrams

The following GIC timing diagrams are divided into Case 1 and Case 2.

### 9.8.1 Case 1: GIC Handling of Multiple Active Interrupts Assigned to the Same VPE

Case 1 describes four sub-cases for various settings of edge vs. level interrupts in non-EIC and EIC modes. [Table 9.61](#) lists the settings of relevant GIC registers for the four sub-cases described in the following subsections.

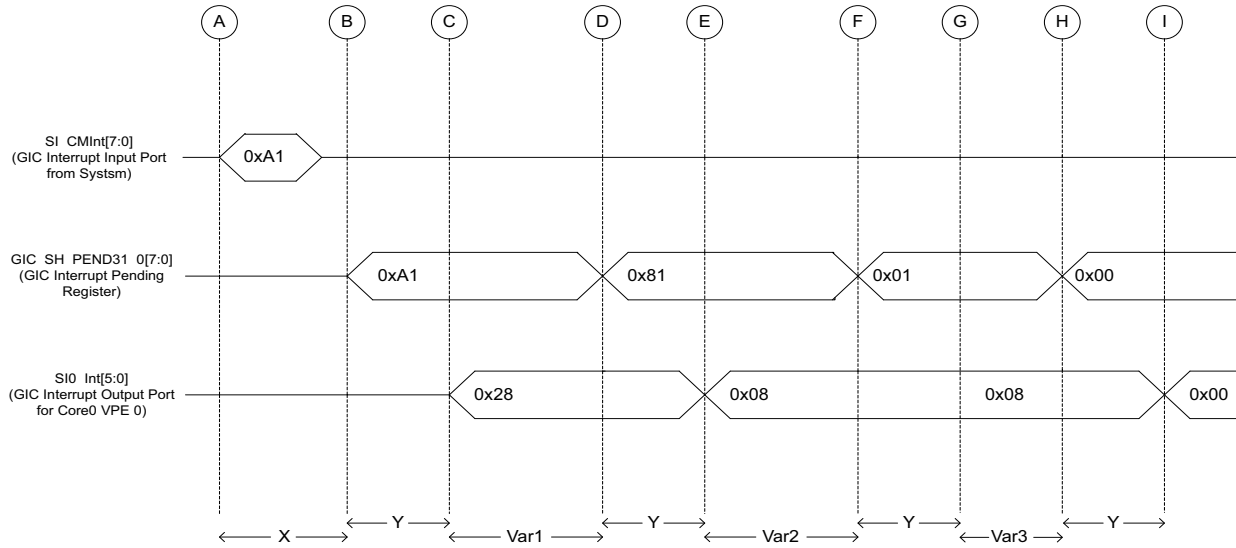
**Table 9.61 GIC Register Settings for Cases 1.1 — 1.4**

Register Name	Case 1.1	Case 1.2	Case 1.3	Case 1.4
GIC_SH_POL31_0	0x000000A1	0x000000A1	0x000000A1	0x000000A1
GIC_SH_TRIG31_0	0x000000A1	0x000000A1	0x00000000	0x00000000
GIC_SH_DUAL31_0	0x00000000	0x00000000	Don't care	Don't care
GIC_SH_MASK31_0	0x000000A1	0x000000A1	0x000000A1	0x000000A1
GIC_SH_MAP0_PIN	0x80000003	0x80000016	0x80000003	0x80000016
GIC_SH_MAP5_PIN	0x80000005	0x8000003E	0x80000005	0x8000003E
GIC_SH_MAP7_PIN	0x80000003	0x80000000	0x80000003	0x80000000
GIC_SH_MAP0_VPE31_0	0x00000001	0x00000002	0x00000004	0x00000008
GIC_SH_MAP5_VPE31_0				
GIC_SH_MAP7_VPE31_0				

### 9.8.1.1 Case 1.1: Multiple Positive Edge-Triggered Interface Routed to the Same VPE in non-EIC Mode

Figure 9.40 shows case 1.1, along with an explanation of what happens at points A through I.

Figure 9.40 Case 1.1 Timing Diagram



X = Cycles required to synchronise and edge detect external interrupts  
 Y = Cycles required to go through interrupt routing stages within GIC  
 Var1, Var2, Var3 = Variable number of cycles required for software to handle and clear each interrupt.

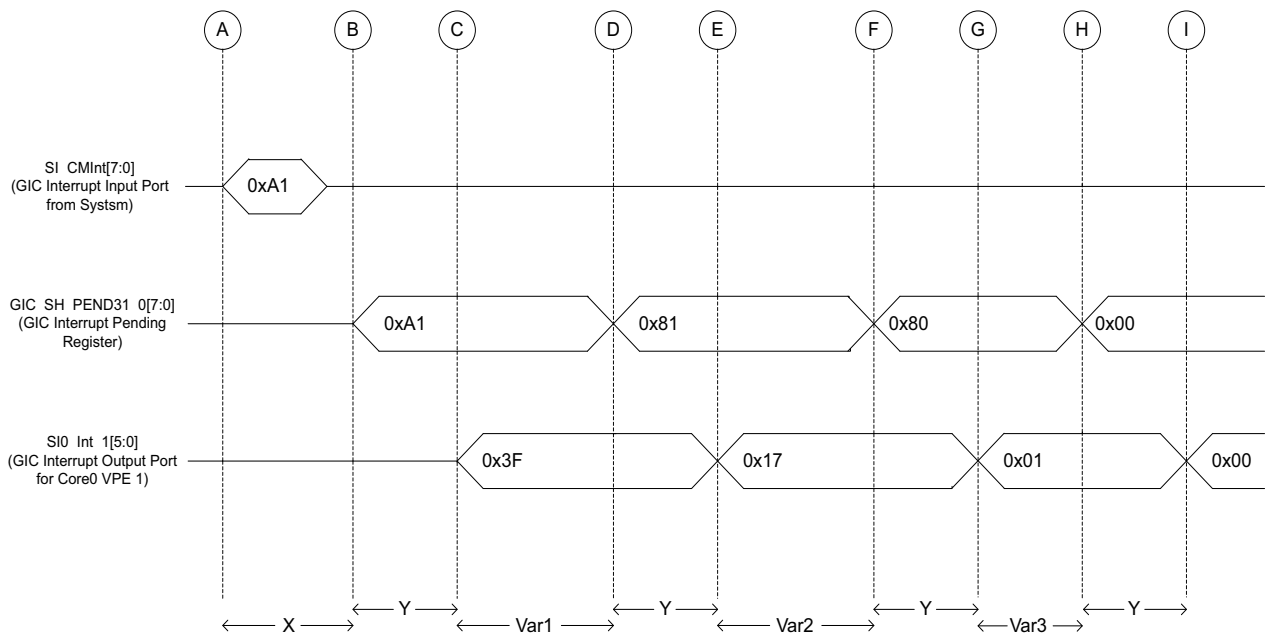
- A. Simultaneous positive edge triggered interrupts generated on external ports 0, 5 and 7.
- B. All three simultaneous interrupt events are captured in GIC pending register *GIC\_SH\_PEND*.
- C. The three external interrupt events are mapped to two hardware interrupt pins 3 and 5 of the VPE according to the *GIC\_SH\_MAP<0,5,7>\_PIN* register settings for case 1.1 in Table 9.61. So the hardware interrupts 3 and 5 appear in Core 0 VPE 0 interrupt bus as 0x28.
- D. The VPE interrupt handling logic prioritizes hardware interrupt pin 5 over pin 3. The interrupt handler software for VPE's hardware interrupt pin 5 clears the corresponding interrupt source from external interrupt port 5.
- E. As a result of external interrupt 5 being cleared, the corresponding hardware interrupt pin 5 of the VPE also gets cleared as there are no more external interrupts mapped to hardware interrupt pin 5 of the VPE.
- F. The interrupt handler software for VPE's hardware interrupt pin 3 clears the corresponding interrupt source from external interrupt port 7. Since both external interrupt sources 0 and 7 are mapped to the same hardware interrupt pin 3, it is the choice of that interrupt handler software to prioritize between external interrupt sources 0 and 7 and in this example, the external interrupt source 7 is chosen first.

- G. The VPE's hardware interrupt pin 3 still stays asserted as the other mapped external interrupt source to it is still not cleared.
- H. The interrupt handler software for VPE's hardware interrupt pin 3 clears the other corresponding interrupt source from external interrupt port 0.
- I. As a result of external interrupts 0 and 7 being cleared, their corresponding hardware interrupt pin 3 of the VPE also gets cleared as there are no more external interrupts mapped to hardware interrupt pin 3 of the VPE.

**9.8.1.2 Case 1.2: Multiple Positive Edge Triggered Interrupts Routed to the Same VPE in EIC Mode**

Figure 9.41 shows case 1.2, along with an explanation of what happens at points A through I.

**Figure 9.41 Case 1.2 Timing Diagram**



X = Cycles required to synchronise and edge detect external interrupts  
 Y = Cycles required to go through interrupt routing stages within GIC  
 Var1, Var2, Var3 = Variable number of cycles required for software to handle and clear each interrupt.

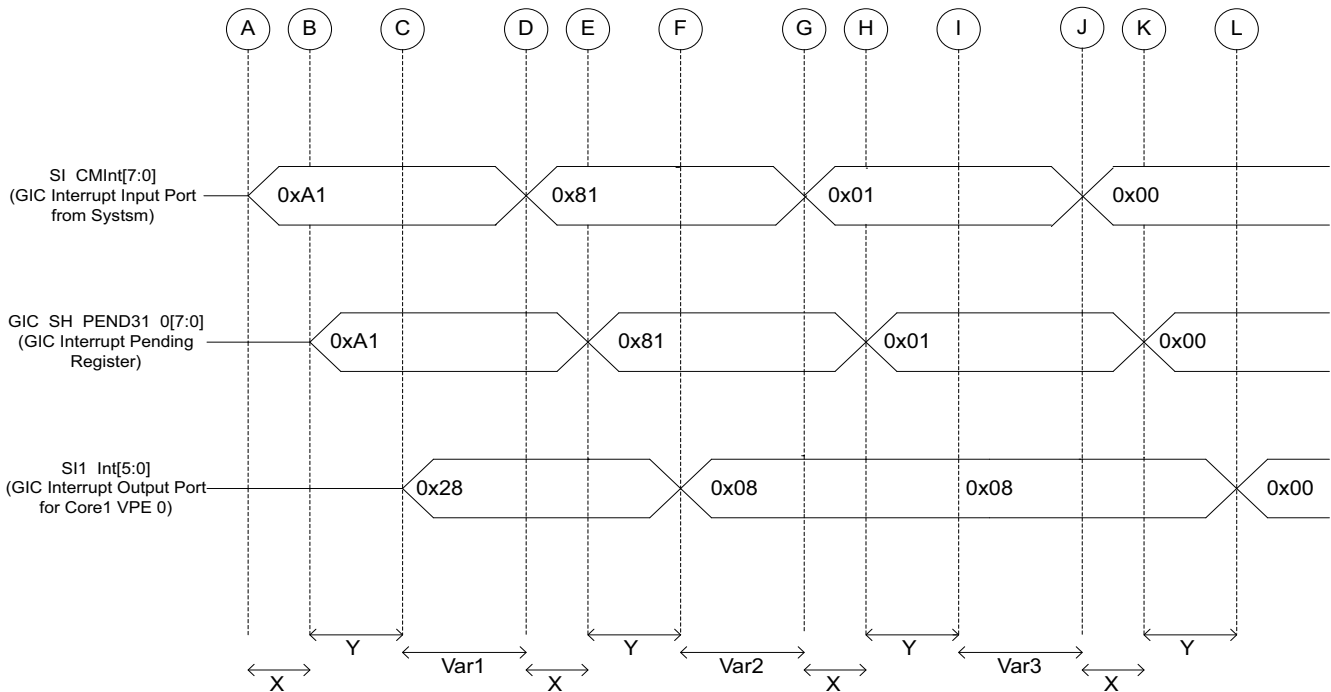
- A. Same description as case 1.1 point A.
- B. Same description as case 1.1 point B.
- C. According to the *GIC\_SH\_MAP<0,5,7>\_PIN* register settings for case 1.2 in Table 9.61, the highest Requesting Interrupt Priority Level (RIPL) assigned is 0x3F (*GIC\_SH\_MAP<5>\_PIN.MAP + 1*). The GIC prioritizes all the active RIPLs for a VPE and sends out the highest RIPL on to VPE's interrupt bus. So the highest RIPL of 0x3F is sent out first on Core 0 VPE 1 interrupt bus as all 3 active interrupts in this case are mapped to VPE 1.

- D. The interrupt handler software associated with RIPL 0x3F clears the corresponding interrupt source from external interrupt port 5.
- E. As a result of external interrupt 5 being cleared, the currently active highest RIPL of 0x3F gets cleared and therefore GIC send out the next highest RIPL 0x17 on VPE's interrupt bus.
- F. The interrupt handler software associated with RIPL 0x17 clears the corresponding interrupt source from external interrupt port 0.
- G. As a result of external interrupt 0 being cleared, the currently active highest RIPL of 0x17 gets cleared and therefore GIC send out the next highest RIPL 0x01 on VPE's interrupt bus.
- H. The interrupt handler software associated with RIPL 0x01 clears the corresponding interrupt source from external interrupt port 7.
- I. As all active interrupts are now cleared, the VPE's interrupt bus goes back to 0x00 which indicates no active interrupt events targeted of that VPE.

**9.8.1.3 Case 1.3 - Multiple Active High-Level Interrupts Routed to the Same VPE in non-EIC Mode**

Figure 9.42 shows case 1.3, along with an explanation of what happens at points A through L.

**Figure 9.42 Case 1.3 Timing Diagram**



X = Cycles required to synchronise external interrupts

Y = Cycles required to go through interrupt routing stages within GIC

Var1, Var2, Var3 = Variable number of cycles required for software to handle and clear each interrupt.

The timing diagram descriptions should be similar to case 1.1. The only difference with level vs edge interrupts is that the *GIC\_SH\_PEND* register would follow the state of the external interrupt ports as shown in points A to B, D to E, G

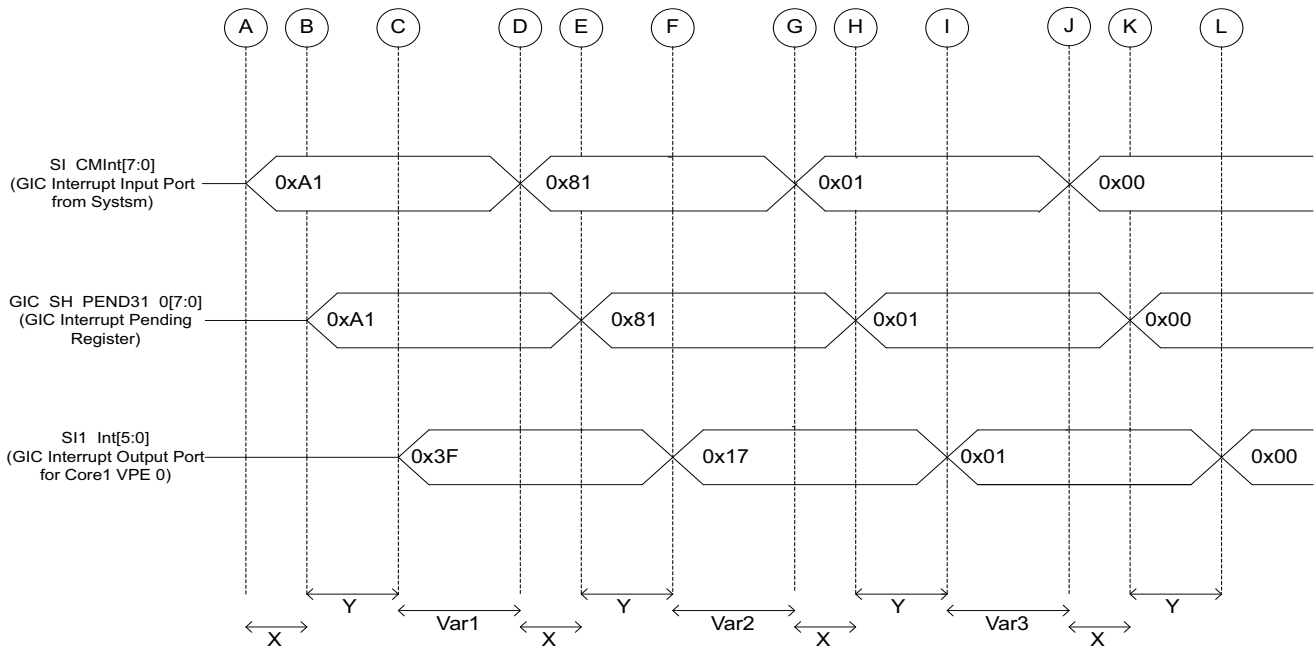


to H and J to K. Also note that the level interrupt clearing would happen at the external interrupt source level and not at the GIC register level as per edge interrupts. These are shown at points D, G and J.

### 9.8.1.4 Case 1.4: Multiple Active High-Level Interrupts Routed to the Same VPE in non-EIC Mode

Figure 9.43 shows case 1.4, along with an explanation of what happens at points A through L.

**Figure 9.43 Case 1.4 Timing Diagram**



X = Cycles required to synchronise external interrupts  
 Y = Cycles required to go through interrupt routing stages within GIC  
 Var1, Var2, Var3 = Variable number of cycles required for software to handle and clear each interrupt.

The timing diagram descriptions would be similar to cases 1.2 and 1.3 for this.

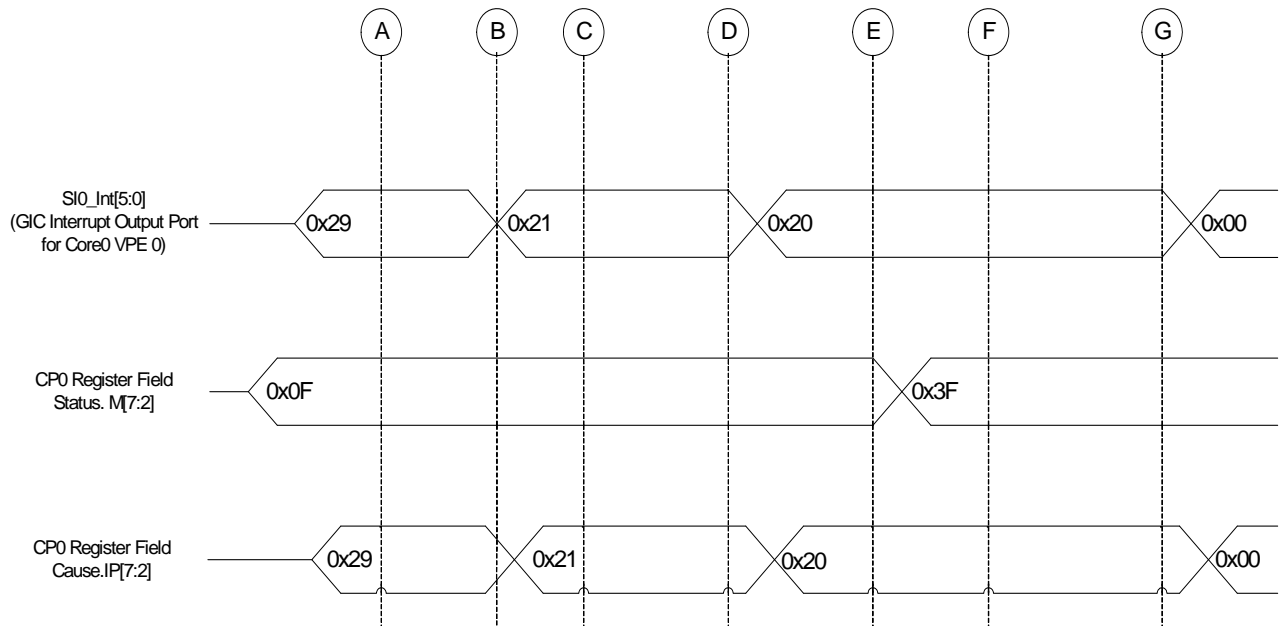
## 9.8.2 Case 2: GIC Passing Interrupts to VPE

This section describes the GIC passing interrupts to the VPE in both EIC and non-EIC modes.

### 9.8.2.1 Case 2.1: VPE Taking Interrupts in non-EIC Mode

Figure 9.44 shows case 2.1, along with an explanation of what happens at points A through G.

**Figure 9.44 Case 2.1 Timing Diagram**




In non-EIC mode, the CP0 *Status.IM[7:2]* register fields perform masking on hardware interrupts 0 to 5. The CP0 *Cause.IP[7:2]* register fields capture the state of the VPE interrupt input port.

- A. Hardware interrupt 3 is taken at this point as it's the highest non-masked active interrupt. The active hardware interrupt 5 is not taken here as its disabled with *Status.IP[7] = 0*.
- B. Finished handling interrupt event from hardware interrupt 3.
- C. Hardware interrupt 0 is taken at this point as it's the next highest non-masked active interrupt.
- D. Finished handling interrupt event from hardware interrupt 0.
- E. Software enables hardware interrupts 4 and 5 as well with *Status.IP[7:2] = 0x3F*.
- F. Now the hardware interrupt 5 is taken as it's the only and highest active interrupt.
- G. Finished handling interrupt event from hardware interrupt 5.

### 9.8.2.2 VPE Taking Interrupts in EIC Mode

Figure 9.45 shows case 2.2, along with an explanation of what happens at points A through G.

**Figure 9.45 Case 2.2 Timing Diagram**



In EIC mode, the Requested Interrupt Priority Level (RIPL) will be taken by the VPE if the RIPL is strictly greater than the CP0 Status.IPL register field value. The CP0 Cause.RIPL register field value will hold the last taken RIPL value for the VPE.

- A. VPE takes the RIPL value 0x39 as its greater than the Status.IPL value of 0x10. The Cause.RIPL is loaded with value 0x39 when that RIPL is taken.
- B. Finished handling interrupt event for RIPL 0x39.
- C. VPE takes the RIPL value 0x11 as its greater than the Status.IPL value of 0x10. Again the Cause.RIPL is loaded accordingly.
- D. Finished handling interrupt event for RIPL 0x11.
- E. Software lowers the Status.IPL value to 0x0F.
- F. Now only the VPE takes the RIPL value 0x10 as it was blocked before the Status.IPL was lowered in step E. Again the Cause.RIPL is loaded accordingly.
- G. Finished handling interrupt event for RIPL 0x10. But the next active RIPL value of 0x09 gets blocked as it's less than Status.IPL value of 0x0F.



## Policy Manager

The Policy Manager (PM) is tasked with giving longer-term hints to the Dispatch Scheduler so as to achieve whatever performance allocation is desired in the system. The Policy Manager will be external to the interAptiv core. MIPS provides a variety of the policy managers depending on the desired performance as described in [Section 10.2, "Policy Managers"](#). In addition, the customer may design their own Policy Manager.

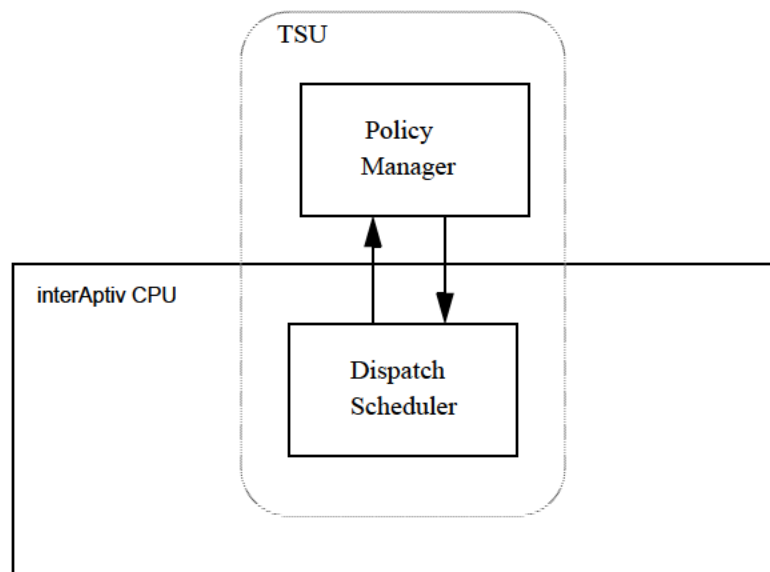
This chapter contains the following sections.

- [Section 10.1, "Thread Scheduling Unit" on page 529](#)
- [Section 10.2, "Policy Managers" on page 530](#)

### 10.1 Thread Scheduling Unit

The interAptiv core contains a unit called the Thread Scheduling Unit (TSU), which has two submodules: an internal Dispatch Scheduler and an external Policy Manager.

**Figure 10.1 TSU Block Diagram**



The Dispatch Scheduler (DS) will make cycle-by-cycle choices on which instructions to issue/dispatch. Since it is internal to the processor core, it cannot be modified by the customer. The DS is designed to be as simple as possible and the system-specific complexity should be put into the Policy Manager so as not to burden the core with extra area/power which is not needed in all configurations.

The Policy Manager Interface incorporates a 4-level priority scheme. Each TC is assigned to one of 4 groups and each group contains a unique priority level. Each cycle the dispatch scheduler chooses to dispatch an instruction from a TC in the highest priority group that contains any runnable TCs. If there are multiple TCs in the selected group, then it chooses among them using a round-robin algorithm.

## 10.2 Policy Managers

MIPS provides the following reference Policy Manager (PM) designs:

- Basic Round Robin (RR)
- Weighted Round Robin (WRR)
- Enhanced Weighted Round Robin (WRR2)

These designs support thread-scheduling capabilities that are common to many systems. For more advanced and/or system-specific capabilities, users can also implement a custom policy manager. The following subsections describe the operation of each of the MIPS-supplied policy managers and the CP0 registers through which they are controlled.

### 10.2.1 Basic Round-Robin Policy Manager

When using the basic round robin PM, all TCs are assigned to the same priority level. Since the internal Dispatch Scheduler implements a simple round-robin among TCs in the same priority level, all TCs are statically given the same weight and bandwidth, and will be fairly allocated amongst all runnable TCs.

This PM does not implement any thread-scheduling CP0 registers. Writes to these registers are ignored. Reads from these registers return -1.

When a new TC is forked, it will begin to participate in the round robin pool. This causes the older TCs to get lower bandwidth allocations.

### 10.2.2 Weighted Round-Robin Policy Manager (WRR)

The main difference between the basic round-robin policy and the weighted round-robin policy manager is software controllability. With the WRR PM, TCs are scheduled round-robin style, but bandwidth given to an individual TC can be adjusted or “weighted” by software so that a TC can get more or less than its fair share of the processor bandwidth.

The WRR PM implements the following CP0 register fields:

- *TCSchedule STP* and *GRP* fields
- *VPESchedule GPO* field
- *TCScheFBack* register.

The WRR PM does not implement the *VPEScheFBack* register.

The group rotation schedule will be implemented. See details in [Section 10.2.4 “Group Rotation Schedule”](#). When a new TC is forked, the GRP of the new TC will be set to be the same as that of its parent.

### 10.2.3 Enhanced Weighted Round-Robin Policy Manager (WRR2)

Internal to the CPU, there are three buffer structures which are shared by all TCs:

- The Load Queue (LDQ). This structure is used for any outstanding load instruction.
- The Fill-Store Buffer (FSB). This structure is used for any in-progress D-cache refills. One FSB entry is generally allocated for each bus transaction.
- The Writeback Buffer (WBB). This structure is used for in-progress D-cache writebacks.

When one of these buffers becomes full, it stalls the pipeline and all TCs. In order to prevent such stalls, the Enhanced Weighted Round-Robin policy manager (WRR2) can automatically deprioritize or throttle non-critical threads when one of these structures gets close to full.

Other than the throttling function, the WRR2 functions the same as the base WRR PM. Please refer to [Section 10.2.2 “Weighted Round-Robin Policy Manager \(WRR\)”](#) for information on the base functionality.

### 10.2.3.1 Throttle Functionality and Operation

There are two programmable throttling functions, `throttle0` and `throttle1`. The functions are as follows:

- Each throttle can be separately enabled for each queue/buffer.
- Software can set the threshold for enabling the throttle function.
- Software can set the group and stop priority override for each throttle.
- If both throttles are activated at the same time, `throttle0` takes priority. Therefore, if both throttles are armed, `throttle0` should generally be the more restrictive one.

The function for each throttle is: When an enabled queue falls below the programmed threshold, the throttle is activated and the effective group and stop priority is overridden with the values from the throttle. When the queue goes back above the threshold, the throttle is deactivated and the group and stop priority return to the values as programmed in `TCSchedule.GRP` and `TCSchedule.STP`, respectively.

NOTE: The throttle function is *armed* when software sets any of the queue enable bits. The throttle is *activated* by hardware dynamically without software intervention.

This function can also be used to boost priority for threads. This is especially useful for the `PM_sys_avail` input - a TC which, when running, helps to decrease the usage of such a resource could be boosted when the resource gets too full.

## 10.2.4 Group Rotation Schedule

When `VPESchedule.GPO` is cleared, the group priorities are rotated as described in this section. These rotations will enable a TC in a higher group to be prioritized higher than TCs in lower groups. The exact weighting function is complicated due to the interplay of non-runnable TCs, but generally, TCs in the next higher group will get at least twice the bandwidth of the TCs in the lower group.

The rotation schedule is described in [Table 10.1](#).

**Table 10.1 Rotation of Group Priority Levels**

Rotation Count	Group3 Priority	Group2 Priority	Group1 Priority	Group0 Priority
4'b0001	P3	P2	P1	P0
4'b0010	P2	P3	P0	P1
4'b0011	P3	P2	P1	P0

**Table 10.1 Rotation of Group Priority Levels**

Rotation Count	Group3 Priority	Group2 Priority	Group1 Priority	Group0 Priority
4'b0100	P2	P0	P3	P1
4'b0101	P3	P2	P1	P0
4'b0110	P2	P3	P0	P1
4'b0111	P3	P1	P2	P0
4'b1000	P0	P2	P1	P3
4'b1001	P3	P1	P2	P0
4'b1010	P2	P3	P0	P1
4'b1011	P3	P2	P1	P0
4'b1100	P2	P0	P3	P1
4'b1101	P3	P2	P1	P0
4'b1110	P2	P3	P0	P1
4'b1111	P3	P1	P2	P0

The group priorities can easily be generated using a 4b counter and a priority encoder.

```
G3_priority = { Cnt [0] | Cnt [1] | Cnt [2], Cnt [0] };

G2_priority =
{ (~Cnt [3] & ~Cnt [2] | ~Cnt [2] & Cnt [1] | Cnt [1] & ~Cnt [0] | ~Cnt [1] & (Cnt [2] ^ ~Cnt [0])),
  (Cnt [2] & Cnt [1] | Cnt [1] & ~Cnt [0] | Cnt [3] & ~Cnt [2] & ~Cnt [1] & Cnt [0]) };

```

The priority values for the other groups can easily be calculated from the above as follows:

```
G1_priority = ~G2_priority;
G0_priority = ~G3_priority;

```

With this mechanism, each group gets successively more slots at the highest priority. Group0 is the highest priority 1/15 slots, Group1 - 2/15, Group2 - 4/15, and Group3 - 8/15. Here are some of the properties of this rotation schedule.

For adjacent groups:

- Group3 is higher priority than group2 10/15 cycles. (G3 has 100% more bandwidth than G2)
- Group2 is higher priority than group1 10/15 cycles. (G2 has 100% more bandwidth than G1)
- Group1 is higher priority than group0 10/15 cycles. (G1 has 100% more bandwidth than G0)

For groups 2 levels apart:

- Group3 is higher priority than group1 12/15 cycles (G3 has 300% more bandwidth than G1)
- Group2 is higher priority than group0 12/15 cycles (G2 has 300% more bandwidth than G0)

And finally, for groups 3 levels apart:

- Group3 is higher priority than group0 14/15 cycles (G3 has 1300% more bandwidth than G0)



The priorities are rotated potentially every cycle. However, when the highest priority group in a given cycle has multiple runnable TCs in it, then that rotation is held for as many cycles as there are TCs in that highest priority group. This mechanism enables the relative bandwidth between groups to be maintained even when one group contains more TCs than another group.

For instance, assume we have a 4 TC system with 3 TCs in group1 and 1 TC in group0. The exact cycle by cycle priority is described in [Table 10.2](#).

**Table 10.2 Priority Level Rotation (3TCs in group1, 1 TC in group0)**

Cycle Count	Rotation Count	Group3 Priority	Group2 Priority	Group1 Priority	Group0 Priority
1	4'b0001	P3	P2	P1	P0
2		P3	P2	P1	P0
3		P3	P2	P1	P0
4	4'b0010	P2	P3	P0	P1
5	4'b0011	P3	P2	P1	P0
6		P3	P2	P1	P0
7		P3	P2	P1	P0
8	4'b0100	P2	P0	P3	P1
9		P2	P0	P3	P1
10		P2	P0	P3	P1
11	4'b0101	P3	P2	P1	P0
12		P3	P2	P1	P0
13		P3	P2	P1	P0
14	4'b0110	P2	P3	P0	P1
15	4'b0111	P3	P1	P2	P0
16		P3	P1	P2	P0
17		P3	P1	P2	P0
18	4'b1000	P0	P2	P1	P3
19	4'b1001	P3	P1	P2	P0
20		P3	P1	P2	P0
21		P3	P1	P2	P0
22	4'b1010	P2	P3	P0	P1
23	4'b1011	P3	P2	P1	P0
24		P3	P2	P1	P0
25		P3	P2	P1	P0
26	4'b1100	P2	P0	P3	P1
27		P2	P0	P3	P1
28		P2	P0	P3	P1
29	4'b1101	P3	P2	P1	P0
30		P3	P2	P1	P0
31		P3	P2	P1	P0

**Table 10.2 Priority Level Rotation (3TCs in group1, 1 TC in group0) (continued)**

Cycle Count	Rotation Count	Group3 Priority	Group2 Priority	Group1 Priority	Group0 Priority
32	4'b1110	P2	P3	P0	P1
33	4'b1111	P3	P1	P2	P0
34		P3	P1	P2	P0
35		P3	P1	P2	P0

As can be seen in the table, the full rotation actually requires 35 cycles to complete. Out of these 35 cycles, group1 is higher priority than group0 for 30 cycles. However, since group1 contains 3 TCs, these will be round-robin'd by the DS, so on average, each of these TCs will get 33% of this group's bandwidth, or 10cycles. (29% of all the issue slots for each of those TCs in group1). The one TC in group0 gets 5 issue slots, or 14%. As can be seen, each of the TCs in group1 gets about double the issue slots of the TC in group0.

### 10.2.5 CP0 Register Interface

The Policy Manager is controlled using the following CP0 registers.

- TCSchedule register (CP0 Register 2, Select 6)
- TCScheFBack register (CP0 Register 2, Select 7)
- VPESchedule register (CP0 Register 1, Select 5)
- VPEScheFBack register (CP0 Register 1, Select 6)

For more information on these registers, refer to the CP0 Registers chapter in this manual.

# Inter-Thread Communication Unit

This chapter describes the Inter-Thread Communication Unit (ITU) included in the interAptiv Multiprocessing System. This chapter contains the following sections:

- [Section 11.1 “Features Overview”](#)
- [Section 11.2 “ITC Storage”](#)
- [Section 11.3 “ITC Views”](#)
- [Section 11.4 “ITC Address Space”](#)

## 11.1 Features Overview

Inter-Thread Communication (ITC) Storage is a gating storage mechanism designed for low-level thread synchronization. Loads and stores to and from gating storage may block until the state of the storage location corresponds to some set of conditions required for completion. A blocked load or store can be precisely aborted if necessary, and restarted later.

In the interAptiv core, the ITC storage is provided by the Inter-Thread Communication Unit (ITU). This block of logic resides outside of the core and connects to the core through the gating storage interface. SoC integrators are free to use the MIPS-supplied reference module, or to implement their own ITU module, or to not use ITC at all. This chapter describes the features of the sample ITU block supplied with the interAptiv core. This block only supports synchronization of TCs within a single interAptiv core.

## 11.2 ITC Storage

References to memory pages which map to ITC storage resolve not to main memory, but to storage locations, or cells, with special attributes. In general, it is possible that behind each ITC storage cell there is more than one memory location. This is useful for mapping hardware queues, stacks, and other structures. The reference ITU supports two kinds of storage cells: four-entry FIFO queues and single-entry Semaphore cells. All ITC cells are composed of the tag and data portions. In the single-entry cells, the data is 32 bits wide. The FIFO cells store four 32-bit data values. Although the memory space allows for 64-bit ITC cells, only the least-significant 32-bit words are present in this implementation. All ITC cells should be accessed as 32-bit memory. Partial-word access such as LH or SB will result in undefined behavior.

The tag of each ITC cell contains a number of control bits that regulate accesses to that cell. The format for the ITC tag is shown in [Table 11.1](#). In addition to the *E* (Empty) and *F* (Full) fields specified by the MT ASE, the tag contains four implementation-specific fields: *T*, *FIFO*, *FIFODepth*, and *FIFOPtr*. The *FIFO* and *FIFODepth* fields indicate whether a cell is a FIFO and its depth. The *FIFOPtr* indicates how many elements are currently in a FIFO; this field is always zero for single-entry cells. The *FIFOPtr* can be reset by writing 1 into the *E* field of a FIFO. Finally, the *T* field indicates whether a Gating Storage exception should be signaled on an E/F or Proberen/Verhogen (P/V) view access to the cell. In the P/V semaphore, Proberen and Verhogen mean ‘test’ and ‘increment’ respectively.

**Table 11.1 ITC Tag Format**

Name	Bit	Description	Read/Write	Reset State
FIFODepth	31:28	Log <sub>2</sub> of the cell depth. This field is set to 0x0 for single- entry cells, and to 0x2 for four-entry FIFO cells.	R	Preset
FIFOPtr	20:18	This field indicates the number of elements in a FIFO cell, and always reads zero for single-entry Semaphore cells.	R	0
FIFO	17	1 for FIFO cells and 0 for single-entry Semaphore cells.	R	Preset
T	16	Trap Bit. When set, this bit causes the processor to take a Gating Storage Exception on PV or EF accesses.	R/W	Undefined
F	1	Full Bit. This bit indicates that the cell is full.	R/W	Undefined
E	0	Empty Bit. This bit indicates that the cell is empty. Writing 1 to this bit also reset FIFOPtr.	R/W	Undefined
0	27:21, 15:2	Must be written as zeros; return zeros on read.	0	0

The number and type of ITC cells implemented in the ITU is configurable. The possible configurations are: 0, 1, 2, 4, 8, or 16 four-entry FIFOs and 0, 1, 2, 4, 8, or 16 single-entry Semaphores. If the implementation includes both types of cell, the FIFO cells will be grouped before the Semaphore cells. N number of FIFO cells will be located at cell addresses 0 to N-1. M number of Semaphore cells will be located at cell addresses N to N+M-1. The actual physical address is dependent on the base address and cell spacing. See [Section 11.4 “ITC Address Space”](#) for more information on addressing.

## 11.3 ITC Views

All ITC cells can be accessed in one of 16 ways, called views, using standard load and store instructions. The view is encoded in bits 6:3 of the memory address, such that the successive views of a cell correspond to successive 64-bit-aligned addresses. [Table 11.2](#) shows the addresses for the various views, and the following sections describe the effects of using each of the views. If the ITC location is of type FIFO, the behavior of some of the views changes, and this is noted in the description of each view.

**Table 11.2 ITC View Addresses**

Address[6:3]	View
0x0	<a href="#">Bypass View</a>
0x1	<a href="#">Control View</a>
0x2	<a href="#">Empty/Full Synchronized View</a>
0x3	<a href="#">Empty/Full Try View</a>
0x4	<a href="#">P/V Synchronized View</a>
0x5	<a href="#">P/V Try View</a>
0x6-0xF	<a href="#">Reserved Views</a>

### 11.3.1 Bypass View

This view of the ITC location implies that a load or a store does not cause the issuing thread to block and does not affect any of the cells state bits. The operation of SC using this view is undefined.

Accesses using Bypass view never result in Gating Storage exceptions.

A Bypass view store to a FIFO ITC location overwrites the newest FIFO entry, while a Bypass view load returns the contents of the oldest entry.

### 11.3.2 Control View

This view of the ITC location can be used to manipulate the tag of the ITC cell. Loads and stores access the entire 32b tag value. Accesses using Control view never cause the issuing thread to block and never result in Gating Storage exceptions.

A Control view store to a FIFO location with the *E* bit set will cause the FIFO to reset its read pointer.

### 11.3.3 Empty/Full Synchronized View

This view of the ITC location implies that a load causes the issuing thread to block if the cell is Empty. Similarly, a store blocks if the cell is full. Accesses using this view cause an automatic update of the *Empty* and *Full* bits to reflect the new state of the cell. The operation of SC using this view is undefined.

If the *T* bit is set, then all E/F Synchronized view accesses, success or failure, cause a gated exception trap.

### 11.3.4 Empty/Full Try View

This view of the ITC location is similar in nature to the previous E/F Synchronized view in most respects other than the waiting policy on an access failure. It is to be used if the issuing thread can potentially find something else to do and does not wish to be blocked if the access fails. A load with this view returns a value of zero if the cell is Empty, regardless of actual data contained. Otherwise the load behaves as in the E/F Synchronized case. Normal Stores to Full locations through the E/F Try view fail silently to update the contents of the cell, rather than block the thread. SC (Store Conditional) instructions referencing the E/F Try view will indicate success or failure based on whether the ITC store succeeds or fails.

If the *T* bit is set, then all E/F Try view accesses, success or failure, cause a gated exception trap.

### 11.3.5 P/V Synchronized View

This view of the ITC location does not modify the Empty and Full bits, both of which are assumed to be cleared as part of the cell initialization routine. Loads with this view return the current cell data value if the value is non-zero, and cause an atomic post-decrement of the value. If the cell value is zero, loads block until the cell takes a non-zero value. Normal Stores cause an atomic increment of the cell value, up to a maximum of 0xFFFF at which point the value saturates. Loads check the least significant 16bits of the cell for a 0x0 irrespective of load size. The operation of SC using this view is undefined.

If the *T* bit is set, then all P/V Synchronized view accesses, success or failure, cause a gated exception trap.

P/V Synchronized view accesses are not allowed to FIFO ITC locations.

### 11.3.6 P/V Try View

This view of the ITC location is similar in nature to the previous P/V Synchronized view in most respects other than the waiting policy on an access failure. It is to be used if the issuing thread can potentially find something else to do and does not wish to be blocked if the access fails. A load with this view returns a value of zero even if the cell contains a data value of 0x0. Otherwise the load behaves as in the E/F Synchronized case. Normal stores using this view cause a saturating atomic increment of the cell value (saturating to 0xFFFF), as described for the P/V Synchronized view, and cannot fail. The operation of SC using this view is undefined.

If the T bit is set, then all P/V Try view accesses, success or failure, will cause a gated exception trap.

P/V Try view accesses are not allowed to FIFO ITC locations.

### 11.3.7 Reserved Views

These views are reserved and should not be used by software.

## 11.4 ITC Address Space

The ITC physical address space is defined by two, 32-bit registers: *ITCAddressMap0* and *ITCAddressMap1*. Together these two registers specify a  $2^N$  aligned block of uncached memory. The *BaseAddress* field of the *ITCAddressMap0* register specifies the starting address of the ITC memory block. The *AddrMask* of the *ITCAddressMap1* register determines the size of the memory block which can be varied from 1KB to 128KB. Within this address space, ITC cells are spread out with a stride specified by the *EntryGrain* field. Tightly spaced cells save on memory space, but widely spaced cells spread across a number of TLB pages, permitting different cells to be mapped to different processes. The number of cells is specified by the *NumEntries* field.

**Table 11.3 ITC AddressMap0 Register Format**

31	10	9	1	0
BaseAddress			0	En

**Table 11.4 ITCAddressMap1 Register Format**

31	30	20	19	17	16	10	9	3	2	0
M	NumEntries		0	AddrMask		0		EntryGrain		

Fields		Description	Read / Write	Reset State
Name	Bit			
BaseAddress	31:10	The top [31:10] bits of the ITC Physical Memory Mapped Block	R/W	Undefined
En	0	ITC enable	R/W	0
0	9:1	Must be written as zeros; return zeros on read	0	0

Fields		Description	Read / Write	Reset State
Name	Bit			
M	31	This bit indicates if another ITC block is defined along with another pair of ITCAddressMap registers. On the interAptiv core, this value is hardcoded to 0.	R	0
NumEntries	30:20	Number of ITC cells present	R	Preset
AddrMask	16:10	Indicates which bits of the BaseAddress field should not participate in determining an ITC memory hit. This field effectively defines the size of the ITC memory block. AddrMask set to zero implies a 1KB ITC address space, and AddrMask set to 0x3f implies a 128KB address space.	R/W	Undefined
EntryGrain	2:0	Cells are spaced at intervals of $128 \times 2^{\text{EntryGrain}}$ bytes, or: 0x0 - 128B 0x1 - 256B 0x2 - 512B 0x3 - 1KB 0x4 - 2KB 0x5 - 4KB 0x6 - 8KB 0x7 - 16KB	R/W	Undefined
0	19:17, 9:3	Must be written as zeros; return zeros on read	0	0

Depending on the setting of the *AddrMask*, *NumEntries*, and *EntryGrain*, it is possible that ITC cells do not fill up the entire ITC address block. If for example, two cells are mapped to a 1KB area with a stride of 256B (*EntryGrain* equal to 0x1), the first cell starts at offset 0x000 and the second at offset 0x100. The remaining two 256B regions starting at offsets 0x200 and 0x300 do not map to any storage. Any access to an address that does not map to an ITC entry will result in undefined behavior. It is also possible to set the ITC registers in a way that makes some of the cells unavailable.





## Instruction and Data Scratch Pad RAM

The instruction scratchpad RAM (ISPRAM) and data scratchpad RAM (DSPRAM) options on the interAptiv™ Multiprocessing System are designed to provide low-latency access to on-chip memories. Separate SPRAM blocks exist for instruction and data references. The SPRAM ports are accessed in parallel with the caches. This saves a number of cycles that would normally be required when going through the BIU and the master OCP interface of the interAptiv core. Throughout this chapter, the term SPRAM is used to refer to the ISPRAM and DSPRAM memories.

This chapter contains details of the SPRAM interfaces and reference designs. The chapter contains the following major sections:

- [Section 12.1 “Scratchpad RAM \(SPRAM\) Features”](#)
- [Section 12.2 “SPRAM Overview”](#)
- [Section 12.3 “SPRAM Initialization”](#)
- [Section 12.4 “SPRAM Clocking”](#)

### 12.1 Scratchpad RAM (SPRAM) Features

The MIPS32® interAptiv core scratchpad has the following features:

- SPRAM is supported for instruction and data references.
- Each SPRAM block occupies one continuous region in the physical address space. The SPRAM wrappers contain the base physical address and size information.
- SPRAM is virtually indexed by the core. There is no hardware support to avoid virtual aliasing.
- Size of SPRAM may range from 4 KB to 1 MB in factors of 2.
- Data Access granularity
  - Read: 64-bit (1 doubleword).
  - Write: maximum write width is 64 bits, minimum write width is 8 bits.
- Instruction Access granularity
  - Read/Write: 64 bits of instruction plus 6 bits of precode. Smaller writes are not supported.
- SPRAM control supports single or multi-cycle access. For maximum frequency, the SPRAM access time should be less than the cache access time. For larger size SPRAMs, the integrator may choose a multi-cycle access
  - For data references, the multi-cycle accesses can be pipelined. SPRAM data needs to be returned in the requested order.
  - For instruction references, requests will be retried if the data is not available at the single cycle point.
- Multi-cycle data scratchpad RAM access is non-blocking.

- Multi-cycle instruction scratchpad RAM access is blocking (within a TC).
- The scratchpad RAM is not required to hold the last read value.
- The data scratchpad has independent tag and data ports. The tag and data arrays are always read together for the instruction scratchpad. The scratchpad RAM does not have a traditional cache tag array. Instead, it has registers holding the SPRAM configuration information.
- SPRAM access hit supersedes cache access hit.
- User may implement a DMA port to the scratchpad RAM. In the reference designs, an OCP slave port is provided.
- The interAptiv core provides integrated BIST support for single-cycle latency SPRAM.
- Optional parity protection is supported for SPRAM.
  - Instruction: 1b per 8b of instruction, 1b for 7b precode
  - Data Parity: 1b per 8b of data

## 12.2 SPRAM Overview

A Scratchpad RAM can be used stand-alone or combined with data or instruction cache. The existence of a scratchpad must be selected at build time.

The SPRAM array, like the cache arrays, is indexed with a virtual address and the “tag comparison” (really just decode logic for the SPRAM) is performed using a physical address. Since the SPRAM size can be larger than the 4 KB minimum page size, it is possible to have virtual aliasing in the SPRAM. Virtual aliasing occurs when a single physical address is accessed via two different virtual addresses that can simultaneously reside in memory. This is not a problem on cores using the Fixed Mapping Translation MMU. For cores with TLB-based MMUs, this can be avoided by accessing the SPRAM through unmapped (kseg0/1) addresses or using using a TLB page. This is not handled by hardware and programmers must be aware of it.

The reference designs contain 8 KB SPRAM arrays, with one cycle latency and a simple DMA port. A user can choose to implement a custom SPRAM with different size, latency, and other desired characteristics.

During normal operation, it will be impossible for a reference to hit in both SPRAM and data cache. If this error condition does occur via manipulation of the cache or SPRAM tags, the SPRAM supersedes the data cache hit. Note that this also means that a CACHE HitInvalidate operation to such a line that exists in both SPRAM and cache will not invalidate the cache entry.

The scratchpad interface consists of a core-side interface as well as an optional DMA interface. MIPS provides a reference design for the external SPRAM module called `imp_sp` and `imp_isp`. These include wrappers that instantiates a SPRAM SRAM array. The reference module can be replaced with a customized SPRAM implementation. For timing reasons, the arbitration logic for the SPRAM DMA interface is located within the `imp_cpu` hierarchy.

### 12.2.1 SPRAM Differences Versus a Cache

SPRAM behavior differs from cache in the following key ways:

- Software must ensure a SPRAM entry has been initialized before it is read, to avoid reading spurious data.

- The SPRAM does not refill automatically. The data SPRAM is normally initialized with stores or DMA writes to the desired address range. The instruction SPRAM can be initialized with Index Store Data CACHE ops or DMA writes.
- Store operations which hit in the data SPRAM do not produce writes to main memory, unlike write-through stores that hit in the cache and write to main memory.

NOTE: The I-Cache Fill and Fetch&Lock cacheops will refill the given line into the I-Cache even if that address hits to the ISPRAM. This is not recommended since normal fetches will hit in the ISPRAM and ignore the I-Cache contents.

## 12.2.2 Uncacheable References to SPRAM

SPRAM can be mapped to either cached or uncached space. The address decode and comparison for SPRAM is done regardless of the cacheability attribute.

## 12.2.3 Independent Tag/Data Accesses

The data SPRAM interface has independent tag and data ports. This is done to aid the efficiency of stores. A store must perform a lookup to determine if/where to write the data, then the actual data must be written. Because the lookup does not need to access the data array, these operations can occur in parallel if the data writes are buffered within the core.

Because there are no stores to the instruction SPRAM, the tag and data ports are linked. Reads will always access both the tag and data port at the same index. Writes will target either the tag or data.

## 12.2.4 SPRAM Tag Reads and Writes

The interface allows for SPRAM “tag” values to be read and written. The tag values are read/written by the CACHE instruction. This can optionally provide a mechanism for software to determine the SPRAM base and size configuration and change it. The reference design shows one possible use for this interface - software can probe the SPRAM to determine the base address and whether it is enabled. These values are also write-able, allowing software to dynamically configure the SPRAM parameters.

## 12.2.5 Multiple Cycle Data SPRAM Access

For a one-cycle latency SPRAM, the scratchpad interface will achieve cache-like access timing. However, the scratchpad interface also supports SPRAM that has a multi-cycle latency.

For a data scratchpad read, when the data from SPRAM is not ready, the processor will register the load in a load buffer and return data to the main pipeline when data is available from SPRAM. No stall of pipeline is necessary unless the result register is used by a following instruction.

## 12.2.6 Multiple Cycle Instruction SPRAM Access

For an instruction scratchpad read, multi-cycle latency causes more of a problem. The instruction fetch pipeline does not have the ability to stall. If the instruction data is not returned at the expected time, the request will be retried. This will add a minimum of 3 cycles of latency to all ISPRAM fetches.

When the core is operating in multi-threaded mode, this gets even trickier. Fetch requests from the different TCs can be intermixed. It is recommended that a multi-cycle ISPRAM have a buffer per TC that holds the last requested fetch to allow TCs to make forward progress.

## 12.2.7 Backstalling the SPRAM Data Port

The backstall mechanism is not really needed if the SPRAM can keep track of all the outstanding requests. If that is not the case, SPRAM is allowed to backstall the core if it is busy, via assertion of the *SP\_ram\_busy* or *ISP\_ram\_busy* signals.

This backstall mechanism may be useful if the customer implements a multi-cycle non-pipelined version of data SPRAM. The scratchpad should assert the *SP\_ram\_busy* when it cannot accept another request in the next cycle. The core will stall its pipeline only if it has a pending SPRAM access and it is about to enter the ER stage of the pipeline.

The above mentioned mechanism only applies to the SPRAM data port, not the tag. The tag port always requires fixed single-cycle latency

## 12.2.8 Access Granularity

A data SPRAM read returns either 1 word (32b) or 1 doubleword (64b) of data (plus parity/ECC). There are two read strobes controlling access to the upper and lower word of the data array. Since many core accesses are 32b or less, banking the SPRAM array and only reading the selected bank can yield power savings. Alternatively, the OR of the read strobes can be used to access a single wider array.

For word accesses, the processor core uses the lower 32b of the read data bus. If only the upper read strobe is asserted, the upper 32b word be returned on the lower 32b of the data bus. Additional alignment and shifting is handled within the core. The maximum write width is 64 bits and partial write is enabled through the byte enable signals, *SP\_data\_wren\_ag[7:0]*.

For DMA access to data SPRAM, the read will always be 64 bits wide. The maximum write width is 64 bits and partial write is enabled using the *OC\_DMA\_MDataByteEn[7:0]* signals.

On the instruction SPRAM, both reads and writes will always be 70 bits wide (64b of data + 6b of precode) (plus parity). The additional precode data will be generated by the core during a DMA write. Data on the OCP Slave bus will only be 64b wide.

## 12.2.9 Connecting I/O Devices to the Data Scratchpad Interface

In addition to, or perhaps instead of, an SRAM array, it is possible to connect I/O devices to the SPRAM interface. Connecting I/O devices to the scratchpad interface allows low latency, high throughput access to critical I/O devices in the system. To accomplish this, the integrator must ensure that the behavior of the I/O devices meets the same requirements as the SPRAM.

Connecting an I/O device to the ISPRAM is not recommended.

## 12.2.10 Null Connection to Unused SPRAM Interface

The presence of scratchpads must be chosen when the core is built. Even if one or both of the SPRAM interfaces are present, there does not need to be arrays connected to them. If a interface is not going to be used, then the *[I]SP\_Present* input signal to the core should be driven low. All other input signals to the core for the unused SPRAM interface should also be tied low, to avoid floating inputs. All output signals from the core related to the unused SPRAM interface can be left unconnected.

## 12.3 SPRAM Initialization

Since the scratchpad is really a RAM-based structure, it must be initialized with valid data before it can be used. Following are few ways to initialize the data SPRAM.

- DMA: The RAM array can be initialized from the system using DMA writes.
- Stores: For data SPRAM, the array can be initialized with normal store instructions that hit in the SPRAM region.
- CACHE Index Store Data instruction: Indexed cache operations can be forced to go to the SPRAM by setting the SPR bit in the Coprocessor0 *ErrCtl* register. When this bit is set, it is possible to use the Index Store Data flavor of the CACHE instruction to move data from the *DataLo/DataHi* Cop0 registers into the SPRAM. This mechanism does not require any backing memory and can even be used to load the SPRAM from an EJTAG probe for early system bringup. For the data SPRAM, using stores to initialize the array is usually a much more efficient mechanism.

### 12.3.1 ISPRAM Boot

Uncached requests can still be serviced by the SPRAMs. On MIPS CPUs, the bootvector is located at an uncacheable address. Since the SPRAM has cache-like timing even when responding to uncached accesses, it can run much faster. This can make booting directly from the ISPRAM an interesting possibility.

Note: When fetching from uncached addresses, even if they hit in the SPRAM, the core will only use 32b at a time instead of 64b. This will reduce the performance versus SPRAM hits in cached space, but will still be much better than normal uncached accesses.

In order to boot from the ISPRAM, the instructions must be loaded into the array before the core can start executing them. The reference design does include some support for this via the *ispram\_boot* internal signal. By default, this signal is statically driven to 0, but commented out RTL shows how to connect it to one of the sideband external signals to allow it to be dynamically controlled. There is example RTL for either directly connecting it or synchronizing the input signal - the latter is recommended.

In order to load the ISPRAM via DMA and boot directly from it:

- Set *ispram\_boot* = 1 while *SI\_Reset* = 1
  - This sets the base address to the physical address of the boot vector (either 0x1fc0\_0000 or *SI\_ExceptionBase* if *SI\_UseExceptionBase* = 1) and sets the enable bit.
- While in reset, the DMA port will be inactive (core deasserts Accept signals)
- After *SI\_Reset* > 0, hold *ispram\_boot* = 1 until the ISPRAM has been loaded via DMA.
  - Note: the DMA port is held inactive while the core is in reset, thus the DMA can only happen after reset has been deasserted
  - This causes *ISP\_dma\_stallreq\_xx* = 1 which gives the DMA priority over core requests .
  - And also sets *ISP\_datavld\_nxt\_if* = 0, which indicates that the data is not available yet and the core will retry any accesses that hit in the ISPRAM.
- Once the ISPRAM has been loaded, *ispram\_boot* should be deasserted, allowing the core accesses to hit out of the ISPRAM.

## 12.3.2 ISPRAM Precode Bits

Six precode bits are included with every 64b of instruction data in the ISPRAM array. These bits contain information about branches and jumps. Having this information allows the fetch unit to quickly react to the potential change of flow and start fetching along the predicted path. These precode bits are not used when executing MIPS16e™ code.

When the ISPRAM array is loaded using Index Store Data CACHE instructions or by using the core's DMA interface, the precode bits are generated automatically and sent out as part of the write data. So, for many systems, nothing special will need to be done with the precode bits. In some custom ISPRAM blocks, however, it may not be possible to utilize the precode blocks within the core. Two examples are if the ISPRAM block contains a ROM array or if the core DMA interface is not used.

**Table 12.1 Precode Bits**

Name	Bit	Description
L/M	6	This bit has two different meanings depending on the state of the B and J bits. 1. If B is set, then this bit is set if the branch is a branch likely instruction. 2. If J is set, then this bit is set if the Jump instruction is a JALX instruction.  If either of the B or J bits are set, then the IFU fetches a delay slot instruction. When both B and J bits are set, this indicates that both instructions decoded to look like branches, Jumps, or ERET instructions. This can happen if one instruction isn't really an instruction, but is instead data.
X	5	0: branch/jump is in bits[31:0] 1: branch/jump is in bits [63:32]
B	4	Branch instruction
J	3	Jump instruction
S	2	Indicates a subroutine call. Return address will be pushed onto return prediction stack
G	1	Indicates Jump Register is not predicted
U/R	0	On branches, indicates an unconditional branch On jumps, indicates a return

**Table 12.2 MIPS32 Control Transfer Instructions**

Instruction	Precoding	Notes
B (BEQ rs==rt)	BU	Does not use branch predictor.
BAL (BGEZAL r0)	BSU	Does not use branch predictor. push PC+8 onto RPS
BC1 [TF]	B	
BC1 [TF]L	B	
BC2 [TF]	B	
BC2 [TF]L	B	
BEQ (rs != rt)	B	
BEQL	B	

**Table 12.2 MIPS32 Control Transfer Instructions (*continued*)**

Instruction	Precoding	Notes
B[GL][ET]Z	B	
BGEZAL (rs != r0)	B	
BLTZAL	B	
B[GE,LT]ZALL	B	
B[GL][ET]ZL	B	
BNE	B	
BNEL	B	
BPOSGE32	B	Instruction from MIPS DSP ASE
DERET	G	decode in IS
ERET	G	decode in IS
J	J	
JAL	JS	push PC+8 onto RPS
JALR[.HB] (rd = \$31)	JSG	push PC+8 onto RPS possible MIPS16e mode change
JALR[.HB] (rd != \$31)	JG	possible MIPS16e mode change
JALX	JS	push PC+8 onto RPS switch to MIPS16e mode
JR (rs = \$31)	JR	possible speculative MIPS16e mode change
JR (rs != \$31)	JGSR	possible MIPS16e mode change
JR.HB (rs = \$31)	JGR	possible MIPS16e mode change pop RPS but don't use
JR.HB (rs != \$31)	JG	possible MIPS16e mode change
ILLEGAL	BJG	If both instructions decode as a branch/jump, set this to let fetch unit know there is a strange situation that needs resolving.  This can happen when data is packed with instructions, when precoding MIPS16e instruction data, or when there is an illegal code sequence.

## 12.4 SPRAM Clocking

The SPRAM block receives two clocks from the core, *[I]SP\_gclk* and *[I]SP\_gfclk*. *[I]SP\_gclk* is a global gated clock, while *[I]SP\_gfclk* is free-running. When top level clock gating is implemented, *[I]SP\_gclk* is not active in the low power sleep mode entered using the WAIT instruction. The processor enters sleep mode only if there are no pending SPRAM transactions. When the processor core is in sleep mode, a DMA request will re-enable *[I]SP\_gclk* in order to process the request.

For the best power management, most of logic in the SPRAM block should reside on the gated *[I]SP\_gclk*. Only the minimal logic needed to detect a DMA request and wake up the core needs to reside on the free-running *[I]SP\_gfclk*.

## 12.4.1 Scratchpad Reference Design

The reference scratchpad design (called `imp_[i]sp`) supports a basic scratchpad implementation. It is configurable within certain constraints:

- SPRAM size can range from 4 KByte to 1 MByte. The supported sizes are 4KB, 8KB, 16KB, 32KB, 64KB, 128 KB, 256 KB, 512 KB, or 1 MB
- SPRAM tag has a base address and a size register used for hit detection. Both base and size register are accessible through the CACHE instruction.
- The address range must be naturally aligned (i.e. a 64KB SPRAM's base address must be on a 64KB boundary).
- The array always returns data in a single cycle.
- The scratchpad contains an OCP DMA slave port.

Following are some considerations of the reference SPRAM design which are not covered in the previous sections.

## 12.4.2 Tag Registers in the Reference SPRAM

Using the CACHE instruction, it is possible to read or write the “tag” value associated with the SPRAM. To provide a common software interface, it is recommended that all SPRAM implementations provide some standard configuration information via this mechanism.

In the reference SPRAM wrapper, the “tag” of SPRAM consists of a base address register and a size register. If the SPR bit in the *ErrCtl* register is set, an Index Load Tag CACHE instruction reads the SPRAM tag and place the contents in the *TagLo* register, while an Index Store Tag CACHE instruction writes the SPRAM tag with the data from *TagLo* register. Bit3 of the index is used to select between base address and size register; when bit3 =1, the size register is selected, otherwise the base address register is selected. The format of the base and size registers are shown in [Table 12.3](#) and [Table 12.4](#), respectively.

**Table 12.3 Format of the Base Address Register in the Reference SPRAM Wrapper**

Field	Description
<code>sp_base_xx[31:12]</code>	Base address of the SPRAM region
<code>sp_base_xx[11]</code>	SPRAM valid



**Table 12.4 Format of the Size Register in the Reference SPRAM Wrapper**

Field	Description																				
sp_size_xx[31:12]	Size of the SPRAM: <table border="1" data-bbox="771 336 1201 709"> <thead> <tr> <th>SPRAM size</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>4KB</td> <td>20'h00001</td> </tr> <tr> <td>8KB</td> <td>20'h00002</td> </tr> <tr> <td>16KB</td> <td>20'h00004</td> </tr> <tr> <td>32KB</td> <td>20'h00008</td> </tr> <tr> <td>64KB</td> <td>20'h00010</td> </tr> <tr> <td>128KB</td> <td>20'h00020</td> </tr> <tr> <td>256KB</td> <td>20'h00040</td> </tr> <tr> <td>512KB</td> <td>20'h00080</td> </tr> <tr> <td>1M</td> <td>20'h00100</td> </tr> </tbody> </table>	SPRAM size	Value	4KB	20'h00001	8KB	20'h00002	16KB	20'h00004	32KB	20'h00008	64KB	20'h00010	128KB	20'h00020	256KB	20'h00040	512KB	20'h00080	1M	20'h00100
SPRAM size	Value																				
4KB	20'h00001																				
8KB	20'h00002																				
16KB	20'h00004																				
32KB	20'h00008																				
64KB	20'h00010																				
128KB	20'h00020																				
256KB	20'h00040																				
512KB	20'h00080																				
1M	20'h00100																				

Software can then query or modify these registers to determine the base and size information. An Index Load Tag will read bits [31:12] from the base address or size register and write them into bits [31:12] of the *TagLo* register. Bit [11] of the base address register serves as a valid bit for SPRAM and will be written into the valid field (bit [7]) of the *TagLo* register. Similarly, an Index Store Tag will write bits [31:12] of the *TagLo* register into bits [31:12] of the base address or size register, and the *TagLo* valid field into the valid bit of base address register.

The reset value of base address information is incorporated into the SPRAM via `define` within the reference SPRAM module. The reset value of the size tag is based on the actual size of the array.

### 12.4.3 Enabling SPRAM Access

After power up, the reference SPRAM is always set to invalid through the reset of bit [11] of the base address register. So a CACHE Index Store Tag instruction is needed to enable the SPRAM. The core comparison logic uses bit [11] of `[I]SP_tag_rdata_xx[31:11]` as the valid bit of SPRAM and an access can hit on SPRAM only when this bit is set.

When a custom SPRAM is implemented, this bit should be set accordingly for the design.

### 12.4.4 SPRAM BIST Support

The core includes an integrated SPRAM BIST controller which can provide BIST support for single-cycle latency SPRAM. The integrated SPRAM BIST controller is capable of supporting two algorithms, March C+ or IFA-13 (IFA-13 includes support for retention testing).

When integrated memory BIST is running, the SPRAM array is tested in parallel with other sub arrays of the instruction and data caches and trace memory.

A custom RAM BIST module is also possible. For a multi-cycle SPRAM implementation, custom BIST is required since the integrated controller only accommodates single-cycle access.

### 12.4.5 SPRAM Parity Support

Parity protection is optionally enabled for SPRAM. A parity error on a SPRAM read will either cause a CacheErr exception (for a load or fetch) or an error response on the OCP bus (DMA access). The CacheErr parity error detec-

tion logic resides in the core. For the reference design, if parity is enabled, it must be supported by the instruction cache, data cache and both SPRAM arrays.

From the reference SPRAM module, the outputs *SP\_parity\_present*, *SP\_ecc\_present*, and *ISP\_parity\_present* indicate whether each SPRAM array is parity protected. If a custom SPRAM module is built, users might choose not to check parity for SPRAM even though parity checking for instruction and data caches is enabled; in this case, the output should be de-asserted and no parity checking will be done for that SPRAM.

## Hardware and Software Initialization

A interAptiv core contains only a minimal amount of hardware initialization and relies on software to fully initialize the device.

This chapter contains the following sections:

- [Section 13.1 “Hardware-Initialized Processor State”](#)
- [Section 13.2 “Software-Initialized Processor State”](#)

### 13.1 Hardware-Initialized Processor State

The interAptiv core is not fully initialized by hardware reset. Only a minimal subset of the processor state is cleared. This is enough to bring the core up while running in unmapped and uncached code space. All other processor state can then be initialized by software. Unlike previous MIPS processors, there is no distinction between cold and warm resets (or hard and soft resets). *SI\_Reset* is used for both power-up reset and soft reset.

#### 13.1.1 Coprocessor 0 State

Much of the hardware initialization occurs in Coprocessor 0:

- *Random* - cleared to maximum value on Reset
- *Wired* - cleared to 0 on Reset
- *Status<sub>BEV</sub>* - set to 1 on Reset
- *Status<sub>TS</sub>* - cleared to 0 on Reset
- *Status<sub>NMI</sub>* - cleared to 0 on Reset
- *Status<sub>ERL</sub>* - set to 1 on Reset
- *Status<sub>RP</sub>* - cleared to 0 on Reset
- *CDMMBase<sub>EN</sub>* - cleared to 0 on Reset
- *WatchLo<sub>I,R,W</sub>* - cleared to 0 on Reset
- *Config* fields related to static inputs - set to input value by Reset
- *Config<sub>K0</sub>* - set to 010 (uncached) on Reset

- *Config<sub>KU</sub>* - set to 010 (uncached) on Reset
- *Config<sub>K23</sub>* - set to 010 (uncached) on Reset
- *Debug<sub>DM</sub>* - cleared to 0 on Reset (unless EJTAGBOOT option is used to boot into Debug Mode, as described in [Chapter 16, “EJTAG Debug Support”](#)).
- *Debug<sub>LSNM</sub>* - cleared to 0 on Reset
- *Debug<sub>IBusEP</sub>* - cleared to 0 on Reset
- *Debug<sub>DBusEP</sub>* - cleared to 0 on Reset
- *Debug<sub>IEXI</sub>* - cleared to 0 on Reset
- *Debug<sub>SSI</sub>* - cleared to 0 on Reset

### 13.1.2 TLB Initialization

Each TLB entry has a “hidden” state bit, which is set by Reset and is cleared when the TLB entry is written. This bit disables matches and prevents “TLB Shutdown” conditions from being generated by the power-up values in the TLB array (when two or more TLB entries match a single address). This bit is not visible to software.

### 13.1.3 Bus State Machines

All pending bus transactions are aborted and the state machines in the bus interface unit are reset when a Reset exception is taken.

### 13.1.4 Static Configuration Inputs

All static configuration inputs (for example, those defining the bus mode and cache size) should only be changed during Reset.

### 13.1.5 Fetch Address

Upon Reset, unless the EJTAGBOOT option is used, the fetch is directed to VA 0xBFC00000 (PA 0x1FC00000). This address is in kseg1, which is unmapped and uncached, so that the TLB and caches do not require hardware initialization.

## 13.2 Software-Initialized Processor State

Software is required to initialize parts of the device, as described below.

### 13.2.1 Register File

The register file powers up in an unknown state with the exception of r0, which is always 0. Initializing the rest of the register file is not required for proper operation. Good code will generally not read a register before writing to it, but the boot code can initialize the register file for added safety.

## 13.2.2 TLB

Because of the hidden bit indicating initialization, the core does not initialize the TLB upon Reset. This is an implementation-specific feature of the interAptiv core and cannot be relied upon if writing generic code for MIPS32/64 processors.

## 13.2.3 Caches

The cache tag and data arrays power up to an unknown state and are not affected by reset. Every tag in the cache arrays should be initialized to an invalid state using the CACHE instruction (typically the Index Invalidate function). This can be a long process, especially because the instruction cache initialization must run in an uncached address region.

## 13.2.4 Coprocessor 0 State

Miscellaneous COP0 states need to be initialized before exiting the boot code. There are various exceptions which are blocked by *ERL*=1 or *EXL*=1, and which are not cleared by Reset. These can be cleared to avoid taking spurious exceptions when leaving the boot code.

- *Cause*: *WP* (Watch Pending), and *SW0* and *SW1* (Software Interrupts) should be cleared.
- *Config*: *K0* should be set to the desired Cache Coherency Algorithm (CCA) prior to accessing *kseg0*.
- *Count*: Should be set to a known value if timer interrupts are used.
- *Compare*: Should be set to a known value if timer interrupts are used. Note that the write to *Compare* will also clear any pending timer interrupts, so *Count* should be set before *Compare* to avoid any unexpected interrupts.
- *Status*: Desired state of the device should be set.
- Other COP0 state: Other registers should be written before they are read. Some registers are not explicitly writable, and are only updated as a by-product of instruction execution or a taken exception. Uninitialized bits should be masked off after reading these registers.

## 13.3 Boot and CMP Bringup

After the system is reset and released, all cores configured in hardware to power up will execute their boot sequence. Typically, CPU0 powers up, while all other CPUs are configured to remain powered down. Alternatively, all CPUs can be hardware configured to remain powered down to be awakened through a hardware signal connected to SOC-specific logic.

After system reset, all caches are in an unknown state and must be initialized. It is advisable for core0 to initialize the L2 cache prior to powering up the other cores, but this is not required if other synchronization methods are utilized. For L1 caches, this is expected to be done using IndexStTag ops running on the same CPU. Prior to the data cache being initialized, processing an intervention would cause unpredictable results, potentially corrupting main memory with random data. Thus, the system starts with all of the cores outside the coherence domain until explicitly enabled by software.

```
Core0:  
Initialize cop0 state  
Initialize L2 Cache  
Initialize GCR state  
Startup other cores if needed  
CoreN:  
Initialize L1 Caches  
Enable Coherence  
Switch to coherent CCA
```

# Floating-Point Unit

This chapter describes the optional MIPS32® Floating-Point Unit (FPU) and contains the following sections:

- [Section 14.1, "Features Overview"](#)
- [Section 14.2 "IEEE Standard 754"](#)
- [Section 14.3 "Enabling the Floating-Point Coprocessor"](#)
- [Section 14.4 "Data Formats"](#)
- [Section 14.5 "Floating-Point General Registers"](#)
- [Section 14.6 "Floating-Point Control Registers"](#)
- [Section 14.7 "Exceptions"](#)
- [Section 14.8 "Latency and Repeat Rates"](#)
- [Section 14.9 "Instruction Overview"](#)
- [Section 14.10 "Alphabetical Listing of Floating Point Instructions"](#)

## 14.1 Features Overview

The FPU is provided via Coprocessor 1 (CP1). Together with its dedicated system software, the FPU fully complies with the ANSI/IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*. The MIPS architecture supports the recommendations of IEEE Standard 754, and the coprocessor implements a precise exception model. The key features of the FPU are listed below.

- Full 64-bit operation is implemented in both the register file and functional units.
- A 32-bit Floating-Point Control register controls the operation of the FPU, and monitors condition codes and exception conditions.
- Like the main processor core, Coprocessor 1 is programmed and operated using a Load/Store instruction set. The processor core communicates with Coprocessor 1 using a dedicated coprocessor interface. The FPU functions as an autonomous unit. The hardware is completely interlocked such that, when writing software, the programmer does not have to worry about inserting delay slots after loads and between dependent instructions.
- Additional arithmetic operations not specified by IEEE Standard 754 (for example, reciprocal and reciprocal square root) are specified by the MIPS architecture and are implemented by the FPU. In order to achieve low latency counts, these instructions satisfy more relaxed precision requirements.
- The MIPS architecture further specifies compound multiply-add instructions. These instructions meet the IEEE accuracy specification, where the result is numerically identical to an equivalent computation using multiply, add, subtract, or subtract from zero instructions.

## 14.2 IEEE Standard 754

The IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*, is referred to in this chapter as “IEEE Standard 754”. IEEE Standard 754 defines the following:

- Floating-point data types
- The basic arithmetic, comparison, and conversion operations
- A computational model

IEEE Standard 754 does not define specific processing resources nor does it define an instruction set.

## 14.3 Enabling the Floating-Point Coprocessor

Coprocessor 1 is enabled by setting the CU1 bit in the CP0 *Status* register. When this bit is cleared, Coprocessor 1 is disabled, and any attempt to execute a floating-point instruction causes a *Coprocessor Unusable* exception.

## 14.4 Data Formats

The FPU provides both floating-point and fixed-point data types, which are described below:

- The single- and double-precision floating-point data types are those specified by IEEE Standard 754.
- The fixed-point types are signed integers provided by the CPU architecture.

### 14.4.1 Floating-Point Formats

The FPU provides the following two floating-point formats:

- A 32-bit single-precision floating point (type S)
- A 64-bit double-precision floating point (type D)

The floating-point data types represent numeric values as well as the following special entities:

- Two infinities,  $+\infty$  and  $-\infty$
- Signaling non-numbers (SNaNs)
- Quiet non-numbers (QNaNs)
- Numbers of the form:  $(-1)^s 2^E b_0.b_1 b_2..b_{p-1}$ , where:
  - $s = 0$  or  $1$
  - $E =$  any integer between  $E_{\min}$  and  $E_{\max}$ , inclusive
  - $b_i = 0$  or  $1$  (the high bit,  $b_0$ , is to the left of the binary point)
  - $p$  is the signed-magnitude precision



The single and double floating-point data types are composed of three fields—sign, exponent, fraction—whose sizes are listed in [Table 14.1](#).


**Table 14.1 Parameters of Floating-Point Data Types**

Parameter	Single	Double
Bits of mantissa precision, $p$	24	53
Maximum exponent, $E_{max}$	+127	+1023
Minimum exponent, $E_{min}$	-126	-1022
Exponent <i>bias</i>	+127	+1023
Bits in exponent field, $e$	8	11
Representation of $b_0$ integer bit	hidden	hidden
Bits in fraction field, $f$	23	52
Total format width in bits	32	64
Magnitude of largest representable number	3.4028234664e+38	1.7976931349e+308
Magnitude of smallest normalized representable number	1.1754943508e-38	2.2250738585e-308


Layouts of these three fields are shown in [Figures 14.1](#) and [14.2](#) below. The fields are:

- 1-bit sign,  $s$
- Biased exponent,  $e = E + bias$
- Binary fraction,  $f = .b_1 b_2 \dots b_{p-1}$  (the  $b_0$  bit is *hidden*; it is not recorded)

**Figure 14.1 Single-Precision Floating-Point Format (S)**



**Figure 14.2 Double-Precision Floating-Point Format (D)**



Values are encoded in the specified format using the unbiased exponent, fraction, and sign values listed in [Table 14.2](#). The high-order bit of the Fraction field, identified as  $b_1$ , is also important for NaNs.

**Table 14.2 Value of Single or Double Floating-Point Data Type Encoding**

Unbiased $E$	$f$	$s$	$b_1$	Value $V$	Type of Value	Typical Single Bit Pattern <sup>1</sup>	Typical Double Bit Pattern <sup>1</sup>
$E_{max} + 1$	$\neq 0$		1	SNaN	Signaling NaN	0x7fffffff	0x7fffffff ffffffff
			0	QNaN	Quiet NaN	0x7fbffff	0x7ff7ffff ffffffff

**Table 14.2 Value of Single or Double Floating-Point Data Type Encoding (continued)**

Unbiased E	f	s	b <sub>1</sub>	Value V	Type of Value	Typical Single Bit Pattern <sup>1</sup>	Typical Double Bit Pattern <sup>1</sup>
$E_{max} + 1$	0	1		$-\infty$	Minus infinity	0xff800000	0xffff00000 00000000
		0		$+\infty$	Plus infinity	0x7f800000	0x7ff00000 00000000
$E_{max}$ to $E_{min}$		1		$-(2^E)(1.f)$	Negative normalized number	0x80800000 through 0xff7fffff	0x80100000 00000000 through 0xffeffffff ffffffff
		0		$+(2^E)(1.f)$	Positive normalized number	0x00800000 through 0x7f7fffff	0x00100000 00000000 through 0x7fefffff ffffffff
$E_{min} - 1$	$\neq 0$	1		$-(2^{E_{min}})(0.f)$	Negative denormalized number	0x807fffff	0x800fffff ffffffff
		0		$+(2^{E_{min}})(0.f)$	Positive denormalized number	0x007fffff	0x000fffff ffffffff
$E_{min} - 1$	0	1		- 0	Negative zero	0x80000000	0x80000000 00000000
		0		+ 0	positive zero	0x00000000	0x00000000 00000000

1. The “Typical” nature of the bit patterns for the NaN and denormalized values reflects the fact that the sign might have either value (NaN) and that the fraction field might have any non-zero value (both). As such, the bit patterns shown are one value in a class of potential values that represent these special values.

**14.4.1.1 Normalized and Denormalized Numbers**

For single and double data types, each representable nonzero numerical value has just one encoding; numbers are kept in normalized form. The high-order bit of the p-bit mantissa, which lies to the left of the binary point, is “hidden,” and not recorded in the *Fraction* field. The encoding rules permit the value of this bit to be determined by looking at the value of the exponent. When the unbiased exponent is in the range  $E_{min}$  to  $E_{max}$ , inclusive, the number is normalized and the hidden bit must be 1. If the numeric value cannot be normalized because the exponent would be less than  $E_{min}$ , then the representation is denormalized, the encoded number has an exponent of  $E_{min} - 1$ , and the hidden bit has the value 0. Plus and minus zero are special cases that are not regarded as denormalized values.

**14.4.1.2 Reserved Operand Values—Infinity and NaN**

A floating-point operation can signal IEEE exception conditions, such as those caused by uninitialized variables, violations of mathematical rules, or results that cannot be represented. If a program does not trap IEEE exception conditions, a computation that encounters any of these conditions proceeds without trapping but generates a result indicating that an exceptional condition arose during the computation. To permit this case, each floating-point format defines representations (listed in the table above) for plus infinity ( $+\infty$ ), minus infinity ( $-\infty$ ), quiet non-numbers (QNaN), and signaling non-numbers (SNaN).

**14.4.1.3 Infinity and Beyond**

Infinity represents a number with magnitude too large to be represented in the given format; it represents a magnitude overflow during a computation. A correctly signed  $\infty$  is generated as the default result in division by zero operations and some cases of overflow as described in [Section 14.7.2 “Exception Conditions”](#).

Once created as a default result,  $\infty$  can become an operand in a subsequent operation. The infinities are interpreted such that  $-\infty < (\text{every finite number}) < +\infty$ . Arithmetic with  $\infty$  is the limiting case of real arithmetic with operands of arbitrarily large magnitude, when such limits exist. In these cases, arithmetic on  $\infty$  is regarded as exact, and exception conditions do not arise. The out-of-range indication represented by  $\infty$  is propagated through subsequent computa-

tions. For some cases, there is no meaningful limiting case in real arithmetic for operands of  $\infty$ . These cases raise the Invalid Operation exception condition as described in [Section 14.7.2.1 “Invalid Operation Exception”](#).

#### 14.4.1.4 Signalling Non-Number (SNaN)

SNaN operands cause an Invalid Operation exception for arithmetic operations. SNaNs are useful values to put in uninitialized variables. An SNaN is never produced as a result value.

IEEE Standard 754 states that “Whether copying a signaling NaN without a change of format signals the Invalid Operation exception is the implementor’s option.” The MIPS architecture makes the formatted operand move instructions (**MOV.fmt**, **MOVT.fmt**, **MOVF.fmt**, **MOVN.fmt**, **MOVZ.fmt**, **ABS.fmt**, **NEG.fmt**) non-arithmetic; they do not signal IEEE 754 exceptions.

#### 14.4.1.5 Quiet Non-Number (QNaN)

QNaNs provide retrospective diagnostic information inherited from invalid or unavailable data and results. Propagation of the diagnostic information requires information contained in a QNaN to be preserved through arithmetic operations and floating-point format conversions.

QNaN operands do not cause arithmetic operations to signal an exception. When a floating-point result is to be delivered, a QNaN operand causes an arithmetic operation to supply a QNaN result. When possible, this QNaN result is one<sup>1</sup> of the operand QNaN values. QNaNs do have effects similar to SNaNs on operations that do not deliver a floating-point result—specifically, comparisons. (For more information, see the detailed description of the floating-point compare instruction, **C.cond.fmt**.)

When certain invalid operations not involving QNaN operands are performed but do not trap (because the trap is not enabled), a new QNaN value is created. [Table 14.3](#) shows the QNaN value generated when no input operand QNaN value can be copied. The values listed for the fixed-point formats are the values supplied to satisfy IEEE Standard 754 when a QNaN or infinite floating-point value is converted to fixed point. There is no other feature of the architecture that detects or makes use of these “integer QNaN” values.

**Table 14.3 Value Supplied When a New Quiet NaN is Created**

Format	New QNaN value
Single floating point	0x7FBF FFFF
Double floating point	0x7FF7 FFFF FFFF FFFF
Word fixed point	0x7FFF FFFF
Longword fixed point	0x7FFF FFFF FFFF FFFF

### 14.4.2 Fixed-Point Formats

The FPU provides two fixed-point data types:


- A 32-bit Word fixed point (type W), shown in [Figure 14.3](#)
- A 64-bit Longword fixed point (type L), shown in [Figure 14.4](#)

---


1. In case of one or more QNaN operands, a QNaN is propagated from one of the operands according to the following priority: 1: fs, 2: ft, 3: fr.

The fixed-point values are held in 2's complement format, which is used for signed integers in the CPU. Unsigned fixed-point data types are not provided by the architecture; application software can synthesize computations for unsigned integers from the existing instructions and data types.

**Figure 14.3 Word Fixed-Point Format (W)**



**Figure 14.4 Longword Fixed-Point Format (L)**



## 14.5 Floating-Point General Registers

This section describes the organization and use of the Floating-Point general Registers (FPRs). The FPU is a 64b FPU, but a 32b register mode for backwards compatibility is also supported. The FR bit in the CP0 *Status* register determines which mode is selected:

- When the FR bit is a 1, the 64b register model is selected, which defines thirty-two 64-bit registers with all formats supported in a register.
- When the FR bit is a 0, the 32b register model is selected, which defines thirty-two 32-bit registers with D-format values stored in even-odd pairs of registers; thus the register file can also be viewed as having sixteen 64-bit registers.
- When configured this way, there are several restrictions for double operation:
  - Any double operations which specify an odd register as a source or destination will cause a ReservedInstruction exception
  - MTHC1/MFHC1 instructions which access an odd FPU register will signal a Reserved Instruction exception.


These registers transfer binary data between the FPU and the system, and are also used to hold formatted FPU operand values.

### 14.5.1 FPRs and Formatted Operand Layout


FPU instructions that operate on formatted operand values specify the Floating-Point Register (FPR) that holds the value. Operands that are only 32 bits wide (*W* and *S* formats) use only half the space in an FPR.

Figures 14.5 and 14.6 show the FPR organization and the way that operand data is stored in them.

**Figure 14.5 Single Floating-Point or Word Fixed-Point Operand in an FPR**



**Figure 14.6 Double Floating-Point or Longword Fixed-Point Operand in an FPR**



## 14.5.2 Formats of Values Used in Floating Point Registers

Unlike the CPU, the FPU neither interprets the binary encoding of source operands nor produces a binary encoding of results for every operation. The value held in a floating-point operand register (FPR) has a format, or type, and it can be used only by instructions that operate on that format. The format of a value is either *uninterpreted*, *unknown*, or one of the valid numeric formats: *single* or *double* floating point, and *word* or *long* fixed point.

The value in an FPR is always set when a value is written to the register as follows:

- When a data transfer instruction writes binary data into an FPR (a load), the FPR receives a binary value that is *uninterpreted*.
- A computational or FP register move instruction that produces a result of type *fmt* puts a value of type *fmt* into the result register.

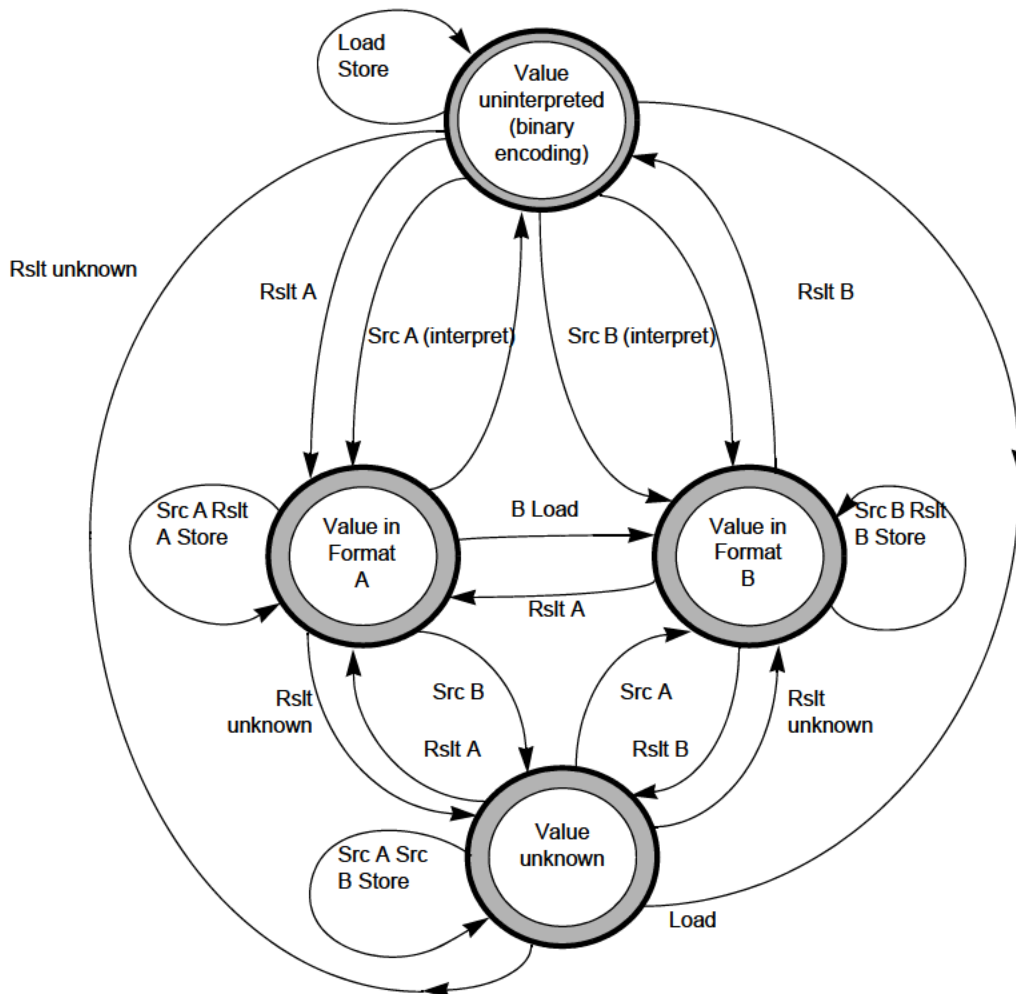
When an FPR with an *uninterpreted* value is used as a source operand by an instruction that requires a value of format *fmt*, the binary contents are interpreted as an encoded value in format *fmt*, and the value in the FPR changes to a value of format *fmt*. The binary contents cannot be reinterpreted in a different format.

If an FPR contains a value of format *fmt*, a computational instruction must not use the FPR as a source operand of a different format. If this case occurs, the value in the register becomes *unknown*, and the result of the instruction is also a value that is *unknown*. Using an FPR containing an *unknown* value as a source operand produces a result that has an *unknown* value.

The format of the value in the FPR is unchanged when it is read by a data transfer instruction (a store). A data transfer instruction produces a binary encoding of the value contained in the FPR. If the value in the FPR is *unknown*, the encoded binary value produced by the operation is not defined.

The state diagram in [Figure 14.7](#) illustrates the manner in which the formatted value in an FPR is set and changed.

**Figure 14.7 Effect of FPU Operations on the Format of Values Held in FPRs**



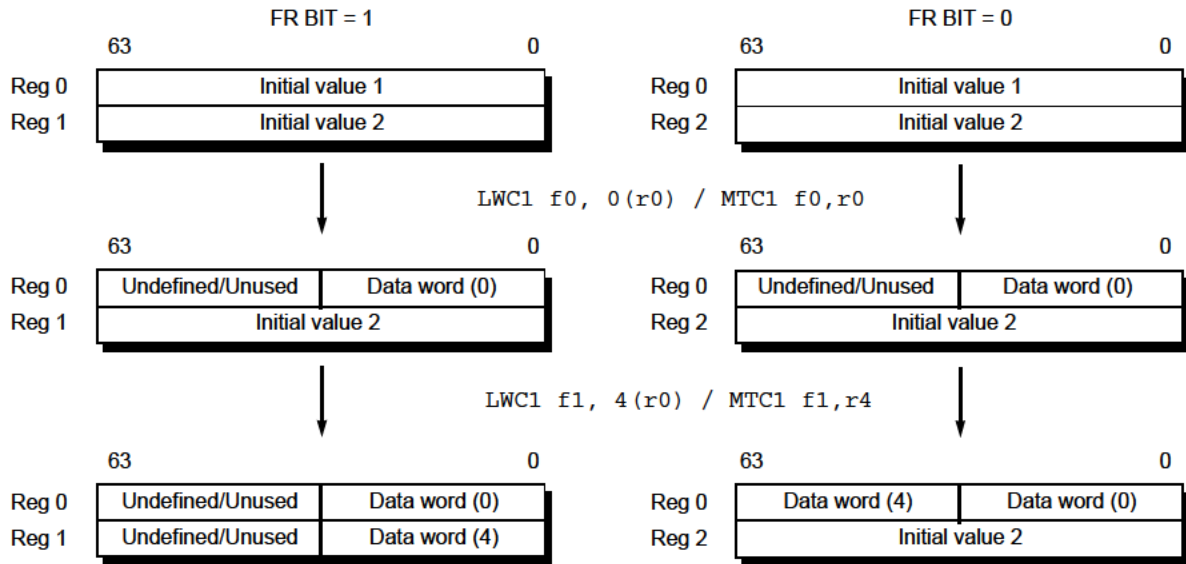
A, B: Example formats  
 Load: Destination of LWC1, LDC1, MTC1 instructions.  
 Store: Source operand of SWC1, SDC1, MFC1 instructions.  
 Src fmt: Source operand of computational instruction expecting format "fmt."  
 Rslt fmt: Result of computational instruction producing value of format "fmt."

### 14.5.3 Binary Data Transfers (32-Bit and 64-Bit)

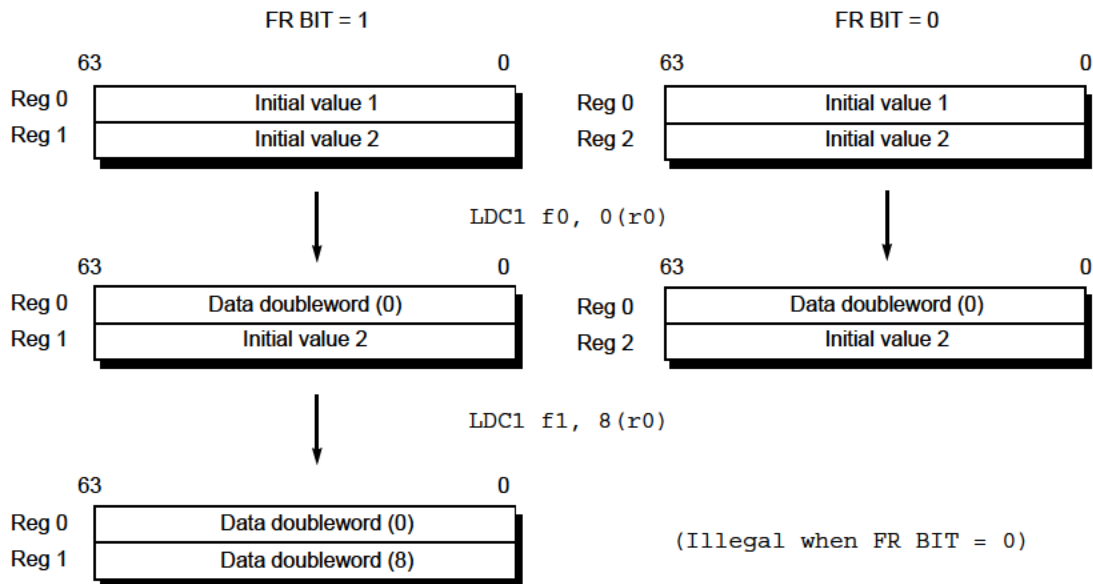
The data transfer instructions move words and doublewords between the FPU FPRs and the remainder of the system. The operations of the word and doubleword load and move-to instructions are shown in [Figure 14.8](#) and [Figure 14.9](#), respectively.

The store and move-from instructions operate in reverse, reading data from the location that the corresponding load or move-to instruction had written.

**Figure 14.8 FPU Word Load and Move-to Operations**



**Figure 14.9 FPU Doubleword Load and Move-to Operations**



## 14.6 Floating-Point Control Registers

The FPU Control Registers (FCRs) identify and control the FPU. The five FPU control registers are 32 bits wide: *FIR*, *FCCR*, *FEXR*, *FENR*, *FCSR*. Three of these registers, *FCCR*, *FEXR*, and *FENR*, select subsets of the floating-point Control/Status register, the *FCSR*. These registers are also denoted Coprocessor 1 (CP1) control registers.

CPI control registers are summarized in [Table 14.4](#) and are described individually in the following subsections of this chapter. Each register’s description includes the read/write properties and the reset state of each field.

**Table 14.4 Coprocessor 1 Register Summary**

Register Number	Register Name	Function
0	FIR	Floating-Point Implementation register. Contains information that identifies the FPU.
25	FCCR	Floating-Point Condition Codes register.
26	FEXR	Floating-Point Exceptions register.
28	FENR	Floating-Point Enables register.
31	FCSR	Floating-Point Control and Status register.

[Table 14.5](#) defines the notation used for the read/write properties of the register bit fields.


**Table 14.5 Read/Write Properties**

Read/Write Notation	Hardware Interpretation	Software Interpretation
R/W	All bits in this field are readable and writable by software and potentially by hardware. Hardware updates of this field are visible by software reads. Software updates of this field are visible by hardware reads. If the reset state of this field is “Undefined,” either software or hardware must initialize the value before the first read returns a predictable value. This definition should not be confused with the formal definition of UNDEFINED behavior.	
R	This field is either static or is updated only by hardware. If the Reset State of this field is either “0” or “Preset”, hardware initializes this field to zero or to the appropriate state, respectively, on powerup. If the Reset State of this field is “Undefined”, hardware updates this field only under those conditions specified in the description of the field.	A field to which the value written by software is ignored by hardware. Software may write any value to this field without affecting hardware behavior. Software reads of this field return the last value updated by hardware. If the Reset State of this field is “Undefined,” software reads of this field result in an UNPREDICTABLE value except after a hardware update done under the conditions specified in the description of the field.
0	Hardware does not update this field. Hardware can assume a zero value.	The value software writes to this field must be zero. Software writes of non-zero values to this field might result in UNDEFINED behavior of the hardware. Software reads of this field return zero as long as all previous software writes are zero. If the Reset State of this field is “Undefined,” software must write this field with zero before it is guaranteed to read as zero.

### 14.6.1 Floating-Point Implementation Register (FIR, CP1 Control Register 0)

The Floating-Point Implementation Register (*FIR*) is a 32-bit read-only register that contains information identifying the capabilities of the FPU, the Floating-Point processor identification, and the revision level of the FPU. [Figure 14.10](#) shows the format of the *FIR*; [Table 14.6](#) describes the *FIR* bit fields.

**Figure 14.10 FIR Format**





**Table 14.6 FIR Bit Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:25	Reserved.	R	0
FC	24	Indicates that full convert ranges are implemented: <ul style="list-style-type: none"> <li>• 0: Full convert ranges not implemented</li> <li>• 1: Full convert ranges implemented</li> </ul> This bit is always 1 to indicate that full convert ranges are implemented. This means that all numbers can be converted to another type by the FPU (If FS bit in FCSR is not set Unimplemented Operation exception can still happen on denormal operands though).	R	1
0	23	These bits must be written as zeros; they return zeros on reads.	0	0
F64	22	Indicates that this is a 64-bit FPU: <ul style="list-style-type: none"> <li>• 0: Not a 64-bit FPU</li> <li>• 1: A 64-bit FPU.</li> </ul> This bit is always 1 to indicate that this is a 64-bit FPU.	R	1
L	21	Indicates that the long fixed point (L) data type and instructions are implemented: <ul style="list-style-type: none"> <li>• 0: Long type not implemented</li> <li>• 1: Long implemented</li> </ul> This bit is always 1 to indicate that long fixed point data types are implemented.	R	1
W	20	Indicates that the word fixed point (W) data type and instructions are implemented: <ul style="list-style-type: none"> <li>• 0: Word type not implemented</li> <li>• 1: Word implemented</li> </ul> This bit is always 1 to indicate that word fixed point data types are implemented.	R	1
3D	19	Indicates that the MIPS-3D ASE is implemented: <ul style="list-style-type: none"> <li>• 0: MIPS-3D not implemented</li> <li>• 1: MIPS-3D implemented</li> </ul> This bit is always 0 to indicate that MIPS-3D is not implemented.	R	0
PS	18	Indicates that the paired-single (PS) floating-point data type and instructions are implemented: <ul style="list-style-type: none"> <li>• 0: PS floating-point not implemented</li> <li>• 1: PS floating-point implemented</li> </ul> This bit is always 0 to indicate that paired-single floating-point data types are not implemented.	R	0
D	17	Indicates that the double-precision (D) floating-point data type and instructions are implemented: <ul style="list-style-type: none"> <li>• 0: D floating-point not implemented</li> <li>• 1: D floating-point implemented</li> </ul> This bit is always 1 to indicate that double-precision floating-point data types are implemented.	R	1


**Table 14.6 FIR Bit Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
S	16	Indicates that the single-precision (S) floating-point data type and instructions are implemented: <ul style="list-style-type: none"> <li>• 0: S floating-point not implemented</li> <li>• 1: S floating-point implemented</li> </ul> This bit is always 1 to indicate that single-precision floating-point data types are implemented.	R	1
Processor ID	15:8	Identifies the floating-point processor.	R	
Revision	7:0	Specifies the revision number of the FPU. This field allows software to distinguish between different revisions of the same floating-point processor type.	R	Hardwired

### 14.6.2 Floating-Point Condition Codes Register (FCCR, CP1 Control Register 25)

The Floating-Point Condition Codes Register (*FCCR*) is an alternative way to read and write the floating-point condition code values that also appear in the *FCSR*. Unlike the *FCSR*, all eight FCC bits are contiguous in the *FCCR*. [Figure 14.11](#) shows the format of the *FCCR*; [Table 14.7](#) describes the *FCCR* bit fields.

**Figure 14.11 FCCR Format**




**Table 14.7 FCCR Bit Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
FCC	7:0	Floating-point condition code. Refer to the description of this field in <a href="#">Section 14.6.5 “Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)”</a> .	R/W	Undefined
0	31:8	These bits must be written as zeros; they return zeros on reads.	0	0

### 14.6.3 Floating-Point Exceptions Register (FEXR, CP1 Control Register 26)

The Floating-Point Exceptions Register (*FEXR*) is an alternative way to read and write the Cause and Flags fields that also appear in the *FCSR*. [Figure 14.12](#) shows the format of the *FEXR*; [Table 14.8](#) describes the *FEXR* bit fields.

**Figure 14.12 FEXR Format**




**Table 14.8 FEXR Bit Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:18	These bits must be written as zeros; they return zeros on reads.	0	0
Cause	17:12	Cause bits. Refer to the description of this field in <a href="#">Section 14.6.5, "Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)"</a> .	R/W	Undefined
0	11:7	These bits must be written as zeros; they return zeros on reads.	0	0
Flags	6:2	Flag bits. Refer to the description of this field in <a href="#">Section 14.6.5, "Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)"</a> .	R/W	Undefined
0	1:0	These bits must be written as zeros; they return zeros on reads.	0	0

#### 14.6.4 Floating-Point Enables Register (FENR, CP1 Control Register 28)

The Floating-Point Enables Register (*FENR*) is an alternative way to read and write the Enables, FS, and RM fields that also appear in the *FCSR*. [Figure 14.13](#) shows the format of the *FENR*; [Table 14.9](#) describes the *FENR* bit fields.

**Figure 14.13 FENR Format**



**Table 14.9 FENR Bit Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:12	These bits must be written as zeros; they return zeros on reads.	0	0
Enables	11:7	Enable bits. Refer to the description of this field in <a href="#">Section 14.6.5, "Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)"</a> .	R/W	Undefined
0	6:3	These bits must be written as zeros; they return zeros on reads.	0	0
FS	2	Flush to Zero bit. Refer to the description of this field in <a href="#">Section 14.6.5, "Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)"</a> .	R/W	Undefined
RM	1:0	Rounding mode. Refer to the description of this field in <a href="#">Section 14.6.5, "Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)"</a> .	R/W	Undefined


## 14.6.5 Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)

The 32-bit Floating-Point Control and Status Register (*FCSR*) controls the operation of the FPU and shows the following status information:

- Selects the default rounding mode for FPU arithmetic operations
- Selectively enables traps of FPU exception conditions
- Controls some denormalized number handling options
- Reports any IEEE exceptions that arose during the most recently executed instruction
- Reports any IEEE exceptions that cumulatively arose in completed instructions
- Indicates the condition code result of FP compare instructions

Access to the *FCSR* is not privileged; it can be read or written by any program that has access to the FPU (via the coprocessor enables in the *Status* register). Figure 14.14 shows the format of the *FCSR*; Table 14.10 describes the *FCSR* bit fields.

**Figure 14.14 FCSR Format**



**Table 14.10 FCSR Bit Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit			
FCC	31:25, 23	Floating-point condition codes. These bits record the result of floating-point compares and are tested for floating-point conditional branches and conditional moves. The FCC bit to use is specified in the compare, branch, or conditional move instruction. For backward compatibility with previous MIPS ISAs, the FCC bits are separated into two non-contiguous fields.	R/W	Undefined
FS	24	Flush to Zero (FS). The FS bit controls the handling of denormalized operands and is encoded as follows:  0: IEEE-compliant mode. Low performance on denormal operands and tiny results. 1: Regular embedded applications. High performance on denormal operands and tiny results.	R/W	Undefined
FO	22	Flush Override (FO). Refer to <a href="#">Section 14.6.6 “Operation of the FS/FO/FN Bits”</a> for more details on this bit.	R/W	Undefined
FN	21	Flush to Nearest (FN). Refer to <a href="#">Section 14.6.6 “Operation of the FS/FO/FN Bits”</a> for more details on this bit.	R/W	Undefined
0	20:18	These bits must be written as zeros; they return zeros on reads.	0	0

**Table 14.10 FCSR Bit Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State
Name	Bit			
Cause	17:12	Cause bits. These bits indicate the exception conditions that arise during execution of an FPU arithmetic instruction. A bit is set to 1 when the corresponding exception condition arises during the execution of an instruction; otherwise, it is cleared to 0. By reading the registers, the exception condition caused by the preceding FPU arithmetic instruction can be determined. Refer to <a href="#">Table 14.11</a> for the meaning of each cause bit.	R/W	Undefined
Enables	11:7	Enable bits. These bits control whether or not a trap is taken when an IEEE exception condition occurs for any of the five conditions. The trap occurs when both an enable bit and its corresponding cause bit are set either during an FPU arithmetic operation or by moving a value to the <i>FCSR</i> or one of its alternative representations. Note that Cause bit E (CauseE) has no corresponding enable bit; the MIPS architecture defines non-IEEE Unimplemented Operation exceptions as always enabled. Refer to <a href="#">Table 14.11</a> for the meaning of each enable bit.	R/W	Undefined
Flags	6:2	Flag bits. This field shows any exception conditions that have occurred for completed instructions since the flag was last reset by software. When an FPU arithmetic operation raises an IEEE exception condition that does not result in a Floating-Point Exception (the enable bit was off), the corresponding bit(s) in the Flags field are set, while the others remain unchanged. Arithmetic operations that result in a Floating-Point Exception (the enable bit was on) do not update the Flags field. Hardware never resets this field; software must explicitly reset this field. Refer to <a href="#">Table 14.11</a> for the meaning of each flag bit.	R/W	Undefined
RM	1:0	Rounding mode. This field indicates the rounding mode used for most floating-point operations (some operations use a specific rounding mode). Refer to <a href="#">Table 14.12</a> for the encoding of this field.	R/W	Undefined

**Table 14.11 Cause, Enables, and Flags Definitions**

Bit Name	Bit Meaning
E	Unimplemented Operation (this bit exists only in the Cause field).
V	Invalid Operations
Z	Divide by Zero
O	Overflow
U	Underflow
I	Inexact

**Table 14.12 Rounding Mode Definitions**

RM Field Encoding	Meaning
0	RN - Round to Nearest Rounds the result to the nearest representable value. When two representable values are equally near, the result is rounded to the value whose least-significant bit is zero (even).
1	RZ - Round Toward Zero Rounds the result to the value closest to but not greater in magnitude than the result.
2	RP - Round Towards Plus Infinity Rounds the result to the value closest to but not less than the result.
3	RM - Round Towards Minus Infinity Rounds the result to the value closest to but not greater than the result.

### 14.6.6 Operation of the FS/FO/FN Bits

The FS, FO, and FN bits in the CP1 *FCSR* register control handling of denormalized operands and *tiny* results (i.e. nonzero result between  $\pm 2^{E-min}$ ), whereby the FPU can handle these cases right away instead of relying on the much slower software handler. The trade-off is a loss of IEEE compliance and accuracy (except for use of the FO bit), because a minimal normalized or zero result is provided by the FPU instead of the more accurate denormalized result that a software handler would give. The benefit is a significantly improved performance and precision.

Use of the FS, FO, and FN bits affects handling of denormalized floating-point numbers and tiny results for the instructions listed below:

FS and FN bit:     ADD, CEIL, CVT, DIV, FLOOR, MADD, MSUB, MUL, NMADD, NMSUB, RECIP, ROUND, RSQRT, SQRT, TRUNC, SUB, ABS, C.cond, and NEG<sup>1</sup>


FO bit:             MADD, MSUB, NMADD, and NMSUB

1. For ABS, C.cond, and NEG, denormal input operands or tiny results do not result in Unimplemented exceptions when FS = 0. Flushing to zero nonetheless is implemented when FS = 1 such that these operations return the same result as an equivalent sequence of arithmetic FPU operations.

Instructions not listed above do not cause Unimplemented Operation exceptions on denormalized numbers in operands or results.

Figure 14.15 depicts how the FS, FO, and FN bits control handling of denormalized numbers. For instructions that are not multiply or add types (such as DIV), only the FS and FN bits apply.

**Figure 14.15 FS/FO/FN Bits Influence on Multiply and Addition Results**



#### 14.6.6.1 Flush To Zero Bit

When the Flush To Zero (FS) bit is set, denormal input operands are flushed to zero. Tiny results are flushed to either zero or the applied format’s smallest normalized number (MinNorm) depending on the rounding mode settings. The following table lists the flushing behavior for tiny results..

**Table 14.13 Zero Flushing for Tiny Results**

Rounding Mode	Negative Tiny Result	Positive Tiny Result
RN (RM=0)	-0	+0
RZ(RM=1)	-0	+0
RP (RM=2)	-0	+MinNorm
RM (RM=3)	-MinNorm	+0

The flushing of results is based on an intermediate result computed by rounding the mantissa using an unbounded exponent range; that is, tiny numbers are not *normalized* into the supported exponent range by shifting in leading zeros prior to rounding.

Handling of denormalized operand values and tiny results depends on the FS bit setting as shown in [Table 14.14](#).

**Table 14.14 Handling of Denormalized Operand Values and Tiny Results Based on FS Bit Setting**

FS Bit	Handling of Denormalized Operand Values
0	An Unimplemented Operation exception is taken.
1	Instead of causing an Unimplemented Operation exception, operands are flushed to zero, and tiny results are forced to zero or MinNorm.

#### 14.6.6.2 Flush Override Bit

When the Flush Override (FO) bit is set, a tiny intermediate result of any multiply-add type instruction is not flushed according to the FS bit. The intermediate result is maintained in an internal normalized format to improve accuracy. FO only applies to the intermediate result of a multiply-add type instruction.

Handling of tiny intermediate results depends on the FO and FS bits as shown in [Table 14.15](#).

**Table 14.15 Handling of Tiny Intermediate Result Based on the FO and FS Bit Settings**

FO Bit	FS Bit	Handling of Tiny Result Values
0	0	An Unimplemented Operation exception is taken.


**Table 14.15 Handling of Tiny Intermediate Result Based on the FO and FS Bit Settings**

FO Bit	FS Bit	Handling of Tiny Result Values
0	1	The intermediate result is forced to the value that would have been delivered for an untrapped underflow (see Table 14.18) instead of causing an Unimplemented Operation exception.
1	Don't care	The intermediate result is kept in an internal format, which can be perceived as having the usual mantissa precision but with unlimited exponent precision and without forcing to a specific value or taking an exception.

### 14.6.6.3 Flush to Nearest

When the Flush to Nearest (FN) bit is set and the rounding mode is Round to Nearest (RN), a tiny final result is flushed to zero or MinNorm. If a tiny number is strictly below MinNorm/2, the result is flushed to zero; otherwise, it is flushed to MinNorm (see Figure 14.16). The flushed result has the same sign as the result prior to flushing. Note that the FN bit takes precedence over the FS bit.

**Figure 14.16 Flushing to Nearest when Rounding Mode is Round to Nearest**



For all rounding modes other than Round to Nearest (RN), setting the FN bit causes final results to be flushed to zero or MinNorm as if the FS bit was set.

Handling of tiny final results depends on the FN and FS bits as shown in Table 14.16.

**Table 14.16 Handling of Tiny Final Result Based on FN and FS Bit Settings**

FN Bit	FS Bit	Handling of Tiny Result Values
0	0	An Unimplemented Operation exception is taken.
0	1	Final result is forced to the value that would have been delivered for an untrapped underflow (see Table 14.18) rather than causing an Unimplemented Operation exception.
1	Don't care	Final result is rounded to either zero or $2^{E_{\min}}$ (MinNorm), whichever is closest when in Round to Nearest (RN) rounding mode. For other rounding modes, a final result is given as if FS was set to 1.

### 14.6.6.4 Recommended FS/FO/FN Settings

Table 14.17 summarizes the recommended FS/FO/FN settings.

**Table 14.17 Recommended FS/FO/FN Settings**

FS Bit	FO Bit	FN Bit	Remarks
0	0	0	IEEE-compliant mode. Low performance on denormal operands and tiny results.



**Table 14.17 Recommended FS/FO/FN Settings**

FS Bit	FO Bit	FN Bit	Remarks
1	0	0	Regular embedded applications. High performance on denormal operands and tiny results.
1	1	1	Highest accuracy and performance configuration. <sup>1</sup>

1. Note that in this mode, MADD might return a different result other than the equivalent MUL and ADD operation sequence.

## 14.6.7 FCSR Cause Bit Update Flow

### 14.6.7.1 Exceptions Triggered by CTC1

Regardless of the targeted control register, the **CTC1** instruction causes the Enables and Cause fields of the *FCSR* to be inspected in order to determine if an exception is to be thrown.

### 14.6.7.2 Generic Flow

Computations are performed in two steps:

1. Compute rounded mantissa with unbound exponent range.
2. Flush to default result if the result from Step #1 above is overflow or tiny (no flushing happens on denorms for instructions supporting denorm results, such as MOV).

The Cause field is updated after each of these two steps. Any enabled exceptions detected in these two steps cause a trap, and no further updates to the Cause field are done by subsequent steps.

Step #1 can set cause bits I, U, O, Z, V, and E. E has priority over V; V has priority over Z; and Z has priority over U and O. Thus when E, V, or Z is set in Step #1, no other cause bits can be set. However, note that I and V both can be set if a denormal operand was flushed (FS = 1). I, U, and O can be set alone or in pairs (IU or IO). U and O never can be set simultaneously in Step #1. U and O are set if the computed unbounded exponent is outside the exponent range supported by the normalized IEEE format.

Step #2 can set I if a default result is generated.

### 14.6.7.3 Multiply-Add Flow

For Multiply-Add type instructions, the computation is extended with two more steps:

1. Compute rounded mantissa with unbound exponent range for the multiply.
2. Flush to default result if the result from Step #1 is overflow or tiny.
3. Compute rounded mantissa with unbounded exponent range for the add.
4. Flush to default result if the result from Step #3 is overflow or tiny.

The Cause field is updated after each of these four steps. Any enabled exceptions detected in these four steps cause a trap, and no further updates to the Cause field are done by subsequent steps.

Step #1 and Step #3 can set a cause bit as described for Step #1 in [14.6.7.2 “Generic Flow”](#).

Step #2 and Step #4 can set I if a default result is generated.

Although U and O can never both be set in Step #1 or Step #3, both U and O might be set after the multiply-add has executed in Step #3 because U might be set in Step #1 and O might be set in Step #3.

#### 14.6.7.4 Cause Update Flow for Input Operands

Denormal input operands to Step #1 or Step #3 always set Cause bit I when FS = 1. For example, SNaN+DeNorm set I (and V) provided that Step #3 was reached (in case of a multiply-add type instruction).

Conditions directly related to the input operand (for example, I/E set due to DeNorm, V set due to SNaN and QNaN propagation) are detected in the step where the operand is logically used. For example, for multiply-add type instructions, exceptional conditions caused by the input operand fr are detected in Step #3.

#### 14.6.7.5 Cause Update Flow for Unimplemented Operations

Note that Cause bit E is special; it clears any Cause updates done in previous steps. For example, if Step #3 caused E to be set, any I, U, or O Cause update done in Step #1 or Step #2 is cleared. Only E is set in the Cause field when an Unimplemented Operation trap is taken.

## 14.7 Exceptions

FPU exceptions are implemented in the MIPS FPU architecture with the Cause, Enables, and Flags fields of the *FCSR*. The flag bits implement IEEE exception status flags, and the cause and enable bits control exception trapping. Each field has a bit for each of the five IEEE exception conditions. The Cause field has an additional exception bit, Unimplemented Operation, used to trap for software emulation assistance. If an exception type is enabled through the Enables field of the *FCSR*, then the FPU is operating in precise exception mode for this type of exception.

### 14.7.1 Precise Exception Mode

In precise exception mode, a trap occurs before the instruction that causes the trap or any following instruction can complete and write its results. If desired, the software trap handler can resume execution of the interrupted instruction stream after handling the exception.

The Cause field reports per-bit instruction exception conditions. The cause bits are written during each floating-point arithmetic operation to show any exception conditions that arise during the operation. A cause bit is set to 1 if its corresponding exception condition arises; otherwise, it is cleared to 0.

A floating-point trap is generated any time both a cause bit and its corresponding enable bit are set. This case occurs either during the execution of a floating-point operation or when moving a value into the *FCSR*. There is no enable bit for Unimplemented Operations; this exception always generates a trap.

In a trap handler, exception conditions that arise during any trapped floating-point operations are reported in the Cause field. Before returning from a floating-point interrupt or exception, or before setting cause bits with a move to the *FCSR*, software first must clear the enabled cause bits by executing a move to the *FCSR* to prevent the trap from being erroneously retaken.

If a floating-point operation sets only non-enabled cause bits, no trap occurs and the default result defined by IEEE Standard 754 is stored (see [Table 14.18](#)). When a floating-point operation does not trap, the program can monitor the exception conditions by reading the Cause field.

The Flags field is a cumulative report of IEEE exception conditions that arise as instructions complete; instructions that trap do not update the flag bits. The flag bits are set to 1 if the corresponding IEEE exception is raised, otherwise the bits are unchanged. There is no flag bit for the MIPS Unimplemented Operation exception. The flag bits are never cleared as a side effect of floating-point operations, but they can be set or cleared by moving a new value into the *FCSR*.

## 14.7.2 Exception Conditions

The subsections below describe the following five exception conditions defined by IEEE Standard 754:

- [Section 14.7.2.1 “Invalid Operation Exception”](#)
- [Section 14.7.2.2 “Division By Zero Exception”](#)
- [Section 14.7.2.3 “Underflow Exception”](#)
- [Section 14.7.2.4 “Overflow Exception”](#)
- [Section 14.7.2.5 “Inexact Exception”](#)
- [Section 14.7.2.6 “Unimplemented Operation Exception”](#)

At the program’s direction, an IEEE exception condition can either cause a trap or not cause a trap. IEEE Standard 754 specifies the result to be delivered in case no trap is taken. The FPU supplies these results whenever the exception condition does not result in a trap. The default action taken depends on the type of exception condition and, in the case of the Overflow and Underflow, the current rounding mode. [Table 14.18](#) summarizes the default results.

**Table 14.18 Result for Exceptions Not Trapped**

Bit	Description	Default Action
V	Invalid Operation	Supplies a quiet NaN.
Z	Divide by zero	Supplies a properly signed infinity.
U	Underflow	Depends on the rounding mode as shown below: <ul style="list-style-type: none"> <li>• 0 (RN) and 1 (RZ): Supplies a zero with the sign of the exact result.</li> <li>• 2 (RP): For positive underflow values, supplies <math>2^{E_{\min}}</math> (MinNorm). For negative underflow values, supplies a positive zero.</li> <li>• 3 (RM): For positive underflow values, supplies a negative zero. For negative underflow values, supplies a negative <math>2^{E_{\min}}</math> (MinNorm).</li> </ul> Note that this behavior is only valid if the <i>FCSR</i> <sub>FN</sub> bit is cleared.
I	Inexact	Supplies a rounded result. If caused by an overflow without the overflow trap enabled, supplies the overflowed result. If caused by an underflow without the underflow trap enabled, supplies the underflowed result.
O	Overflow	Depends on the rounding mode, as shown below: <ul style="list-style-type: none"> <li>• 0 (RN): Supplies an infinity with the sign of the exact result.</li> <li>• 1 (RZ): Supplies the format’s largest finite number with the sign of the exact result.</li> <li>• 2 (RP): For positive overflow values, supplies positive infinity. For negative overflow values, supplies the format’s most negative finite number.</li> <li>• 3 (RM): For positive overflow values, supplies the format’s largest finite number. For negative overflow values, supplies minus infinity.</li> </ul>

### 14.7.2.1 Invalid Operation Exception

An Invalid Operation exception is signaled when one or both of the operands are invalid for the operation to be performed. When the exception condition occurs without a precise trap, the result is a quiet NaN.

The following operations are invalid:

- One or both operands are a signaling NaN (except for the non-arithmetic MOV.fmt, MOVT fmt, MOVF fmt, MOVN fmt, and MOVZ.fmt instructions).
- Addition or subtraction: magnitude subtraction of infinities, such as  $(+\infty) + (-\infty)$  or  $(-\infty) - (-\infty)$ .
- Multiplication:  $0 \times \infty$ , with any signs.
- Division:  $0/0$  or  $\infty/\infty$ , with any signs.
- Square root: An operand of less than 0 (-0 is a valid operand value).
- Conversion of a floating-point number to a fixed-point format when either an overflow or an operand value of infinity or NaN precludes a faithful representation in that format.
- Some comparison operations in which one or both of the operands is a QNaN value.

#### 14.7.2.2 Division By Zero Exception

The divide operation signals a Division By Zero exception if the divisor is zero and the dividend is a finite nonzero number. When no precise trap occurs, the result is a correctly signed infinity. Divisions ( $0/0$  and  $\infty/0$ ) do not cause the Division By Zero exception. The result of ( $0/0$ ) is an Invalid Operation exception. The result of ( $\infty/0$ ) is a correctly signed infinity.

#### 14.7.2.3 Underflow Exception

Two related events contribute to underflow:

- Tininess: The creation of a tiny, nonzero result between  $\pm 2^{E_{min}}$  which, because it is tiny, might cause some other exception later such as overflow on division. IEEE Standard 754 allows choices in detecting tininess events. The MIPS architecture specifies that tininess be detected after rounding, when a nonzero result computed as though the exponent range were unbounded would lie strictly between  $\pm 2^{E_{min}}$ .
- Loss of accuracy: The extraordinary loss of accuracy occurs during the approximation of such tiny numbers by denormalized numbers. IEEE Standard 754 allows choices in detecting loss of accuracy events. The MIPS architecture specifies that loss of accuracy be detected as inexact result, when the delivered result differs from what would have been computed if both the exponent range and precision were unbounded.

The way that an underflow is signaled depends on whether or not underflow traps are enabled:

- When an underflow trap is not enabled, underflow is signaled only when both tininess and loss of accuracy have been detected. The delivered result might be zero, denormalized, or  $\pm 2^{E_{min}}$ .
- When an underflow trap is enabled (through the *FCSR* Enables field), underflow is signaled when tininess is detected regardless of loss of accuracy.

#### 14.7.2.4 Overflow Exception

An Overflow exception is signaled when the magnitude of a rounded floating-point result (if the exponent range is unbounded) is larger than the destination format's largest finite number.

When no precise trap occurs, the result is determined by the rounding mode and the sign of the intermediate result.

#### 14.7.2.5 Inexact Exception

An Inexact exception is signaled when one of the following occurs:

- The rounded result of an operation is not exact.
- The rounded result of an operation overflows without an overflow trap.
- When a denormal operand is flushed to zero.

#### 14.7.2.6 Unimplemented Operation Exception

The Unimplemented Operation exception is a MIPS-defined exception that provides software emulation support. This exception is not IEEE-compliant and is used to signal a need for software emulation of an instruction. Normally an IEEE arithmetic operation can cause only one exception condition; the only case in which two exceptions can occur at the same time are *Inexact With Overflow* and *Inexact With Underflow*.

The MIPS architecture is designed so that a combination of hardware and software can implement the architecture. Operations not fully supported in hardware cause an Unimplemented Operation exception, allowing software to perform the operation.

There is no enable bit for this condition; it always causes a trap (but the condition is effectively masked for all operations when FS=1). After the appropriate emulation or other operation is done in a software exception handler, the original instruction stream can be continued.

An Unimplemented Operation exception is taken in the following situations:

- when denormalized operands or tiny results are encountered for instructions not supporting denormal numbers and where such are not handled by the FS bit.

## 14.8 Latency and Repeat Rates

Table 14.19 shows the repeat rate and latency for the FPU instructions. Note that cycles related to floating point operations are listed in terms of FPU clocks.

**Table 14.19 interAptiv FPU Latency and Repeat Rates**

Opcode <sup>1</sup>	Latency (cycles)	Repeat Rate (cycles)
ABS.[S,D], NEG.[S,D], ADD.[S,D], SUB.[S,D], MUL.S, MADD.S, MSUB.S, NMADD.S, NMSUB.S	4	1
MUL.D, MADD.D, MSUB.D, NMADD.D, NMSUB.D	5	2
RECIP.S	13	10
RECIP.D	25	21
RSQRT.S	17	14
RSQRT.D	35	31
DIV.S, SQRT.S	17	14
DIV.D, SQRT.D	32	29
C.cond.[S,D] to MOVF fmt and MOVT fmt instruction / MOVT, MOVN, BC1 instruction	1 / 2	1
CVT.D.S, CVT.[S,D].[W,L]	4	1
CVT.S.D	6	1

**Table 14.19 interAptiv FPU Latency and Repeat Rates (continued)**

Opcode <sup>1</sup>	Latency (cycles)	Repeat Rate (cycles)
CVT.[W,L].[S,D], CEIL.[W,L].[S,D], FLOOR.[W,L].[S,D], ROUND.[W,L].[S,D], TRUNC.[W,L].[S,D]	5	1
MOV.[S,D], MOVF.[S,D], MOVN.[S,D], MOVT.[S,D], MOVZ.[S,D]	4	1
LWC1, LDC1, LDXC1, LUXC1, LWXC1	3	1
MTC1, MFC1	2	1

1. Format: S = Single, D = Double, W = Word, L = Longword.

## 14.9 Instruction Overview

The functional groups into which the FPU instructions are divided are described in the following subsections:

- [Section 14.9.1 “Data Transfer Instructions”](#)
- [Section 14.9.2 “Arithmetic Instructions”](#)
- [Section 14.9.3 “Conversion Instructions”](#)
- [Section 14.9.4 “Formatted Operand-Value Move Instructions”](#)
- [Section 14.9.5 “Conditional Branch Instructions”](#)
- [Section 14.9.6 “Miscellaneous Instructions”](#)

### 14.9.1 Data Transfer Instructions

The FPU has two separate register sets: floating point coprocessor general registers (FPRs) and floating point coprocessor control registers (FCRs). The FPU has a load/store architecture; all computations are done on data held in coprocessor general registers. The control registers are used to control FPU operation. Data is transferred between registers and the rest of the system with dedicated load, store, and move instructions. The transferred data is treated as unformatted binary data; no format conversions are performed, and therefore no IEEE floating-point exceptions can occur.

[Table 14.20](#) lists the supported transfer operations.

**Table 14.20 FPU Data Transfer Instructions**

Transfer Direction	Data Transferred
FPU general register ↔ Memory	Word/doubleword load/store
FPU general register ↔ CPU general register	Word move
FPU control register ↔ CPU general register	Word move

#### 14.9.1.1 Data Alignment in Loads, Stores, and Moves

All coprocessor loads and stores operate on naturally aligned data items. An attempt to load or store to an address that is not naturally aligned for the data item causes an Address Error exception. Regardless of byte ordering (the endianness),

ness), the address of a word or doubleword is the smallest byte address in the object. For a big-endian machine, this is the most-significant byte; for a little-endian machine, this is the least-significant byte.

### 14.9.1.2 Addressing Used in Data Transfer Instructions

The FPU has loads and stores using the same register+offset addressing as that used by the CPU. Moreover, for the FPU only, there are load and store instructions using *register+register* addressing.

Tables 14.21 through 14.23 list the FPU data transfer instructions.

**Table 14.21 FPU Loads and Stores Using Register+Offset Address Mode**

Mnemonic	Instruction
LDC1	Load Doubleword to Floating Point
LWC1	Load Word to Floating Point
SDC1	Store Doubleword to Floating Point
SWC1	Store Word to Floating Point

**Table 14.22 FPU Loads and Stores Using Register+Register Address Mode**

Mnemonic	Instruction
LDXC1	Load Doubleword Indexed to Floating Point
LUXC1	Load Doubleword Indexed Unaligned to Floating Point
LWXC1	Load Word Indexed to Floating Point
SDXC1	Store Doubleword Indexed to Floating Point
SUXC1	Store Doubleword Indexed Unaligned to Floating Point
SWXC1	Store Word Indexed to Floating Point

**Table 14.23 FPU Move To and From Instructions**

Mnemonic	Instruction
CFC1	Move Control Word From Floating Point
CTC1	Move Control Word To Floating Point
MFC1	Move Word From Floating Point
MTC1	Move Word To Floating Point

## 14.9.2 Arithmetic Instructions

Arithmetic instructions operate on formatted data values. The results of most floating-point arithmetic operations meet IEEE Standard 754 for accuracy—a result is identical to an infinite-precision result that has been rounded to the specified format using the current rounding mode. The rounded result differs from the exact result by less than one Unit in the Least-significant Place (ULP).

Table 14.24 lists the FPU IEEE compliant arithmetic operations.

**Table 14.24 FPU IEEE Arithmetic Operations**

Mnemonic	Instruction
ABS.fmt	Floating-Point Absolute Value
ADD.fmt	Floating-Point Add
C.cond.fmt	Floating-Point Compare
DIV.fmt	Floating-Point Divide
MUL.fmt	Floating-Point Multiply
NEG.fmt	Floating-Point Negate
SQRT.fmt	Floating-Point Square Root
SUB.fmt	Floating-Point Subtract

The two low latency operations, Reciprocal Approximation (RECIP) and Reciprocal Square Root Approximation (RSQRT), might be less accurate than the IEEE specification:

- The result of RECIP differs from the exact reciprocal by no more than one ULP.
- The result of RSQRT differs from the exact reciprocal square root by no more than two ULPs.

Table 14.25 lists the FPU-approximate arithmetic operations.

**Table 14.25 FPU-Approximate Arithmetic Operations**

Mnemonic	Instruction
RECIP.fmt	Floating-Point Reciprocal Approximation
RSQRT.fmt	Floating-Point Reciprocal Square Root Approximation

Four compound-operation instructions perform variations of multiply-accumulate operations; that is, multiply two operands, accumulate the result to a third operand, and produce a result. These instructions are listed in Table 14.26. The product is rounded according to the current rounding mode prior to the accumulation. This model meets the IEEE accuracy specification; the result is numerically identical to an equivalent computation using multiply, add, or subtract instructions.

**Table 14.26 FPU Multiply-Accumulate Arithmetic Operations**

Mnemonic	Instruction
MADD.fmt	Floating-Point Multiply Add
MSUB.fmt	Floating-Point Multiply Subtract
NMADD.fmt	Floating-Point Negative Multiply Add
NMSUB.fmt	Floating-Point Negative Multiply Subtract

### 14.9.3 Conversion Instructions

These instructions perform conversions between floating-point and fixed-point data types. Each instruction converts values from a number of operand formats to a particular result format. Some conversion instructions use the rounding mode specified in the Floating Control/Status register (*FCSR*), while others specify the rounding mode directly.



Table 14.27 and Table 14.28 list the FPU conversion instructions according to their rounding mode.

**Table 14.27 FPU Conversion Operations Using the FCSR Rounding Mode**

Mnemonic	Instruction
CVT.D fmt	Floating-Point Convert to Double Floating Point
CVT.L fmt	Floating-Point Convert to Long Fixed Point
CVT.S fmt	Floating-Point Convert to Single Floating Point
CVT.W fmt	Floating-Point Convert to Word Fixed Point

**Table 14.28 FPU Conversion Operations Using a Directed Rounding Mode**

Mnemonic	Instruction
CEIL.L fmt	Floating-Point Ceiling to Long Fixed Point
CEIL.W fmt	Floating-Point Ceiling to Word Fixed Point
FLOOR.L fmt	Floating-Point Floor to Long Fixed Point
FLOOR.W fmt	Floating-Point Floor to Word Fixed Point
ROUND.L fmt	Floating-Point Round to Long Fixed Point
ROUND.W fmt	Floating-Point Round to Word Fixed Point
TRUNC.L fmt	Floating-Point Truncate to Long Fixed Point
TRUNC.W fmt	Floating-Point Truncate to Word Fixed Point

#### 14.9.4 Formatted Operand-Value Move Instructions

These instructions move formatted operand values among FPU general registers. A particular operand type must be moved by the instruction that handles that type. There are three kinds of move instructions:

- Unconditional move
- Conditional move that tests an FPU true/false condition code
- Conditional move that tests a CPU general-purpose register against zero

Conditional move instructions operate in a way that might be unexpected. They always force the value in the destination register to become a value of the format specified in the instruction. If the destination register does not contain an operand of the specified format before the conditional move is executed, the contents become undefined. (For more information, see the individual descriptions of the conditional move instructions in the *MIPS32 Architecture Reference Manual, Volume II*.)

Table 14.29 through Table 14.31 list the formatted operand-value move instructions.

**Table 14.29 FPU Formatted Operand Move Instructions**

Mnemonic	Instruction
MOV fmt	Floating-Point Move

**Table 14.30 FPU Conditional Move on True/False Instructions**

Mnemonic	Instruction
MOV.F fmt	Floating-Point Move Conditional on FP False

**Table 14.30 FPU Conditional Move on True/False Instructions(continued)**

Mnemonic	Instruction
MOVT fmt	Floating-Point Move Conditional on FP True

**Table 14.31 FPU Conditional Move on Zero/Non-Zero Instructions**

Mnemonic	Instruction
MOVN.fmt	Floating-Point Move Conditional on Nonzero
MOVZ fmt	Floating-Point Move Conditional on Zero

## 14.9.5 Conditional Branch Instructions

The FPU has PC-relative conditional branch instructions that test condition codes set by FPU compare instructions (C.cond fmt).

All branches have an architectural delay of one instruction. When a branch is taken, the instruction immediately following the branch instruction is said to be in the branch delay slot; it is executed before the branch to the target instruction takes place. Conditional branches come in two versions, depending upon how they handle an instruction in the delay slot when the branch is not taken and execution falls through:

- Branch instructions execute the instruction in the delay slot.
- Branch likely instructions do not execute the instruction in the delay slot if the branch is not taken (they are said to nullify the instruction in the delay slot).

**Although the Branch Likely instructions are included, software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.**

[Table 14.32](#) lists the conditional branch (branch and branch likely) FPU instructions; [Table 14.33](#) lists the deprecated conditional branch likely instructions.

**Table 14.32 FPU Conditional Branch Instructions**

Mnemonic	Instruction
BC1F	Branch on FP False
BC1T	Branch on FP True

**Table 14.33 Deprecated FPU Conditional Branch Likely Instructions**

Mnemonic	Instruction
BC1FL	Branch on FP False Likely
BC1TL	Branch on FP True Likely

## 14.9.6 Miscellaneous Instructions

The MIPS32 architecture defines various miscellaneous instructions that conditionally move one CPU general register to another, based on an FPU condition code.

Table 14.34 lists these conditional move instructions.

**Table 14.34 CPU Conditional Move on FPU True/False Instructions**

<b>Mnemonic</b>	<b>Instruction</b>
MOVN	Move Conditional on FP False
MOVZ	Move Conditional on FP True

## 14.10 Alphabetical Listing of Floating Point Instructions

Table 14.35 shows an alphabetical listing of the floating point unit instruction set, along with the associated instruction group, the page number location of the actual instruction. The actual instruction can be viewed by clicking on either the instruction or the page number reference in the table below. For the definition of each instruction, refer to Table 14.21 through Table 14.34 above.

**Table 14.35 Alphabetical Listing of FPU Instructions**

Instruction Name	Instruction Group
ABS.fmt	Move
ADD.fmt	Arithmetic
BC1F	Conditional Branch
BC1FL	Conditional Branch
BC1T	Conditional Branch
BC1TL	Conditional Branch
C.cond.fmt	Arithmetic
CEIL.L.fmt	Conversion
CEIL.W.fmt	Conversion
CFC1	Move
CTC1	Move
CVT.D.fmt	Conversion
CVT.L.fmt	Conversion
CVT.S.fmt	Conversion
CVT.W.fmt	Conversion
DIV.fmt	Arithmetic
FLOOR.L.fmt	Conversion
FLOOR.W.fmt	Conversion
LDC1	Load/Store
LDXC1	Load/Store
LUXC1	Load/Store
LWC1	Load/Store
LWXC1	Load/Store
MADD.fmt	Multiply-Accumulate
MFC1	Move
MFHC1	Move
MOV.fmt	Move
MOVE.fmt	Move
MOVN.fmt	Move
MOVT.fmt	Move
MOVZ.fmt	Move
MSUB.fmt	Multiply-Accumulate

**Table 14.35 Alphabetical Listing of FPU Instructions(continued)**

<b>Instruction Name</b>	<b>Instruction Group</b>
MTC1	Move
MUL fmt	Arithmetic
NEG.fmt	Move
NMADD fmt	Multiply-Accumulate
NMSUB fmt	Multiply-Accumulate
RECIP.fmt	Arithmetic
ROUND.L fmt	Conversion
ROUND.W fmt	Conversion
RSQRT.fmt	Arithmetic
SDC1	Load/Store
SDXC1	Load/Store
SQRT fmt	Arithmetic
SUB.fmt	Arithmetic
SUXC1	Load/Store
SWC1	Load/Store
SWXC1	Load/Store
TRUNC.L fmt	Conversion
TRUNC.W fmt	Conversion



## MIPS DSP-R2 Application Specific Extension

The interAptiv core includes support for the MIPS DSP ASE Revision 2 that provides enhanced performance capabilities for a wide range of signal-processing applications, with computational support for fractional data types, SIMD, saturation, and other operations that are commonly used in these applications. The DSP instruction set is a collection of special-case instructions, in many cases aimed at the known ‘hot-spots’ of important algorithms that are common in DSP applications. The DSP Revision 2 (DSP-R2) is a superset of DSP Revision 1 and includes all of the instructions in revision 1.

This chapter contains the following sections:

- [Section 15.1 “MIPS32® DSP-R2 ASE Features”](#)
- [Section 15.2 “Common Applications”](#)
- [Section 15.3 “Software Detection of the DSP ASE Revision 2”](#)
- [Section 15.4 “DSP-R2 Registers”](#)
- [Section 15.5 “DSP-R2 Instruction Types”](#)
- [Section 15.6 “DSP-R2 Code Optimization for Performance”](#)
- [Section 15.7 “Compiler Usage Model for MIPS DSP-R2”](#)
- [Section 15.8 “Code Optimizations for the MIPS DSP-R2”](#)
- [Section 15.9 “Programming Examples”](#)
- [Section 15.10 “MIPS32 DSP-R2 Intrinsics”](#)
- [Section 15.11 “DSP-R2 ASE Instruction Groups”](#)
- [Section 15.12 “Listing of DSP-R2 ASE Instruction Groups”](#)
- [Section “Repeat rate is measured as number of independent instructions that can be sent in 1 cycle.”](#)

### 15.1 MIPS32® DSP-R2 ASE Features

The DPS2 ASE contains features that support multiple DSP-R2 functions as described below:

- Q31 and Q15 (signed 16-bit) fractions
- Saturating arithmetic
- Multiplying fractions
- Rounding
- Multiply-accumulate sequences

- Single Instruction Multiple Data (SIMD) operations

### 15.1.1 Q31 and Q15 Signed 16-bit Fractions

DSP-R2 applications use fixed-point fractional data types. Such a fraction is a signed integer that represents an integer divided by some power of two. A 32-bit fractional format where the implicit divisor is  $2^{16}$  (65536) would be referred to as a Q15.16 format; that's because there are 16 bits devoted to fractional precision and 15 bits to the whole number range (the highest bit does duty as a sign bit and isn't counted).

Using this notation, Q31.0 is a conventional signed integer, and Q0.31 is a fraction representing numbers between -1 and 1. The Q0.31 notation is the most popular 32-bit format for DSP-R2 applications since it won't overflow when multiplied. Q0.31 is often abbreviated to Q31.

### 15.1.2 Saturating Arithmetic

The DSP-R2 instruction set in the interAptiv core provides instruction that perform both saturated and non-saturated arithmetic. When a calculation overflows, the saturating (**\_SA**) instructions make the result the most positive or most negative representable value.

### 15.1.3 Multiplying Fractions

Multiplying two Q31 fractions by re-using a full-precision integer multiplier results in a 64-bit result that consists of a Q62 result with (in the very highest bit) a second copy of the sign bit. A left-shift-by-1 must then be performed on this value to produce a Q63 format. Similarly, Q15 multiplies that generate a Q31 value must also perform a shift-left. This is the function of the **MULQ** instructions.

### 15.1.4 Rounding

Some fractional operations implicitly discard the least-significant bits. To get a better approximation, increment the truncated result by one when the discarded bits represent more than a half of the value of a 1 in the new LS position. That is how the term '*rounding*' is defined in this chapter.

### 15.1.5 Multiply-Accumulate Sequences

For enhanced performance performing fractional and saturating operations, the interAptiv core contains four accumulators for multiply-accumulate sequences (with fixed-point types, sometimes saturating). For backward compatibility with previous generation MIPS processors that have only one accumulator, the new *ac0* accumulator functions as the previous generation *H//LO*.

### 15.1.6 SIMD Operations

Many DSP-R2 calculations are a good match for Single Instruction Multiple Data (SIMD) or *vector* operations, where the same arithmetic operation is applied in parallel to several sets of operands.

In the MIPS DSP-R2 ASE, some operations are SIMD type - two 16-bit operations or four 8-bit operations are carried out in parallel on operands packed into a single 32-bit general-purpose register. Instructions operating on vectors can be recognized because the name includes **.PH** (paired-half, usually signed, often fractional) or **.QB** (quad-byte, always unsigned, only occasionally fractional).



## 15.2 Common Applications

Different target applications generally need different data size and precision. The DSP-R2 ASE can be used by the following applications.

- *32-bit data*: audio (non-hand-held) decoding/encoding - a wide range of “hi-fi” standards for consumer audio or television sound.
- *16-bit data*: digital voice for telephony. International telephony code/decode standards include G.723.1 (8Ksample/s, 5-6Kbit/s data rate, 37ms delay), G.729 (8Kbit/s, 15ms delay) and G.726 (16-40Kbit/s, computationally simpler and higher quality, good for carrying analogue modem tones). Application-specific filters are used for echo cancellation, noise cancellation, and channel equalization. Also used for soft modems and general ‘DSP’ work such as filters, correlation, and convolution.
- *8-bit data*: processing of printer images, JPEG (still) images and video data.

## 15.3 Software Detection of the DSP ASE Revision 2

The presence of the MIPS DSP-R2 ASE in the interAptiv core is indicated by two static bits in the *Config3* register: the ‘DSP Present’ bit (*Config3\_DSPP*) indicates the presence of the DSP-R2 ASE, and the ‘DSP Rev2 Present’ bit (*Config3\_DSP2P*) indicates the presence of the MIPS DSP ASE Rev2. Because all members of the interAptiv family support both ASEs, these bits are always set to 1.

The CP0 Status register (*Status\_MX*) must be set to enable access to the extra instructions defined by the DSP-R2 module, as well as to the **MTLO/HI**, **MFLO/HI** instructions that access accumulators ac1, ac2, and ac3. Executing a DSP-R2 ASE instruction or the **MTLO/HI**, **MFLO/HI** instructions with this bit set to zero causes a DSP-R2 State Disabled Exception (exception code 26 in the CP0 *Cause* register). This exception can be used by system software to do lazy context-switching.

## 15.4 DSP-R2 Registers

The DSP-R2 ASE defines three additional accumulator registers and one additional control/status register, as described below. These registers require the operating system to recognize the presence of the DSP-R2 ASE and to include these additional registers in the context save and restore operations.

### 15.4.1 DSP-R2 Accumulator Registers


Whereas a standard MIPS32 architecture CPU has just one 64-bit multiply unit accumulator (accessible as *hi/lo*), the DSP-R2 ASE in the interAptiv core provides four 64-bit accumulators. Instructions accessing the extra accumulators specify a 2-bit field as 0 - 3 (0 selects the original accumulator).

The DSP-R2 ASE includes three HI/LO accumulator register pairs (ac1, ac2, and ac3) in addition to the HI/LO register pair (ac0) in the standard MIPS32 architecture. These registers improve the parallelization of independent accumulation routines—for example, filter operations, convolutions, etc. DSP-R2 instructions that target the accumulators use two instruction bits to specify the destination accumulator, with the zero value referring to the original accumulator.

### 15.4.2 DSP-R2 ASE Control Register

This is a part of the user-mode programming model for the DSP-R2 ASE, and is a 32-bit value read and written with the **RDDSP/WRDSP** instructions. It holds state information for some DSP-R2 sequences.

**Figure 15.1 DSP-R2 Control Register**



**Table 15.1 Field Descriptions for DSP-R2 Control Register**

Name	Bit(s)	Description	Read/Write	Reset State
0	31:28	Reserved. Write as zero. Undefined on read.	R/W	Undefined
<i>CCOND</i>	27:24	Condition bits set by compare instructions (there have to be four to report on compares between vector types). "Compare" operations on scalars or vectors of length two only touch the lower-numbered bits.	R/W	0
<i>OUFLAG</i>	23:16	Overflow/underflow flags.  One of these bits may be set when a result overflows (whether or not the result is saturated depends on the instruction - the flag is set in either case). Underflow indicates a value that is negative but with excessive absolute value. Any overflowed/underflowed result produced by any DSP-R2 ASE instruction sets an <i>ouflag</i> bit, <i>except</i> for <b>ADDSC/ADDWC</b> and <b>SHILO/SHILOV</b> instructions.  See <a href="#">Table 15.2</a> for a full list of which bits are set by what instructions.	R/W	
0	15	Reserved. Write as zero. Undefined on read.	R/W	
<i>EFI</i>	14	Extract field of instruction.  This bit is set by any of the accumulator-to-register bit field extract instructions ( <b>EXTP</b> , <b>EXTPV</b> , <b>EXTPDP</b> , or <b>EXTPDP</b> ) only if the instruction finds there are insufficient bits to extract. In other words, if <i>DSPControl</i> <sub>POS</sub> , which is supposed to mark the highest-numbered bit of the field being extracting, is less than the size value specified by the instruction.  Note that this bit is not sticky, so each invocation of one of the four instructions will reset the bit depending on whether or not the instruction failed.	R/W	
<i>C</i>	13	Carry bit for 32-bit add/carry instructions <b>ADDSC</b> and <b>ADDWC</b> .  This bit is set and used by special add instructions that implement a 64-bit add across two GPRs. The <b>ADDSC</b> instruction sets the bit and the <b>ADDWC</b> instruction uses this bit.	R/W	
<i>SCOUNT</i>	12:7	Size count.  This field specifies the size of the bit field to be inserted, while <i>DSPControl</i> <sub>POS</sub> specifies the insert position.  This field is used by "variable" bit field insert and extract instructions such as <b>INSV</b> (the normal MIPS32 <b>INS/EXT</b> instructions have the field size and position hard-coded in the instruction).	R/W	
0	6	Reserved. Write as zero. Undefined on read.	R/W	

**Table 15.1 Field Descriptions for DSP-R2 Control Register**

Name	Bit(s)	Description	Read/Write	Reset State
<i>POS</i>	5:0	<p>Insertion position.</p> <p>This field is used by the variable insert instruction <b>INSV</b> and specifies the position of the bit field to be inserted, while <i>DSPControl</i><b>SCOUNT</b> specifies the size of the bit field to be inserted. This field to specify the insert position.</p> <p>In most insertions (following the lead of the standard MIPS32 insert/extract instructions) this field is set to the lowest bit number. However, in the DSP-R2 ASE extract-from-accumulator instructions (<b>EXTP</b>, <b>EXTPV</b>, <b>EXTPDP</b> and <b>EXTPDPV</b>), this field identifies the <i>highest</i>-numbered bit in the field.</p> <p>The <b>EXTPDP</b> and <b>EXTPDPV</b> instructions post-decrement this field (by the bit-field length <i>Size</i>) to help software which is unpacking a series of bit fields from a dense data structure.</p> <p>The <b>MTHLIP</b> instruction increments the value in this field by 32 after copying the value of <i>lo</i> to <i>hi</i>.</p>	R/W	

The bits of the overflow flag *OUFLAG* field in the *DSPControl* register are set by a number of instructions, as described in Table 15.2. These bits are sticky and can be reset only by an explicit write to these bits in the register (using the **WRDSP** instruction). Refer to the following section for more information on the DSP-R2 instructions.

**Table 15.2 Instructions that set the DSPControl OUFLAG Bits**

Bit Number	Description
16	This bit is set when the destination is accumulator ( <i>HI-LO</i> pair) zero, and an operation overflow or underflow occurs. These instructions are: <b>DPAQ_S</b> , <b>DPAQ_SA</b> , <b>DPSQ_S</b> , <b>DPSQ_SA</b> , <b>DPAQX_S</b> , <b>DPAQX_SA</b> , <b>DPSQX_S</b> , <b>DPSQX_SA</b> , <b>MAQ_S</b> , <b>MAQ_SA</b> and <b>MULSAQ_S</b> .
17	Same instructions as above, when the destination is accumulator ( <i>HI-LO</i> pair) one.
18	Same instructions as above, when the destination is accumulator ( <i>HI-LO</i> pair) two.
19	Same instructions as above, when the destination is accumulator ( <i>HI-LO</i> pair) three.
20	Instructions that set this bit on an overflow/underflow: <b>ABSQ_S</b> , <b>ADDQ</b> , <b>ADDQ_S</b> , <b>ADDU</b> , <b>ADDU_S</b> , <b>ADDWC</b> , <b>SUBQ</b> , <b>SUBQ_S</b> , <b>SUBU</b> and <b>SUBU_S</b> .
21	Instructions that set this bit on an overflow/underflow: <b>MUL</b> , <b>MUL_S</b> , <b>MULEQ_S</b> , <b>MULEU_S</b> , <b>MULQ_RS</b> , and <b>MULQ_S</b> .
22	Instructions that set this bit on an overflow/underflow: <b>PRECRQ_RS</b> , <b>SHLL</b> , <b>SHLL_S</b> , <b>SHLLV</b> , and <b>SHLLV_S</b> .
23	Instructions that set this bit on an overflow/underflow: <b>EXTR</b> , <b>EXTR_S</b> , <b>EXTR_RS</b> , <b>EXTRV</b> , and <b>EXTRV_RS</b> .

## 15.5 DSP-R2 Instruction Types

The DSP-R2 instruction set in the interAptiv core is a collection of special-case instructions aimed at addressing known ‘hot-spots’ of important DSP algorithms.

In this section, the DSP-R2 instructions have been divided into the following subsections, which represent their likely usage and application and type of the result derived. Most of the multiplication instructions have multiple uses and are divided based on the most obvious use.

- [Arithmetic - 64-bit](#)
- [Arithmetic - saturating and/or SIMD Types](#)
- [Bit-shifts - saturating and/or SIMD types](#)
- [Comparison and "conditional-move" operations on SIMD types](#) - includes **PICK** instructions.
- [Conversions to and from SIMD types](#)
- [Multiplication - SIMD types with result in GP register](#)
- [Multiply Q15s from paired-half and accumulate](#)
- [Load with register+register address](#)
- [DSPControl register access](#)
- [Accumulator access instructions](#)
- [Dot products and building blocks for complex multiplication](#) - includes full-word (Q31) multiply-accumulate
- [Other DSP ASE instructions](#) - everything else...

A complete alphabetical list of DSP-R2 instructions is shown in [Section 15.11 “DSP-R2 ASE Instruction Groups”](#), followed by a description of each individual instruction.

### 15.5.1 Hints in instruction names

An instruction’s name may have some suffixes which are often informative:

**Q**: generally means it treats operands as fractions (which isn’t important for adds and subtracts, but is important for multiplications and convert operations);

**\_S**: usually means the full-precision result is saturated to the size of the destination; **\_SA** is used for instructions which saturate intermediate results before accumulating; and **R**: denotes rounding (see above);

**.W**, **.PH**, **.QB**: suggest the operation is dealing with 32-bit word, paired-halfword, or quad-byte values respectively. Where there are two of these (as in **MAQ\_S.W.PHL**) the first one suggests the type of the result, and the second the type of the operand(s).

**V**: (in a shift instruction) suggests that the shift amount is defined in a register, rather than being encoded in a field of the instruction.

### 15.5.2 Arithmetic — 64-bit

**ADDSC/ADDWC** generate and use a carry bit, for efficient 64-bit add.

### 15.5.3 Arithmetic — Saturating and/or SIMD Types

- *32-bit signed saturating arithmetic*: **ADDQ\_S.W**, **SUBQ\_S.W** and **ABSQ\_S.W**.
- *Paired-half and quad-byte SIMD arithmetic*: perform the same operation simultaneously on both 16-bit halves or all four 8-bit bytes of a 32-bit register. The “Q” in the instruction mnemonic for the PH operations here is cosmetic: Q15 and signed 16-bit integer add/subtract operations are bit-identical - Q15 only behaves very differently when converted or multiplied.

The paired half operations are: **ADDQ.PH/ADDQ\_S.PH**, **SUBQ.PH/SUBQ\_S.PH** and **ABSQ\_S.PH**.

The quad-byte operations (all unsigned) are: **ADDU.QB/ADDU\_S.QB**, **SUBU.QB/SUBU\_S.QB**.

- *Sum of quad-byte vector*: **RADDU.W.QB** does an unsigned sum of the four bytes found in a register, zero extends the result and delivers it as a 32-bit value.

### 15.5.4 Bit-shifts — Saturating and/or SIMD Types

All shifts can either have a shift amount encoded in the instruction, or - indicated by a trailing “V” in the instruction name - provided as a register operand. **PH** and 32-bit shifts have optional forms which saturate the result.

- *32-bit signed shifts*: include a saturating version of shift left, **SHLL\_S.W**; and an auto-rounded shift right (just the “arithmetic”, sign-propagating form): **SHRA\_R.W**. Recall from above that rounding can be imagined as pre-adding a half to the least significant surviving bit.
- *Paired-half and quad-byte SIMD shifts*: **SHLL.PH/SHLLV.PH/SHLL\_S.PH/SHLLV\_S** are as above. For **PH** only there’s a shift-right-arithmetic instruction (“arithmetic” means it propagates the sign bit downward) **SHRA.PH**, which has a variant which rounds the result **SHRA\_R.PH**.

The quad-byte shifts are unsigned and don’t round or saturate: **SHLL.QB/SHLLV.QB**, **SHRL.QB/SHRLV.QB**.

### 15.5.5 Comparison and “Conditional-Move” Operations on SIMD Types

The “**cmp**” operations simultaneously compare and set flags for two or four values packed in a vector (with equality, less-than and less-than-or-equal tests). For PH that’s **CMP.EQ.PH**, **CMP.LT.PH** and **CMP.LE.PH**. The result is left in the two LS bits of *DSPControl[ccond]*.

For quad-byte values **CMPU.EQ.QB**, **CMPU.LT.QB** and **CMPU.LE.QB** simultaneously compare and set flags for four bytes in *DSPControl[ccond]* - the flag relating to the bytes found in the low-order bits of the source register is in the lowest-numbered bit (and so on). There’s an alternative set of instructions **CMPGU.EQ.QB**, **CMPGU.LT.QB** and **CMPGU.LE.QB** which leave the 4-bit result in a specified general-purpose register.

**PICK.PH** uses the two LS bits of *DSPControl[ccond]* (usually the outcome of a paired-half compare instruction, see above) to determine whether corresponding halves of the result should come from the first or second source register. Among other things, this can implement a paired-half conditional move. You can reverse the order of your conditional inputs to do a move dependent on the complementary condition, too.

**PICK.QB** does the same for QB types, this time using four bits of *DSPControl[ccond]*.

## 15.5.6 Conversions to and from SIMD Types

Conversion operations from larger to smaller fractional types have names which start “**PRECRQ...**” for “precision reduction, fractional”. Conversion operations from smaller to larger have names which start “**PRECE...**” for “precision expansion”.

- *Form vector from high/low parts of two other paired-half values:* **PACKRL.PH** makes a paired-half vector from two half vectors, swapping the position of each sub-vector. It can be used to acquire a properly formed sub-vector from a non-aligned data stream.
- *One Q15 from a paired-half to a Q31 value:* **PRECEQ.W.PHL/PRECEQ.W.PHR** select respectively the “left” (high bit numbered) or “right” (low bit numbered) Q15 value from a paired-half register, and load it into the result register as a Q31 (that is, it’s put in the high 16 bits and the low 15 bits are zeroed).
- *Two bytes from a quad-byte to paired-half:* **PRECEQU.PH.QBL/PRECEQU.PH.QBR** picks two bytes from either the “left” (high bit numbered) or “right” (low bit numbered) halves of a quad-byte value, and unpacks to a pair of Q15 fractions.

**PRECEQU.PH.QBLA** does the same, except that it picks two “alternate” bytes from bits 31-24 and 15-8, while **PRECEQU.PH.QBRA** picks bytes from bits 23-16 and 7-0.

Similar instructions without the **q** - **PRECEU.PH.QBL**, **PRECEU.PH.QBR**, **PRECEU.PH.QBLA**” and **PRECEU.PH.QBRA** - work on the same register fields, but treat the quantities as integers, so the 16-bit results get their low bits set.

- *2×Q31 to a paired-half:* both operands and result are assumed to be signed fractions, so **PRECRQ.PH.W** just takes the high halves of the two source operands and packs them into a paired-half; **PRECRQ\_RS.PH.W** rounds and saturates the results to Q15.
- *2×paired-half to quad-byte:* you need two source registers to provide four paired-half values, of course. This is a fractional operation, so it’s the low bits of the 16-bit fractions which are discarded.

**PRECRQ.QB.PH** treats the paired-half operands as unsigned fractions, retaining just the 8 high bits of each 16-bit component.

**PRECRQU\_S.QB.PH** treats the paired-half operands as Q15 signed fractions and both rounds and saturates the result (in particular, a negative Q15 fraction produces a zero byte, since zero is the lowest representable quantity).

- *Replicate immediate or register value to paired-half:* in **REPL.PH** the value to be replicated is a 10-bit signed immediate value (that’s in the range  $-512 \leq x \leq 511$ ) which is sign-extended to 16 bits, whereas in **REPLV.PH** the value - assumed to be already a Q15 value - is in a register.
- *Replicate single value to quad-byte:* there’s both a register-to-register form **REPLV.QB** and an immediate form **REPL.QB**.

## 15.5.7 Multiplication - SIMD Types with Result in GP Register

When a multiply’s destination is a general-purpose register, the operation is still done in the multiply unit, and you should expect it to overwrite the *hi/lo* registers (otherwise known as *ac0*.)

- *8-bit×16-bit 2-way SIMD multiplication:* **MULEU\_S.PH.QBL/MULEU\_S.PH.QBR** picks the “left” (high bit numbered) or “right” (low bit numbered) pair of byte values from one source register and a pair of 16-bit values from the other. Two unsigned integer multiplications are done at once, the results unsigned-saturated and delivered to the two 16-bit halves of the destination.

The asymmetric use of the source operands is not a bit like a Q15 operation. But 8×16 multiplies are heavily used in imaging and video processing (JPEG image encode/decode, for example).

- *Paired-half SIMD multiplication*: **MULQ\_RS.PH** multiplies two Q15s at once and delivers it to a paired-half value in a general-purpose register, with rounding and saturation.
- *Multiply half-PH operands to a Q31 result*: **MULEQ\_S.W.PHL/MULEQ\_S.W.PHR** pick the “left”/“right” Q15 value respectively from each operand, multiply and store a Q31 value.

“Precision-doubling” multiplications like this *can* overflow, but only in the extreme case where you multiply -1×-1, and can’t represent 1 exactly.

### 15.5.8 Multiply Q15s from Paired-Half and Accumulate

**MAQ\_S.W.PHL/MAQ\_S.W.PHR** picks either the left/high or right/low Q15 value from each operand, multiplies them to Q31 and accumulates to a Q32.31 result. The multiply is saturated only when it’s -1×-1.

**MAQ\_SA.W.PHL/MAQ\_SA.W.PHR** differ in that the final result is saturated to a Q31 value held in the low half of the accumulator (required by some ITU voice encoding standards).

### 15.5.9 Load with Register + Register Address

Previously available only for floating point data<sup>1</sup>: **LWX** for 32-bit loads, **LHX** for 16-bit loads (sign-extended) and **LBUX** for 8-bit loads, zero-extended.

### 15.5.10 DSP-R2 Control Register Access

**WRDSP RS, MASK** sets *DSPControl* fields, but only those fields which are enabled by a 1 bit in the 6-bit mask.

**RDDSP** reads *DSPControl* into a GPR; but again it takes a mask field. Bit fields in the GPR corresponding to *DSPControl* fields which are not enabled will be set all-zero.

The mask bits tie up with fields like this:

**Table 15.3 Mask bits for instructions accessing the DSPControl register**

<i>Mask Bit</i>	<i>DSPControl field</i>
0	pos
1	scount
2	c
3	ouflag
4	ccond
5	EFI

### 15.5.11 Accumulator Access Instructions

- *Historical instructions which now access new accumulators*: the familiar **MFHI/MFLO/MTHI/MTLO** instructions now take an optional extra accumulator-number parameter.

1. Well, an integer instruction is also included in the MIPS SmartMIPS™ ASE.

- *Shift and move to general register:* **EXTR.W/EXTR\_R.W/EXTR\_RS.W** gets a 32-bit field from an accumulator (starting at bit 0 up to 31) and puts the value in a general purpose register. At your option you can specify rounding and signed 32-bit saturation.

**EXTRV.W/EXTRV\_R.W/EXTRV\_RS.W** do the same but specify the field's starting bit number with a register.

- *Extract bit field from accumulator:* **EXTP/EXTPV** takes a bit field (up to 32 bits) from an accumulator and moves it to a GPR. The length of the field can be an immediate value or from a register. The position of the field is determined by *DSPControl[pos]*, which holds the bit number of the most significant bit.

**EXTPDP/EXTPDPV** do the same, but also auto-decrement *DSPControl[pos]* to the bit-number just below the field you extracted.

- *Accumulator rearrangement:* **SHILO/SHILOV** has a *signed* shift value between -32 and +31, where positive numbers shift right, and negative ones shift left. The “v” version, as usual, takes the shift value from a register. The right shift is a “logical” type so the result is zero extended.
- *Fill accumulator pushing low half to high:* **MTHLIP** moves the low half of the accumulator to the high half, then writes the GPR value in the low half. Generally used to bring 32 more bits from a bit stream into the accumulator for parsing by the various **EXT...** instructions.

## 15.5.12 Dot Products and Building Blocks for Complex Multiplication

In 2-dimensional vector math (or in any doubled-up step of a multiply-accumulate sequence which has been optimized for 2-way SIMD) you're often interested in the *dot product* of two vectors:

$$v[0] * w[0] + v[1] * w[1]$$

In many cases you take the dot product of a series of vectors and add it up, too.

Some algorithms use complex numbers, represented by 2D vectors. Complex numbers use *i* to stand for “the square root of -1”, and a vector  $[a, b]$  is interpreted as  $a + ib$  (mathematicians leave out the multiply sign and use single-letter variables, habits which would not be appreciated in C programming!) Complex multiplication just follows the rules of multiplying out sums, remembering that  $i * i = -1$ , so:

$$(a + ib) * (c + id) = (a*c - b*d) + i(a*d + b*c)$$

Or in vector format:

$$[a, b] * [c, d] = [a*c - b*d, a*d + b*c]$$

The first element of the result (the “real component”) is like a dot product but with a subtraction, and the second (the “imaginary component”) is like a dot product but with the vectors crossed.

- *Q15 dot product from paired-half, and accumulate:* **DPAQ\_S.W.PH** does a SIMD multiply of the Q15 halves of the operands, then adds the results and saturates to form a Q31 fraction, which is accumulated into a Q32.31 fraction in the accumulator.

**DPSQ\_S.W.PH** does the same but subtracts the dot product from the accumulator.

For the imaginary component of a complex multiply, first swap the Q15 numbers in one of the register operands with a **ROT** (bit-rotate) instruction.



For the real component of a complex Q15 multiply, you have the difference-of-products instruction **MULSAQ\_S.W.PH**, which parallel-multiplies both Q15 halves of the PH operands, then computes the difference of the two results and leaves it in an accumulator in Q32.31 format (beware: this does not accumulate the result).

- *16-bit integer dot-product from paired-half, and accumulate:* **DPAU.H.QBL/DPAU.H.QBR** picks two QB values from each source register, parallel-multiplies the corresponding pairs to integer 16-bit values, adds them together and then adds the whole lot into an accumulator. **DPSU.H.QBL/DPSU.H.QBR** do the same sum-of-products, but the result is then subtracted from the accumulator. In both cases, note this is integer (not fractional) arithmetic.
- *Q31 saturated multiply-accumulate:* is the nearest thing you can get to a dot-product for Q31 values. **DPAQ\_SA.L.W** does a Q31 multiplication and saturates to produce a Q63 result, which is added to the accumulator and saturated again. **DPSQ\_SA.L.W** does the same, except that the multiply result is subtracted from the accumulator (again, useful for the real component of a complex number).

### 15.5.13 Other DSP-R2 ASE Instructions

- *Branch on DSPControl field:* **BPOSGE32** branches if  $DSPControl[pos] \geq 32$ .

Typically the test is for “is it time to load another 32 bits of data from the bit stream yet?”

- *Circular buffer index update:* **MODSUB** takes an operand which packs both a maximum index value and an index step, and uses it to decrement a “buffer index” by the step value, but arranging to step from zero to the provided maximum.
- *Bit field insert with variable size/position:* **INSV** is a bit-insert instruction. It acts like the MIPS32 standard instruction **INS** except that the position and size of the inserted field are specified not as immediates inside the instruction, but are obtained from  $DSPControl[pos]$  (which should be set to the lowest numbered bit of the field you want) and  $DSPControl[scount]$  respectively.
- *Bit-order reversal:* **BITREV** reverses the bits in the low 16 bits of the register. The high half of the destination is zero.

The bit-reverse operation is a computationally crucial step in buffer management for FFT algorithms, and a 16-bit operation supports up to a 32K-point FFT, which is much more than enough. A full 32-bit reversal would be expensive and slow.

## 15.6 DSP-R2 Code Optimization for Performance

Some optimization methods are seldom used for general-purpose software, but are often seen in DSP-R2 code. A typical example is a technique called zipping, which reduces the number of data loads in algorithms like FIR filters. Consider the calculation of the first two output values of an 8-tap FIR filter. The illustration below shows how the coefficients get multiplied by the data samples:

Input data samples:	d0	d1	d2	d3	d4	d5	d6	d7	d8	...
Coefficients for y0:	c0	c1	c2	c3	c4	c5	c6	c7		
Coefficients for y1:		c0	c1	c2	c3	c4	c5	c6	c7	
First output (y0):	$c0d0 + c1d1 + c2d2 + c3d3 + c4d4 + c5d5 + c6d6 + c7d7$									
Second output (y1):	$c0d1 + c1d2 + c2d3 + c3d4 + c4d5 + c5d6 + c6d7 + c7d8$									

A very naive implementation will load each coefficient and data sample from memory every time they are needed. A more optimized implementation will load the coefficients just once and keep them in registers. It will load data sam-

ples d0-d7 first, to compute the first output, and then data samples d1-d8 to compute the second output. With zipping, an even more optimized implementation will load d0-d8 once and use the loaded values for both output calculations. The relatively large number of general-purpose registers in the MIPS architecture is useful for applying this technique. An even larger number of samples can be kept in registers if the SIMD features of the DSP-R2 ASE are used. In this case, another slightly rearranged set of coefficients may be needed, as illustrated below for the case of 16-bit coefficients and samples packed into 32-bit words:

```

Input data samples:      d0:d1  d2:d3  d4:d5  d6:d7  d8:d9
Coefficients for y0:    c0:c1  c2:c3  c4:c5  c6:c7
Coefficients for y1:    00:c0  c1:c2  c3:c4  c5:c6  c7:00

```

The first set of coefficients is used to compute y0 and all even-numbered output samples. The second set is used for y1 and all odd-numbered output samples.

Because of the large number of general-purpose registers, especially considering the SIMD features offered by the DSP-R2 ASE, the interAptiv cores lend themselves well to algorithms processing more data elements at a time. For example, it is easy to meet the requirements of a radix-4 FFT implementation, which is faster than a radix-2 implementation but needs to keep a large number of values in registers during the calculation.

Many algorithms can be implemented in a variety of different ways and often some of these algorithm transformations offer performance advantages. The output results are similar in all cases, but an optimal implementation strikes the best balance between number of registers needed, number of memory operations, number and type of arithmetic operations, regularity of data access patterns, result delays, etc. Make sure the selected algorithm implementation approach is the best match to the architecture.

It should be noted that a reasonable degree of familiarity with the DSP-R2 instructions will allow the programmer to extract the best performance. The architecture specification and the core programming guide provide the necessary information.

If the data type is 16-bit or 8-bit, then attempting to rewrite the algorithm in SIMD style using operations directly supported by the DSP-R2 ASE instruction set will yield a lower cycle count due to the obtained parallelism. Once the number of instructions required to implement the algorithm is minimized, the instructions must be scheduled taking into account their result delays. When evaluating the resulting performance, obtaining a trace from one execution will illuminate the cause of stalls and also facilitate optimization.

### 15.6.1 Pixel Unpacking Example

As a simple illustration of efficiently using the SIMD capabilities of the processor, consider the task of unpacking YUV image data prior to processing. The YUV color space is commonly used in image and video processing. The color components (U and V) are subsampled with respect to the luminosity (Y) by a factor of two or four in the horizontal and/or vertical dimension. A commonly used format is YUV 4:2:2, which has the color components subsampled horizontally by a factor of two. Hence there is one luminosity value for each pixel, but a UV pair is shared between two adjacent pixels. Video data in YUV 4:2:2 format is commonly transmitted in the following order:

```

Pixel data:      U0  Y0  V0  Y1  U2  Y2  V2  Y3  U4  Y4  V4  Y5  ...

```

Video processing algorithms usually perform different tasks on each of the Y, U, and V components. In order to use the SIMD capabilities of the interAptiv core, the pixel data needs to be unpacked, i.e., each of the Y, U, and V values should be separated out and grouped together. Examining the DSP-R2 ASE instruction set reveals that some instructions intended for precision reduction and expansion can be used to implement data unpacking:

```

lw          $t0, 0($a0)      # U0:Y0:V0:Y1 - two pixels in YUV 4:2:2
lw          $t1, 4($a0)      # U2:Y2:V2:Y3 - next two pixels

precequ.ph.qbra $t2, $t0      # 00:Y0:00:Y1 - half the unpacked Ys
precrq.qb.ph    $t4, $t0, $t1 # U0:V0:U2:V2 - interleaved Us and Vs
precequ.ph.qbra $t3, $t1      # 00:Y2:00:Y3 - the other unpacked Ys
precequ.ph.qbra $t5, $t4      # 00:V0:00:V2 - unpacked Vs
precequ.ph.qbla $t5, $t4      # 00:U0:00:U2 - unpacked Us

```

Unpacking the YUV data as illustrated above also has the advantage of converting each data item from 8-bit unsigned to 16-bit unsigned format. This ensures there is enough room for performing the video processing calculations with enhanced precision.

Note that the example above as well as the following examples have not been scheduled to the interAptiv core pipelines. In a real application other instructions surrounding the illustrated code fragments can be used to fill-in the delays caused by result dependencies.

### 15.6.2 Sum of Absolute Differences Example

Another interesting DSP-R2 ASE example shows the kernel performing the sum of absolute differences (SAD) function used in motion estimation algorithms for video compression. The function accumulates the absolute difference between the pixels from a reference 8x8-pixel block and those from a similar block inside the current video frame. Using the DSP-R2 ASE, here is how the SAD of four pixels at a time can be calculated and accumulated:

```

SUBUH_R.QB    $t0, $s0, $s1    # subtract 4 pixels with halving
ABSQ_S.QB     $t0, $t0         # find the 4 absolute values
RADDU.W.QB    $t0, $t0         # sum the absolute values
ADDU          $v0, $v0, $t0     # accumulate the result

```

The above sequence is four instructions long and will execute in four cycles when properly scheduled. Performing the same calculation using the MIPS32 Release 2 instruction set will require approximately 20 instructions. The performance advantages offered by the DSP-R2 ASE are obvious.

### 15.6.3 Bit Stream Unpacking Example

The last example of efficient use of DSP-R2 ASE instructions presented here is bit stream unpacking. Many audio and video compression algorithms pack various parameters of different bit widths in a continuous compressed bit-stream. The decoder has to first unpack-or extract-the individual values from the bit stream before further processing them. Recognizing the importance of this task, the DSP-R2 ASE instruction set includes instructions that facilitate and accelerate bit stream unpacking. One of the accumulator registers is used as a 64-bit data buffer. The EXTP instruction variants extract a specified number of bits and optionally decrement the pos field of the DSPControl register. The pos field holds the number of remaining bits in the bit buffer. The BPOSGE32 instructions checks this number and branches if there are at least 32 bits left. And finally, the MTHLIP instruction is used to reload the bit buffer with a new 32-bit word and at the same time increment the number of available bits by 32. This process is illustrated in the code fragment below:

```

loop:
lbu          $t0, 0($a0)      # size of the data field to extract
extpdpv     $v0, ac0, $t0    # extract a value from ac0, pos -= size
addiu       $a1, $a1, 4      # increment output data pointer
addiu       $a0, $a0, 1      # increment size table pointer

```

```

bposge32  loop           # loop back if pos >= 32
sw        $v0, -4($a1)   # store the extracted value

lw        $t1, 0($a2)    # load next bitstream word
addiu    $a2, $a2, 4     # increment bitstream pointer
bne      $a2, $a3, loop  # loop until no more data available
mthlip   $t1, ac0       # update ac0 bit buffer, pos += 32

```

The illustrated code loads the size (bit width) of each field to be extracted from memory. The performance of the loop can be further improved if the sizes are static and known in advance.

## 15.7 Compiler Usage Model for MIPS DSP-R2

The MIPS® DSP-R2 ASE allows for performance optimization of signal processing and multimedia applications running on the interAptiv core. The DSP-R2 ASE includes a set of instructions that provide computational support for fractional data types, SIMD, saturation, and other operations commonly used in DSP applications. This section describes a compiler usage model for the MIPS DSP-R2.

Typical DSP applications include certain loops or kernels that take a large percentage of the execution time. Because of the sensitivity of DSP application performance to these kernels, programmers have tended to write these functions in assembly, hand-scheduling the code for optimal pipeline scheduling. But with the introduction of MIPS DSP-R2 extensions, compilers like GCC can be used to reduce, and perhaps eliminate, the need to write assembly code. However, to obtain the best optimizations, a particular coding style and usage must be followed, as explained in this section.

For reference, the GCC Compiler 6.03.00-rc3, based on FSF (Free Software Foundation) GCC 3.4, was used for the examples in this section. Programmers should use the newest compilers available to ensure that the DSP-R2 ASE is supported and to enable the highest performance instruction scheduling. Throughout the section, the term “GCC compiler” refers to a compiler that incorporates the DSP-R2 ASE intrinsics, such as the GCC compiler.

Using a high-level language such as C to develop applications has many advantages:

1. C programmers do not have to manually allocate registers—the compiler can allocate registers for variables.
2. C programmers do not have to manually schedule instructions—the compiler can schedule instructions based on given latency information for specific CPU pipelines.
3. C programmers do not have to consider function calling conventions when writing modular programs—the compiler saves and restores registers at entries and exits of functions per a pre-defined convention.
4. C programmers can declare SIMD variables and use generic C operators or intrinsics (built-in functions) to manage SIMD variables in order to achieve parallelism.
5. Development and debug time is shortened when using a high-level language.
6. Applications are more maintainable when written in a high-level language.

### 15.7.1 Enabling the DSP-R2 ASE in the Compiler

To enable the MIPS32 DSP-R2 ASE in the GCC compiler, the “-mdsp2” command-line option is required. In addition to “-mdsp2”, “-mips32r2” is recommended for better performance in accessing elements in SIMD variables, because this allows the use of the MIPS32 Release 2 instruction **INS** for more efficient code.

## 15.7.2 Data Types

In the GCC compiler, the Q15 data type is represented by 16-bit integer data type (short), and the Q31 data type is represented by 32-bit integer data type (int). Typedefs can be created for Q15 and Q31 as follows:

```
typedef short q15;
typedef int q31;
```

The MIPS32 DSP-R2 ASE implements four 64-bit accumulators which can be represented by a “long long” data type.

```
typedef long long a64;
```

To declare SIMD data types, typedefs with special vector\_size attributes are required. For example,

```
typedef signed char v4i8 __attribute__((vector_size(4)));
typedef short v2q15 __attribute__((vector_size(4)));
```

where “v4i8” defines a SIMD data type containing four 8-bit integer data. “v2q15” defines a type containing two Q15 fractional data (which is the same as two 16-bit integer data).

## 15.7.3 Initialization of Q15 and Q31 Variables

To initialize Q15 variables, programmers can multiply the fractional value (e.g., 0.1234) by 32768.0. To initialize Q31 variables, programmers can multiply the fractional value by 2147483648.0.

```
Ex: /* Q15 Example */
typedef short q15;
q15 a = 0.1234 * 32768.0;
/* ----- */

Ex: /* Q31 Example */typedef int q31;
q31 b = 0.2468 * 2147483648.0;
```

## 15.7.4 Initialization of SIMD Variables


To initialize SIMD variables is similar to initializing aggregate data. The following examples show how to initialize SIMD variables.

```
Ex: /* v4i8 Example */
v4i8 a = {1, 2, 3, 4};
v4i8 b;
b = (v4i8) {5, 6, 7, 8};
/* ----- */

Ex: /* v2q15 Example */
v2q15 a = {0x0fcb, 0x3a75};
v2q15 b;
b = (v2q15) {0.1234 * 32768.0, 0.4567 * 32768.0};
```

Note that CPU endianness affects the ordering of data stored in a register. In a Big Endian CPU, data is stored from the left-to-right location of a register. In a Little Endian CPU, data is stored from the right-to-left location of a register. For the example of v4i8 a = {1, 2, 3, 4}, a Big Endian CPU stores 1, 2, 3, and 4 from the left-to-right location, but a Little Endian CPU stores 1, 2, 3, and 4 from the right-to-left location as shown in [Figure 15.2](#).

**Figure 15.2 Register Values for v4i8 a = {1, 2, 3, 4} in Big Endian and Little Endian CPUs**



Most arithmetic operations will simply work on the SIMD operands in the register irrespective of endianness. But the programmer must beware of such instructions in the DSP-R2 ASE that directly refer to the left or right portions of a GPR. For example, MAQ\_SA.W.PHL.

### 15.7.5 Accessing Elements in SIMD Variables

The use of SIMD variables enables operations on multiple data in parallel. However, in certain situations, programmers need to access elements inside a SIMD variable. This can be done by using a union type that unites a SIMD type and an array of a basic type as follows.

```
typedef union
{
    v4i8 a;
    unsigned char b[4];
} v4i8_union;

typedef short q15;
typedef union
{
    v2q15 a;
    q15 b[2];
} v2q15_union;
```

As shown in [Figure 15.2](#) for a v4i8 variable, b[0] is used in both big-endian and Little Endian CPUs to access the first element in the variable. However, b[0] is stored in the left-most position in a Big Endian CPU, but it is stored in the right-most position in a little-endian CPU. The following examples show how to extract or assign elements.

```
Ex: /* v4i8 Example */
v4i8 i;
unsigned char j, k, l, m;
v4i8_union temp;

/* Assume we want to extract from i. */
temp.a = i;
j = temp.b[0];
k = temp.b[1];
l = temp.b[2];
m = temp.b[3];

/* Assume we want to assign j, k, l, m to i. */
temp.b[0] = j;
```

```

temp.b[1] = k;
temp.b[2] = l;
temp.b[3] = m;
i = temp.a;
/* ----- */

Ex: /* v2q15 Example */
v2q15 i;
q15 j, k;
v2q15_union temp;

/* Assume we want to extract from i. */
temp.a = i;
j = temp.b[0];
k = temp.b[1];

/* Assume we want to assign j, k to i. */
temp.b[0] = j;
temp.b[1] = k;
i = temp.a;

```

## 15.7.6 C Operators

Addition and subtraction on fractional data are similar to addition and subtraction with integer data, but multiplication requires a post-multiply shift to align the resulting values appropriately. Because fractional data uses integer data types in the GCC compiler, users must be very cautious when applying operators upon fractional data. The GCC compiler accepts all operators upon Q15 and Q31 data, but only addition and subtraction generate the expected results for Q15 and Q31. Note that operators other than “+” and “-” upon Q15 and Q31 are treated as integer arithmetic.

```

Ex: /* Q15 Example */
typedef short q15;
q15 i, j, k, l;
i = k + l;
j = k - l;
/* ----- */

Ex: /* Q31 Example */
typedef int q31;
q31 i, j, k, l;
i = k + l;
j = k - l;

```

Certain C operators can be applied to SIMD variables. They are +, -, \*, /, unary minus, ^, |, &, ~. The MIPS32 DSP-R2 ASE provides SIMD addition and subtraction instructions for v4i8 and v2q15, allowing the GCC compiler to generate appropriate instructions for addition and subtraction of v4i8 and v2q15 variables. For other operators, the GCC compiler synthesizes a sequence of instructions. The examples here show compiler-generated SIMD instructions when the appropriate operator is applied to SIMD variables.

```

Ex: /* v4i8 Addition */
v4i8 test1 (v4i8 a, v4i8 b)
{
    return a + b;
}

# Generated Assembly

```

```

test1:
    j            $31
    addu.qb $2, $4, $5
# -----

Ex: /* v4i8 Subtraction */
v4i8 test2 (v4i8 a, v4i8 b)
{
    return a - b;
}

# Generated Assembly
test2:
    j            $31
    subu.qb $2, $4, $5
# -----

Ex: /* v2q15 Addition */
v2q15 test3 (v2q15 a, v2q15 b)
{
    return a + b;
}

# Generated Assembly
test3:
    j            $31
    addq.ph $2, $4, $5
# -----

Ex: /* v2q15 Subtraction */
v2q15 test4 (v2q15 a, v2q15 b)
{
    return a - b;
}

# Generated Assembly
test4:
    j            $31
    subq.ph $2, $4, $5

```

In situations where special integer and fractional calculations are needed and the compiler cannot generate them automatically, C intrinsics can be directly used by the programmer as described in the next section.

### 15.7.7 C Intrinsics for the MIPS32 DSP-R2 ASE

Intrinsics are very similar to function calls in syntax. Programmers need to pass parameters to intrinsics, and intrinsics return results to variables. The difference between intrinsics and functions is that the compiler directly maps intrinsics to instructions for better performance. Intrinsics for the MIPS32 DSP-R2 ASE use the following data types:

```

typedef signed char v4i8 __attribute__((vector_size(4)));
typedef short v2i16 __attribute__((vector_size(4)));
typedef short v2q15 __attribute__((vector_size(4)));
typedef int q31;
typedef int i32;
typedef long long a64;
typedef unsigned int ui32;

```



NOTE: “q31” and “i32” are actually the same as “int”, but the intrinsic that accepts “q31” processes data as Q31 fractional data, and the intrinsic that accepts “i32” processes data as 32-bit integer data.

NOTE: “a64” is the same as “long long”, but the compiler allocates “a64” variables to accumulators (\$ac0, \$ac1, \$ac2, \$ac3) ready to be operated on relevant DSP-R2 instructions.

The list of all intrinsics can be found in the [MIPS32 DSP-R2 Intrinsics](#), and [Section 15.7.9](#) will introduce each MIPS32 DSP-R2 intrinsic. Programmers should be familiar with the semantics of all MIPS32 DSP-R2 instructions so that the corresponding intrinsic can be used appropriately in C programs to achieve better performance.

One example of an intrinsic for the **ADDQ.PH** instruction is “v2q15 \_\_builtin\_mips\_addq\_ph (v2q15, v2q15).” Two v2q15 variables are required to be passed to “\_\_builtin\_mips\_addq\_ph” and one v2q15 variable is needed to receive the returned result from this intrinsic. The following C code demonstrates the usage of “\_\_builtin\_mips\_addq\_ph”.

```
Ex:
v2q15 test5 ()
{
    v2q15 a, b, c;
    a = (v2q15) {0.12 * 32768.0, 0.34 * 32768.0};
    b = (v2q15) {0.56 * 32768.0, 0.78 * 32768.0};
    c = __builtin_mips_addq_ph (a, b);
    return c;
}

# Generated Assembly
    .file 1 "test5.c"
    .section .mdebug.abi32
    .previous
    .section .rodata.cst4,"aM",@progbits,4
    .align 2
.LC0:
    .half 3921
    .half 11141
    .align 2
.LC1:
    .half 18299
    .half 25559
    .text
    .align 2
    .align 3
    .globl test5
    .set nomips16
    .ent test5
test5:
    .frame $sp,0,$31 # vars= 0, regs= 0/0, args= 0, gp= 0
    .mask 0x00000000,0
    .fmask 0x00000000,0
    .set noreorder
    .set nomacro

    lui $5,%hi(.LC0)
    lui $4,%hi(.LC1)
    lw $2,%lo(.LC0)($5)
    lw $3,%lo(.LC1)($4)
    j $31
    addq.ph $2,$2,$3
```

```

.set    macro
.set    reorder
.end    test5
.ident  "GCC: (GNU) 3.4.4 mipssde-6.03.00-20051020"

```

## 15.7.8 Compiler Usage and the DSPControl Register

The MIPS32 DSP-R2 ASE includes a new DSP control register that has six fields as described in [Section 15.4.2 “DSP-R2 ASE Control Register”](#). These fields are:

- CCOND (condition code bits)
- OUFLAG (overflow/underflow bits)
- EFI (extract fail indicator bit)
- C (carry bit)
- SCOUNT (size count bits)
- POS (position bits).

The compiler treats the SCOUNT and POS fields as global variables, such that instructions that modify SCOUNT or POS are never optimized away. These instructions include **WRDSP**, **EXTPDP**, **EXTPDPV**, and **MTHLIP**. A function call that jumps to a function containing **WRDSP**, **EXTPDP**, **EXTPDPV**, or **MTHLIP** is also never deleted by the compiler.

For correctness, programmers must assume that a function call clobbers all fields of the DSP control register. That is, programmers cannot depend on the values in CCOND, OUFLAG, EFI or C across a function-call boundary. They must re-initialize the values of CCOND, OUFLAG, EFI or C before using them. Note that because SCOUNT and POS fields are treated as global variables, the values of SCOUNT and POS are always valid across function-call boundaries and can be used without re-initialization.

The following example shows possibly incorrect code. The first intrinsic “`__builtin_mips_addsc`” sets the carry bit (C) in the DSP control register, and the second intrinsic “`__builtin_mips_addwc`” reads the carry bit (C) from the DSP control register. However, a function call “`func`” inserted between “`__builtin_mips_addsc`” and “`__builtin_mips_addwc`” may change the carry bit to affect the correct result of “`__builtin_mips_addwc`”.

```

Ex:
int test6 (int a, int b, int c, int d)
{
    __builtin_mips_addsc (a, b);
    func();
    return __builtin_mips_addwc (c, d);
}

```

The previous example may be corrected by moving “`func`” before the first intrinsic or after the second intrinsic as follows.

```

Ex:
int test7 (int a, int b, int c, int d)
{
    func();
    __builtin_mips_addsc (a, b);
    return __builtin_mips_addwc (c, d);
}
/* ----- */

```

```

int test8 (int a, int b, int c, int d)
{
    int i;
    __builtin_mips_addsc (a, b);
    i = __builtin_mips_addwc (c, d);
    func();
    return i;
}

```

## 15.7.9 C-Based Intrinsic for the MIPS32® DSP-R2 ASE

This section provides a basic introduction to all the intrinsics supported for the MIPS32 DSP-R2 ASE. The intrinsics are illustrated using examples, and some usage tips are provided as well. They are categorized by function and data size type as follows:

- Intrinsics to access and use the DSPControl register
- Intrinsics for signed and unsigned 8-bit integers
- Intrinsics for Q15 data
- Intrinsics for Q31 data
- Intrinsics for mixed data types of 8-bit integers and Q15/16-bit integers
- Intrinsics for mixed data types of Q15 and Q31
- Intrinsics for 64-bit accumulators
- Intrinsics for 32-bit integers
- Intrinsics for 16-bit integers
- Intrinsics for mixed data types of 16-bit integers and 32-bit integers

### 15.7.9.1 Intrinsics for Instructions that Access and Use the DSPControl Register

#### *Read/Write the DSPControl Register*

```

i32 __builtin_mips_rddsp (imm0_63);
void __builtin_mips_wrdsp (i32, imm0_63);

```

The immediate parameter, `imm0_63` used in the two intrinsics here is a mask value used to specify which fields of the DSPControl register should be read or written respectively. The correspondence of the specific bits of the mask to specific fields in the DSPControl register is shown in [Figure 15.3](#). As shown, bit 0 of `imm0_63` is for the POS field, bit 1 of `imm0_63` is for the SCOUNT field, bit 2 of `imm0_63` is for the C field, bit 3 of `imm0_63` is for the OUFLAG, bit 4 of `imm0_63` is for the CCOND flag, and bit 5 of `imm0_63` is for the EFI field. For example, to read the SCOUNT field, `imm0_63` must be set to 2. To read all fields, `imm0_63` must be set to 63 (1 + 2 + 4 + 8 + 16 + 32).

Ex:

```

int the_scount = (__builtin_mips_rddsp (2)) >> 7; // Read SCOUNT
int all_fields = __builtin_mips_rddsp (63); // Read all fields

```

To write the DSPControl register, programmers must pass a 32-bit integer as the first parameter to “\_\_builtin\_mips\_wrdsb”, as well as the imm0\_63 mask value that determines which fields are to be updated. The first parameter should be a 32-bit value that mimics the format of the DSPControl register fields. Then, based on the mask value, the corresponding fields will be copied from this 32-bit value to the DSPControl register. For example, to set the SCOUNT field to 63, (63<<7) is passed as the first parameter and second parameter imm0\_63 must be 2 so that an update of the SCOUNT field is done to the value 63 from the first input. To update all bits of all fields to 1, the first parameter can be 0xFFFFFFFF with a second parameter value of 63.

Ex:  
 \_\_builtin\_mips\_wrdsb (63<<7, 2); // Update SCOUNT to 63  
 \_\_builtin\_mips\_wrdsb (0xFFFFFFFF, 63); // Update all bits of fields to 1

**Figure 15.3 Mask Value to Access the MIPS32 DSPControl Register**

<b>Bit 31</b>	<b>28 27</b>	<b>24 23</b>	<b>16 15 14 13 12</b>	<b>7 6 5</b>	<b>0</b>
<b>0</b>	<b>CCOND</b>	<b>OUFLAG</b>	<b>0 EFC SCOUNT</b>	<b>0</b>	<b>POS</b>
<b>IMM0_63 =&gt;</b>	<b>16</b>	<b>8</b>	<b>32 4 2</b>		<b>1</b>

**Branch on Greater Than or Equal to 32 in POS**

```
i32 __builtin_mips_bposge32 ();
```

This intrinsic returns 1 if the value of the POS field is greater than or equal to 32. Otherwise, the intrinsic returns 0. Programmers can use “\_\_builtin\_mips\_bposge32” inside a condition test, and the compiler will optimize the code to generate the “bposge32” instruction as follows.

Ex:  
 void test9 ()  
 {  
 if (\_\_builtin\_mips\_bposge32())  
 result\_is\_true();  
 else  
 result\_is\_false();  
 }  
 # Generated Assembly  
 test9:  
 .set noreorder  
 .set nomacro  
 bposge32 .L3  
 nop  
 j result\_is\_false  
 nop  
 .align 3  
 .L3:  
 j result\_is\_true  
 nop

**15.7.9.2 Using Intrinsics for Signed and Unsigned 8-bit Integers**

This section introduces intrinsics that operate on signed and unsigned 8-bit integers in register SIMD fashion by using a “v4i8” data type. If the programmer wants to perform an operation such as add on a single 8-bit item, then these intrinsics can still be used by ignoring the other three un-used elements inside the “v4i8” variable. Each set of intrinsics for operations are listed below, with simple examples of their usage.

### **Unsigned Add/Subtract with Optional Saturation**

```
v4i8 __builtin_mips_addu_qb (v4i8, v4i8);
v4i8 __builtin_mips_addu_s_qb (v4i8, v4i8);
v4i8 __builtin_mips_subu_qb (v4i8, v4i8);
v4i8 __builtin_mips_subu_s_qb (v4i8, v4i8);
#-----
Ex:
v4i8 a = {1, 2, 3, 0xFF};
v4i8 b = {2, 4, 6, 8};
v4i8 r1, r2, r3, r4;
r1 = __builtin_mips_addu_qb (a, b); // r1 will be {3, 6, 9, 7}
r2 = __builtin_mips_addu_s_qb (a, b); // r2 will be {3, 6, 9, 0xFF}
r3 = __builtin_mips_subu_qb (a, b); // r3 will be {0xFF, 0xFE, 0xFD, 0xF7}
r4 = __builtin_mips_subu_s_qb (a, b); // r4 will be {0, 0, 0, 0xF7}
```

### **Unsigned Reduction Add**

```
i32 __builtin_mips_raddu_w_qb (v4i8);
#-----
Ex:
v4i8 a = {1, 2, 3, 4};
int sum = __builtin_mips_raddu_w_qb (a); // sum will be 1 + 2 + 3 + 4 = 10
```

### **Shift Left/Right Logical**

```
v4i8 __builtin_mips_shll_qb (v4i8, imm0_7);
v4i8 __builtin_mips_shll_qb (v4i8, i32);
v4i8 __builtin_mips_shrl_qb (v4i8, imm0_7);
v4i8 __builtin_mips_shrl_qb (v4i8, i32);
#-----
Ex:
v4i8 a = {1, 2, 3, 4};
v4i8 b = {128, 64, 32, 16};
v4i8 r1, r2, r3, r4;
int shift_amount = 2;
r1 = __builtin_mips_shll_qb (a, 1); // r1 will be {2, 4, 6, 8}
r2 = __builtin_mips_shll_qb (a, shift_amount); // r2 will be {4, 8, 12, 16}
r3 = __builtin_mips_shrl_qb (b, 3); // r3 will be {16, 8, 4, 2}
r4 = __builtin_mips_shrl_qb (b, shift_amount); // r4 will be {32, 16, 8, 4}
```

### **Dot Product with Accumulate/Subtract**

```
a64 __builtin_mips_dpau_h_qbl (a64, v4i8, v4i8);
a64 __builtin_mips_dpau_h_qbr (a64, v4i8, v4i8);
a64 __builtin_mips_dpsu_h_qbl (a64, v4i8, v4i8);
a64 __builtin_mips_dpsu_h_qbr (a64, v4i8, v4i8);
```

#### **NOTES:**

1. The result will be a 64-bit integer.
2. The processor endianness of the data affects the format of the result.
3. Using the same “a64” variable for both the target and the first parameter could result in better performance.

```
#-----
Ex: /* Assume a big-endian CPU */
v4i8 a = {1, 2, 3, 4};
v4i8 b = {4, 5, 6, 7};
a64 ac1, ac2, ac3, ac4;
```

```

ac1 = ac2 = ac3 = ac4 = 0;
ac1 = __builtin_mips_dpau_h_qbl (ac1, a, b); // ac1 will be 0 + 1*4 + 2*5 = 14
ac2 = __builtin_mips_dpau_h_qbr (ac2, a, b); // ac2 will be 0 + 3*6 + 4*7 = 46
ac3 = __builtin_mips_dpsu_h_qbl (ac3, a, b); // ac3 will be 0 - (1*4 + 2*5) = -14
ac4 = __builtin_mips_dpsu_h_qbr (ac4, a, b); // ac4 will be 0 - (3*6 + 4*7) = -46

```

### **Replicate a Fixed Byte Value into SIMD Elements**

```

v4i8 __builtin_mips_repl_qb (imm0_255);
v4i8 __builtin_mips_repl_qb (i32);
#-----
Ex:
v4i8 a, b;
int value = 100;
a = __builtin_mips_repl_qb (10); // a will be {10, 10, 10, 10}
b = __builtin_mips_repl_qb (value); // b will be {100, 100, 100, 100}

```

### **Compare Unsigned**

```

void __builtin_mips_cmpu_eq_qb (v4i8, v4i8);
void __builtin_mips_cmpu_lt_qb (v4i8, v4i8);
void __builtin_mips_cmpu_le_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgu_eq_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgu_lt_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgu_le_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgdu_eq_qb (v4i8, v4i8); // DSPR2
i32 __builtin_mips_cmpgdu_lt_qb (v4i8, v4i8); // DSPR2
i32 __builtin_mips_cmpgdu_le_qb (v4i8, v4i8); // DSPR2

```

Note that the first three intrinsics update the condition code bits of the DSPControl register, but the middle three intrinsics write the condition code bits to a specified general purpose register. The last three intrinsics do the both.

```

#-----
Ex: /* Assume a big-endian CPU */
v4i8 a = {1, 4, 10, 8};
v4i8 b = {1, 2, 100, 8};
int r1, r2, r3;
__builtin_mips_cmpu_eq_qb (a, b); // CCOND bits will be 9 (= 1001b)
__builtin_mips_cmpu_lt_qb (a, b); // CCOND bits will be 2 (= 0010b)
__builtin_mips_cmpu_le_qb (a, b); // CCOND bits will be 11 (= 1011b)
r1 = __builtin_mips_cmpgu_eq_qb (a, b); // r1 will be 9
r2 = __builtin_mips_cmpgu_lt_qb (a, b); // r2 will be 2
r3 = __builtin_mips_cmpgu_le_qb (a, b); // r3 will be 11
r1 = __builtin_mips_cmpgdu_eq_qb (a, b); // Both CCOND bits and r1 will be 9
r2 = __builtin_mips_cmpgdu_lt_qb (a, b); // Both CCOND bits and r2 will be 2
r3 = __builtin_mips_cmpgdu_le_qb (a, b); // Both CCOND bits and r3 will be 11

```

### **Pick Based on Condition Code Bits**

```

v4i8 __builtin_mips_pick_qb (v4i8, v4i8);

```

Note that this intrinsic is usually used together with the first three compare intrinsics in [Section](#) .

```

#-----
Ex:
v4i8 a = {1, 4, 10, 8};
v4i8 b = {1, 2, 100, 8};
v4i8 r;

```

```

__builtin_mips_cmpu_eq_qb (a, b); // CCOND bits will be 9 (= 1001b)
r = __builtin_mips_pick_qb (a, b); // r will be {1, 2, 100, 8}

```

### **Find Absolute Value**

```

v4i8 __builtin_mips_absq_s_qb (v4i8); // DSPR2
#-----
Ex:
v4i8 a = {-1, -128, 1, 127};
v4i8 r;
r = __builtin_mips_absq_s_qb (a); // r will be {1, 127, 1, 127}
/* Note that the absolute value of -128 is 128 that is represented by the maximum
value as 127. */

```

### **Unsigned Add and Right Shift to Halve Results with Optional Rounding**

```

v4i8 __builtin_mips_adduh_qb (v4i8, v4i8); // DSPR2
v4i8 __builtin_mips_adduh_r_qb (v4i8, v4i8); // DSPR2
#-----
Ex:
v4i8 a = {1, 2, 3, 4};
v4i8 b = {0x80, 0x80, 0x80, 0x80};
v4i8 r1, r2;
r1 = __builtin_mips_adduh_qb (a, b); // r1 will be {0x40, 0x41, 0x41, 0x42}
r2 = __builtin_mips_adduh_r_qb (a, b); // r2 will be {0x41, 0x41, 0x44, 0x42}

```

### **Shift Right Arithmetic with Optional Rounding**

```

v4i8 __builtin_mips_shra_qb (v4i8, imm0_7); // DSPR2
v4i8 __builtin_mips_shra_r_qb (v4i8, imm0_7); // DSPR2
v4i8 __builtin_mips_shra_qb (v4i8, i32); // DSPR2
v4i8 __builtin_mips_shra_r_qb (v4i8, i32); // DSPR2
#-----
Ex:
v4i8 a = {0x40, 0x20, 0x10, 0x0F};
v4i8 r1, r2, r3, r4;
int shift_amount = 2;
r1 = __builtin_mips_shra_qb (a, 2); // r1 will be {0x10, 0x08, 0x04, 0x3}
r2 = __builtin_mips_shra_r_qb (a, 2); // r2 will be {0x10, 0x08, 0x04, 0x4}
r3 = __builtin_mips_shra_qb (a, shift_amount); // r3 will be {0x10, 0x08, 0x04, 0x3}
r4 = __builtin_mips_shra_r_qb (a, shift_amount); // r4 will be {0x10, 0x08, 0x04, 0x4}

```

### **Unsigned Subtract and Right Shift to Halve Results with Optional Rounding**

```

v4i8 __builtin_mips_subuh_qb (v4i8, v4i8); // DSPR2
v4i8 __builtin_mips_subuh_r_qb (v4i8, v4i8); // DSPR2
#-----
Ex:
v4i8 a = {0x80, 0x80, 0x80, 0x80};
v4i8 b = {1, 2, 3, 4};
v4i8 r1, r2;
r1 = __builtin_mips_subuh_qb (a, b); // r1 will be {0x3F, 0x3F, 0x3E, 0x3E}
r2 = __builtin_mips_subuh_r_qb (a, b); // r2 will be {0x40, 0x3F, 0x3F, 0x3E}

```

### 15.7.9.3 Using Intrinsics for Q15 Data Type

This section introduces intrinsics that operate on Q15 data present in register SIMD fashion by using a “v2q15” data type. If the programmer wants to perform the specified operation on a single data in the register, then these intrinsics can still be used while ignoring the other element inside the “v2q15” variable.

#### **Add/Subtract with Optional Saturation**

```
v2q15 __builtin_mips_addq_ph (v2q15, v2q15);
v2q15 __builtin_mips_addq_s_ph (v2q15, v2q15);
v2q15 __builtin_mips_subq_ph (v2q15, v2q15);
v2q15 __builtin_mips_subq_s_ph (v2q15, v2q15);
#-----
Ex:
v2q15 a = {0x0000, 0x8000};
v2q15 b = {0x8000, 0x8000};
v2q15 r1, r2, r3, r4;
r1 = __builtin_mips_addq_ph (a, b); // r1 will be {0x8000, 0x0000}
r2 = __builtin_mips_addq_s_ph (a, b); // r2 will be {0x8000, 0x8000}
r3 = __builtin_mips_subq_ph (a, b); // r3 will be {0x8000, 0x0000}
r4 = __builtin_mips_subq_s_ph (a, b); // r4 will be {0x7FFF, 0x0000}
```

#### **Find Absolute Value**

```
v2q15 __builtin_mips_absq_s_ph (v2q15);
#-----
Ex:
v2q15 a = {0xFFFF, 0x8000};
v2q15 r;
r = __builtin_mips_absq_s_ph (a); // r will be {0x0001, 0x7FFF}
/* Note that the value of 0x8000 is -1 in Q15. The absolute value of -1 is 1 that
is represented by the maximum value as 0x7FFF in Q15. */
```

#### **Shift Left Logical with Optional Saturation**

```
v2q15 __builtin_mips_shll_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shll_ph (v2q15, i32);
v2q15 __builtin_mips_shll_s_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shll_s_ph (v2q15, i32);
#-----
Ex:
v2q15 a = {0x0001, 0x8000};
v2q15 r1, r2, r3, r4;
int shift_amount = 2;
r1 = __builtin_mips_shll_ph (a, 1); // r1 will be {0x0002, 0x0000}
r2 = __builtin_mips_shll_ph (a, shift_amount); // r2 will be {0x0004, 0x0000}
r3 = __builtin_mips_shll_s_ph (a, 1); // r3 will be {0x0002, 0x8000}
r4 = __builtin_mips_shll_s_ph (a, shift_amount); // r4 will be {0x0004, 0x8000}
```

#### **Shift Right Arithmetic with Optional Rounding**

```
v2q15 __builtin_mips_shra_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shra_ph (v2q15, i32);
v2q15 __builtin_mips_shra_r_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shra_r_ph (v2q15, i32);
#-----
Ex:
```



```

v2q15 a = {0x7FFF, 0x8000};
v2q15 r1, r2, r3, r4;
int shift_amount = 2;
r1 = __builtin_mips_shra_ph (a, 1); // r1 will be {0x3FFF, 0xC000}
r2 = __builtin_mips_shra_ph (a, shift_amount); // r2 will be {0x1FFF, 0xE000}
r3 = __builtin_mips_shra_r_ph (a, 1); // r3 will be {0x4000, 0xC000}
r4 = __builtin_mips_shra_r_ph (a, shift_amount); // r4 will be {0x2000, 0xE000}

```

### **Multiply with Rounding and Saturation (Q15 x Q15 => Q15)**

```

v2q15 __builtin_mips_mulq_rs_ph (v2q15, v2q15);
#-----
Ex:
v2q15 a = {0x7FFF, 0x8000};
v2q15 b = {0x7FFF, 0x8000};
v2q15 r;
r = __builtin_mips_mulq_rs_ph (a, b); // r will be {0x7FFE, 0x7FFF}

```

### **Dot Product with Accumulate/Subtract (Q15 x Q15 => Q32.31)**

```

a64 __builtin_mips_dpaq_s_w_ph (a64, v2q15, v2q15);
a64 __builtin_mips_dpsq_s_w_ph (a64, v2q15, v2q15);

```

#### NOTES:

1. The result will be in the Q32.31 format.
2. Using the same “a64” variable for the target and the first parameter could lead to better performance.

```

#-----
Ex:
v2q15 a = {0x0001, 0x8000};
v2q15 b = {0x0002, 0x8000};
a64 ac1, ac2;
ac1 = ac2 = 0;
ac1 = __builtin_mips_dpaq_s_w_ph (ac1, a, b); // ac1 will be 0 + (1*2)<<1 +
// 0x7FFFFFFF = 0x0000000080000003
ac2 = __builtin_mips_dpsq_s_w_ph (ac2, a, b); // ac2 will be 0 - (1*2)<<1 -
// 0x7FFFFFFF = 0xFFFFFFFF7FFFFFFFD

```

### **Multiply and Subtract and Accumulate (Q15 x Q15 => Q32.31)**

```

a64 __builtin_mips_mulsaq_s_w_ph (a64, v2q15, v2q15);

```

#### NOTES:

1. The result will be in the Q32.31 format.
2. The processor endianness affects the format of the result.
3. Using the same “a64” variable for the target and the first parameter could lead to better performance.

```

#-----
Ex: /* Assume a big-endian CPU */
v2q15 a = {0x0001, 0x8000};
v2q15 b = {0x0002, 0x8000};
a64 ac1 = 0;
ac1 = __builtin_mips_mulsaq_s_w_ph (ac1, a, b); // ac1 will be 0 + (1*2)<<1 -
// 0x7FFFFFFF = 0xFFFFFFFF80000005

```

### **Multiply with Accumulate a Single Element (Q15 x Q15 => Q31)**

```

a64 __builtin_mips_maq_s_w_ph1 (a64, v2q15, v2q15);
a64 __builtin_mips_maq_s_w_phr (a64, v2q15, v2q15);

```

```
a64 __builtin_mips_maq_sa_w_phl (a64, v2q15, v2q15);
a64 __builtin_mips_maq_sa_w_phr (a64, v2q15, v2q15);
```

NOTES:

1. The result will be in the Q31 format.
2. The processor endianness affects the format of the result.
3. Using the same “a64” variable for the target and the first parameter could lead to better performance.

```
#-----
Ex: /* Assume a big-endian CPU */
v2q15 a = {0x0001, 0x8000};
v2q15 b = {0x0002, 0x8000};
a64 ac1, ac2, ac3, ac4;
ac1 = ac2 = 0;
ac3 = ac4 = 0x7FFFFFF0;
ac1 = __builtin_mips_maq_s_w_phl (ac1, a, b); // ac1 will be 0 + (1*2)<<1 =
// 0x4
ac2 = __builtin_mips_maq_s_w_phr (ac2, a, b); // ac2 will be 0 + 0x7FFFFFFF =
// 0x7FFFFFFF
ac3 = __builtin_mips_maq_sa_w_phl (ac3, a, b); // ac3 will be 0x7FFFFFF0 +
// (1*2)<<1 = 0x7FFFFFF4
ac4 = __builtin_mips_maq_sa_w_phr (ac4, a, b); // ac4 will be 0x7FFFFFF0 +
// 0x7FFFFFFF = 0x7FFFFFFF
```

**Multiply Vector Fractional Left/Right Half-Words to Expanded Width Product with Saturation (Q15 x Q15 => Q31)**

```
q31 __builtin_mips_muleq_s_w_phl (v2q15, v2q15);
q31 __builtin_mips_muleq_s_w_phr (v2q15, v2q15);
```

NOTES:

1. The result will be in the Q31 format.
2. The processor endianness affects the format of the result.

```
#-----
Ex: /* Assume a big-endian CPU */
v2q15 a = {0x1234, 0x8000};
v2q15 b = {0x5678, 0x8000};
q31 r1, r2;
r1 = __builtin_mips_muleq_s_w_phl (a, b); // r1 will be 0x0C4C00C0
r2 = __builtin_mips_muleq_s_w_phr (a, b); // r2 will be 0x7FFFFFFF
```

**Replicate a Fixed Half-word into Elements**

```
v2q15 __builtin_mips_repl_ph (imm_n512_511);
v2q15 __builtin_mips_repl_ph (i32);
```

Note that for the immediate version, imm\_n512\_511 will be sign-extended to a 16-bit value and replicated into each SIMD element.

```
#-----
Ex:
v2q15 r1, r2;
int value = 0x1234;
r1 = __builtin_mips_repl_ph (-512); // r1 will be {0xFE00, 0xFE00};
r2 = __builtin_mips_repl_ph (value); // r2 will be {0x1234, 0x1234};
```

## Compare

```
void __builtin_mips_cmp_eq_ph (v2q15, v2q15);
void __builtin_mips_cmp_lt_ph (v2q15, v2q15);
void __builtin_mips_cmp_le_ph (v2q15, v2q15);
Ex:
v2q15 a = {0x1111, 0x1234};
v2q15 b = {0x4444, 0x1234};
__builtin_mips_cmp_eq_ph (a, b); // CCOND bits will be 1 (= 01b)
__builtin_mips_cmp_lt_ph (a, b); // CCOND bits will be 2 (= 10b)
__builtin_mips_cmp_le_ph (a, b); // CCOND bits will be 3 (= 11b)
```

## Pick Based on Condition Code Bits

```
v2q15 __builtin_mips_pick_ph (v2q15, v2q15);
```

Note that this intrinsic is usually used together with the compare intrinsics in [Section](#) .

```
#-----
Ex:
v2q15 a = {0x1111, 0x1234};
v2q15 b = {0x4444, 0x1234};
v2q15 r;
__builtin_mips_cmp_eq_ph (a, b); // CCOND bits will be 1 (= 01b)
r = __builtin_mips_pick_ph (a, b); // r will be {0x4444, 0x1234}
```

## Pack from the Right and Left

```
v2q15 __builtin_mips_packrl_ph (v2q15, v2q15);
```

Note that the endianness affects the result.

```
#-----
Ex: /* Assume a big-endian CPU */
v2q15 a = {0x1111, 0x2222};
v2q15 b = {0x3333, 0x4444};
v2q15 r;
r = __builtin_mips_packrl_ph (a, b); // r will be {0x2222, 0x3333}
```

## Multiply with Saturation (Q15 x Q15 => Q15)

```
v2q15 __builtin_mips_mulq_s_ph (v2q15, v2q15); // DSPR2
#-----
Ex:
v2q15 a = {0x7FFF, 0x8000};
v2q15 b = {0x7FFF, 0x8000};
v2q15 r;
r = __builtin_mips_mulq_s_ph (a, b); // r will be {0x7FFE, 0x7FFF}
```

## Add and Right Shift to Halve Results with Optional Rounding

```
v2q15 __builtin_mips_addqh_ph (v2q15, v2q15); // DSPR2
v2q15 __builtin_mips_addqh_r_ph (v2q15, v2q15); // DSPR2
#-----
Ex:
v2q15 a = {0x1000, 0x1000};
v2q15 b = {0x1001, 0x1000};
v2q15 r1, r2;
r1 = __builtin_mips_addqh_ph (a, b); // r1 will be {0x1000, 0x1000}
```

```
r2 = __builtin_mips_addqh_r_ph (a, b); // r1 will be {0x1001, 0x1000}
```

### **Subtract and Right Shift to Halve Results with Optional Rounding**

```
v2q15 __builtin_mips_subqh_ph (v2q15, v2q15); // DSPR2
v2q15 __builtin_mips_subqh_r_ph (v2q15, v2q15); // DSPR2
```

```
#-----
```

Ex:

```
v2q15 a = {0x1000, 0x1000};
v2q15 b = {0x1001, 0x1000};
v2q15 r1, r2;
r1 = __builtin_mips_subqh_ph (a, b); // r1 will be {0xFFFF, 0x0000}
r2 = __builtin_mips_subqh_r_ph (a, b); // r1 will be {0x0000, 0x0000}
```

### **Cross Dot Product with Accumulate/Subtract (Q15 x Q15 => Q32.31)**

```
a64 __builtin_mips_dpaqx_s_w_ph (a64, v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpaqx_sa_w_ph (a64, v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpsqx_s_w_ph (a64, v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpsqx_sa_w_ph (a64, v2q15, v2q15); // DSPR2
```

NOTES:

1. The result will be in the Q32.31 format.
2. Using the same “a64” variable for the target and the first parameter could lead to better performance.

```
#-----
```

Ex:

```
v2q15 a = {0x0002, 0x8000};
v2q15 b = {0x8000, 0x0003};
a64 ac1, ac2, ac3, ac4;
ac1 = ac2 = ac3 = ac4 = 0;
ac1 = __builtin_mips_dpaqx_s_w_ph (ac1, a, b); // ac1 will be
// 0 + (2*3)<<1 + 0x7FFFFFFF
// = 0x000000008000000B
ac2 = __builtin_mips_dpaqx_sa_w_ph (ac2, a, b); // ac2 will be saturated to
// 0x000000007FFFFFFF
ac3 = __builtin_mips_dpsqx_s_w_ph (ac3, a, b); // ac3 will be
// 0 - (2*3)<<1 - 0x7FFFFFFF
// = 0xFFFFFFFF7FFFFFF5
ac4 = __builtin_mips_dpsqx_sa_w_ph (ac4, a, b); // ac4 will be saturated
// to 0xFFFFFFFF80000000
```

## **15.7.9.4 Using Intrinsics for Q31 Data Type**

### **Add/Subtract with Saturation**

```
q31 __builtin_mips_addq_s_w (q31, q31);
q31 __builtin_mips_subq_s_w (q31, q31);
```

```
#-----
```

Ex:

```
q31 a = 0x12345678;
q31 b = 0x7FFFFFFF;
q31 r1, r2;
r1 = __builtin_mips_addq_s_w (a, b); // r1 will be 0x7FFFFFFF
r2 = __builtin_mips_subq_s_w (a, b); // r2 will be 0x92345679
```

### **Find Absolute Value**

```
q31 __builtin_mips_absq_s_w (q31);
#-----
Ex:
q31 a = 0x80000000;
q31 r;
r = __builtin_mips_absq_s_w (a); // r will be 0x7FFFFFFF
/* Note that the value of 0x80000000 is -1 in Q31. The absolute value of -1 is 1
that is represented by the maximum value as 0x7FFFFFFF in Q31. */
```

### **Shift Left Logical with Saturation**

```
q31 __builtin_mips_shll_s_w (q31, imm0_31);
q31 __builtin_mips_shll_s_w (q31, i32);
#-----
Ex:
q31 a = 0x70000000;
q31 r1, r2;
int shift_amount = 2;
r1 = __builtin_mips_shll_s_w (a, 1); // r1 will be 0x7FFFFFFF
r2 = __builtin_mips_shll_s_w (a, shift_amount); // r2 will be 0x7FFFFFFF
```

### **Shift Right Arithmetic with Rounding**

```
q31 __builtin_mips_shra_r_w (q31, imm0_31);
q31 __builtin_mips_shra_r_w (q31, i32);
#-----
Ex:
q31 a = 0x7FFFFFFF;
q31 r1, r2;
int shift_amount = 2;
r1 = __builtin_mips_shra_r_w (a, 1); // r1 will be 0x40000000
r2 = __builtin_mips_shra_r_w (a, shift_amount); // r2 will be 0x20000000
```

### **Dot Product with Accumulate/Subtract (Q31 x Q31 => Q63)**

```
a64 __builtin_mips_dpaq_sa_l_w (a64, q31, q31);
a64 __builtin_mips_dpsq_sa_l_w (a64, q31, q31);
```

#### NOTES:

1. The result will be in the Q63 format.
2. Using the same “a64” variable for the target and the first parameter could lead to better performance.

```
#-----
Ex:
q31 a = 0x80000000;
q31 b = 0x80000000;
a64 ac1, ac2;
ac1 = ac2 = 1;
ac1 = __builtin_mips_dpaq_sa_l_w (ac1, a, b); // ac1 will be 0x7FFFFFFFFFFFFFFF
ac2 = __builtin_mips_dpsq_sa_l_w (ac2, a, b); // ac2 will be 0x8000000000000002
```

### **Multiply with Rounding and Saturation (Q31 x Q31 => Q31)**

```
q31 __builtin_mips_mulq_rs_w (q31, q31); // DSPR2
#-----
Ex:
```

```

q31 a = 0x7FFFFFFF;
q31 b = 0x00000001;
q31 r;
r = __builtin_mips_mulq_rs_w (a, b); // r will be 0x00000001

```

### **Multiply with Saturation (Q31 x Q31 => Q31)**

```

q31 __builtin_mips_mulq_s_w (q31, q31); // DSPR2
#-----
Ex:
q31 a = 0x80000000;
q31 b = 0x00000001;
q31 r;
r = __builtin_mips_mulq_s_w (a, b); // r will be 0x00000000

```

### **Add and Right Shift to Halve Results with Optional Rounding**

```

q31 __builtin_mips_addqh_w (q31, q31); // DSPR2
q31 __builtin_mips_addqh_r_w (q31, q31); // DSPR2
#-----
Ex:
q31 a = 0x10000000;
q31 b = 0x10000001;
q31 r1, r2;
r1 = __builtin_mips_addqh_w (a, b); // r1 will be 0x10000000
r2 = __builtin_mips_addqh_r_w (a, b); // r2 will be 0x10000001

```

### **Subtract and Right Shift to Halve Results with Optional Rounding**

```

q31 __builtin_mips_subqh_w (q31, q31); // DSPR2
q31 __builtin_mips_subqh_r_w (q31, q31); // DSPR2
#-----
Ex:
q31 a = 0x10000000;
q31 b = 0x10000001;
q31 r1, r2;
r1 = __builtin_mips_subqh_w (a, b); // r1 will be 0xFFFFFFFF
r2 = __builtin_mips_subqh_r_w (a, b); // r2 will be 0x00000000

```

## **15.7.9.5 Using Intrinsics for Mixed Data Types: 8-bit Integers and Q15/16-bit Integers**

### **Precision Reduce Four Fractional Half-words to Four Bytes**

```

v4i8 __builtin_mips_preocrq_qb_ph (v2q15, v2q15);

```

Note that the processor endianness affects the format of the result.

```

#-----
Ex: /* Assume a big-endian CPU */
v2q15 a = {0x1234, 0x5678};
v2q15 b = {0x1111, 0x2222};
v4i8 r;
r = __builtin_mips_preocrq_qb_ph (a, b); // r will be {0x12, 0x56, 0x11, 0x22}

```

### **Precision Reduce Unsigned Four Fractional Half-words to Four Bytes with Saturation**

```

v4i8 __builtin_mips_preocrqu_s_qb_ph (v2q15, v2q15);

```

Note that the processor endianness affects the format of the result.

```
#-----  
Ex: /* Assume a big-endian CPU */  
v2q15 a = {0x7F79, 0xFFFF};  
v2q15 b = {0x7F81, 0x2000};  
v4i8 r;  
r = __builtin_mips_preocrqu_s_qb_ph (a, b); // r will be {0xFE, 0x00, 0xFF, 0x40}
```

### **Precision Expand Two Unsigned Bytes to Fractional Half-word Values**

```
v2q15 __builtin_mips_precequ_ph_qbl (v4i8);  
v2q15 __builtin_mips_precequ_ph_qbr (v4i8);  
v2q15 __builtin_mips_precequ_ph_qbla (v4i8);  
v2q15 __builtin_mips_precequ_ph_qbra (v4i8);
```

Note that the processor endianness affects the format of the result.

```
#-----  
Ex: /* Assume a big-endian CPU */  
v4i8 a = {0x12, 0x34, 0x56, 0x78};  
v2q15 r1, r2, r3, r4;  
r1 = __builtin_mips_precequ_ph_qbl (a, b); // r1 will be {0x0900, 0x1A00}  
r2 = __builtin_mips_precequ_ph_qbr (a, b); // r2 will be {0x2B00, 0x3C00}  
r3 = __builtin_mips_precequ_ph_qbla (a, b); // r3 will be {0x0900, 0x2B00}  
r4 = __builtin_mips_precequ_ph_qbra (a, b); // r4 will be {0x1A00, 0x3C00}
```

### **Precision Expand Two Unsigned Bytes to Unsigned Integer Half-words**

```
v2q15 __builtin_mips_preceu_ph_qbl (v4i8);  
v2q15 __builtin_mips_preceu_ph_qbr (v4i8);  
v2q15 __builtin_mips_preceu_ph_qbla (v4i8);  
v2q15 __builtin_mips_preceu_ph_qbra (v4i8);
```

Note that the processor endianness affects the format of the result.

```
#-----  
Ex: /* Assume a big-endian CPU */  
v4i8 a = {0x12, 0x34, 0x56, 0x78};  
v2q15 r1, r2, r3, r4;  
r1 = __builtin_mips_preceu_ph_qbl (a, b); // r1 will be {0x0012, 0x0034}  
r2 = __builtin_mips_preceu_ph_qbr (a, b); // r2 will be {0x0056, 0x0078}  
r3 = __builtin_mips_preceu_ph_qbla (a, b); // r3 will be {0x0012, 0x0056}  
r4 = __builtin_mips_preceu_ph_qbra (a, b); // r4 will be {0x0034, 0x0078}
```

### **Multiply Unsigned Vector Left/Right Bytes with Half-Words to Half Word Products with Saturation (Int8 x Q15 => Q15)**

```
v2q15 __builtin_mips_muleu_s_ph_qbl (v4i8, v2q15);  
v2q15 __builtin_mips_muleu_s_ph_qbr (v4i8, v2q15);
```

Note that the processor endianness affects the format of the result.

```
#-----  
Ex: /* Assume a big-endian CPU */  
v4i8 a = {0x1, 0x3, 0x5, 0x7};  
v2q15 b = {0x1234, 0x5678};  
v2q15 r1, r2;  
r1 = __builtin_mips_muleu_s_ph_qbl (a, b); // r1 will be {0x1234, 0xFFFF}  
r2 = __builtin_mips_muleu_s_ph_qbr (a, b); // r2 will be {0x5B04, 0xFFFF}
```

### **Precision Reduce Four Integer Half-words to Four Bytes**

```
v4i8 __builtin_mips_preocr_qb_ph (v2i16, v2i16); // DSPR2
```

Note that the processor endianness affects the format of the result.

```
#-----  
Ex: /* Assume a big-endian CPU */  
v2i16 a = {0x7F79, 0xFFFF};  
v2i16 b = {0x7F81, 0x2000};  
v4i8 r;  
r = __builtin_mips_preocr_qb_ph (a, b); // r will be {0x79, 0xFF, 0x81, 0x00}
```

### **15.7.9.6 Using Intrinsics for Mixed Data Types: Q15 and Q31**

#### **Precision Reduce Two Fractional Words to Two Half-Words**

```
v2q15 __builtin_mips_preocrq_ph_w (q31, q31);  
#-----  
Ex:  
q31 a = {0x12345678};  
q31 b = {0x11112222};  
v2q15 r;  
r = __builtin_mips_preocrq_ph_w (a, b); // r will be {0x1234, 0x1111}
```

#### **Precision Reduce Two Fractional Words to Two Half-Words with Rounding and Saturation**

```
v2q15 __builtin_mips_preocrq_rs_ph_w (q31, q31);  
#-----  
Ex:  
q31 a = {0x7000FFFF};  
q31 b = {0x80000000};  
v2q15 r;  
r = __builtin_mips_preocrq_rs_ph_w (a, b); // r will be {0x7001, 0x8000}
```

#### **Precision Expand a Fractional Half-word to a Fractional Word Value**

```
q31 __builtin_mips_preocrq_w_phl (v2q15);  
q31 __builtin_mips_preocrq_w_phr (v2q15);  
  
Note that the endianness affects the result.  
#-----  
Ex: /* Assume a big-endian CPU */  
v2q15 a = {0x1234, 0x5678};  
q31 r1, r2;  
r1 = __builtin_mips_preocrq_w_phl (a, b); // r1 will be 0x12340000  
r2 = __builtin_mips_preocrq_w_phr (a, b); // r2 will be 0x56780000
```

### **15.7.9.7 Using Intrinsics for 64-bit Accumulators**

#### **Extract a Value with Right Shift**

```
i32 __builtin_mips_extr_w (a64, imm0_31);  
i32 __builtin_mips_extr_w (a64, i32);  
i32 __builtin_mips_extr_r_w (a64, imm0_31);  
i32 __builtin_mips_extr_r_w (a64, i32);  
i32 __builtin_mips_extr_rs_w (a64, imm0_31);  
i32 __builtin_mips_extr_rs_w (a64, i32);
```



```
#-----
Ex:
a64 ac1 = 0x8123456712345678;
i32 shift_amount = 31;
i32 r1, r2, r3, r4, r5, r6;
r1 = __builtin_mips_extr_w (ac1, 1); // r1 will be 0x891A2B3C
r2 = __builtin_mips_extr_w (ac1, shift_amount); // r2 will be 0x02468ACE
r3 = __builtin_mips_extr_r_w (ac1, 4); // r3 will be 0x71234568
r4 = __builtin_mips_extr_r_w (ac1, shift_amount); // r4 will be 0x02468ACE
r5 = __builtin_mips_extr_rs_w (ac1, 4); // r5 will be 0x80000000
r6 = __builtin_mips_extr_rs_w (ac1, shift_amount); // r6 will be 0x80000000
```

### **Extract Half-word with Right Shift and Saturate**

```
i32 __builtin_mips_extr_s_h (a64, imm0_31);
i32 __builtin_mips_extr_s_h (a64, i32);
```

Note that the 16-bit result is sign-extended to a 32-bit result.

```
#-----
Ex:
a64 ac1 = 0xFFFFF81230000000;
i32 shift_amount = 4;
i32 r1, r2;
r1 = __builtin_mips_extr_s_h (ac1, 28); // r1 will be 0xFFFF8123
r2 = __builtin_mips_extr_s_h (ac1, shift_amount); // r2 will be 0xFFFF8000
```

### **Extract Bit from an Arbitrary Position**

```
i32 __builtin_mips_extp (a64, imm0_31);
i32 __builtin_mips_extp (a64, i32);
```

Note that the “imm0\_31” + 1 bits between “POS” and “POS” - “imm0\_31” are extracted and zero-extended to a 32-bit result. So, if X bits are extracted, X-1 should be used as the second parameter. The POS field can be set by using “\_\_builtin\_mips\_wrdsp”.

```
#-----
Ex:
a64 ac1 = 0x1234567887654321;
i32 r1, r2;
int the_size = 3;
__builtin_mips_wrdsp (35, 1); // Write 35 to the POS field
r1 = __builtin_mips_extp (ac1, 31); // r1 will be 0x88765432
r2 = __builtin_mips_extp (ac1, the_size); // r2 will be 0x8
```

### **Extract Bit from an Arbitrary Position and Decrement POS**

```
i32 __builtin_mips_extpdp (a64, imm0_31);
i32 __builtin_mips_extpdp (a64, i32);
```

Note that this intrinsic is the same as the ones in [Section](#), except that in addition, the POS field is decremented by the number of extracted bits, “imm0\_31” + 1 or “i32” + 1.

```
#-----
Ex:
a64 ac1 = 0x123456789ABCDEF0;
i32 r1, r2;
int the_size = 7;
__builtin_mips_wrdsp (35, 1); // Write 35 to the POS field
r1 = __builtin_mips_extpdp (ac1, 3); // r1 will be 0x8, and POS will be 31
```

```
r2 = __builtin_mips_extpdp (ac1, the_size); // r2 will be 0x9a, and POS will be 23
```

### **Shift an Accumulator Value**

```
a64 __builtin_mips_shilo (a64, imm_n32_31);  
a64 __builtin_mips_shilo (a64, i32);
```

Note that using the same a64 variable for the target and the first parameter could lead to better performance.

```
#-----  
Ex:  
a64 ac1 = 0x1234567887654321;  
int shift_amount = -8;  
ac1 = __builtin_mips_shilo (ac1, 8); // ac1 will be 0x0012345678876543  
ac1 = __builtin_mips_shilo (ac1, shift_amount); // ac1 will be 0x1234567887654300
```

### **Copy the LO to HI and a Value to LO and Increment POS by 32**

```
a64 __builtin_mips_mthlip (a64, i32);
```

Note that using the same a64 variable for the target and the first parameter could lead to better performance.

```
#-----  
Ex:  
a64 ac1 = 0x1234567887654321;  
int b = 0x11112222  
__builtin_mips_wrdsp (0, 1); // Write 0 to the POS field  
ac1 = __builtin_mips_mthlip (ac1, b); // ac1 will be 0x8765432111112222,  
// and POS will be 32
```

## **15.7.9.8 Using Intrinsics for 32-bit Integers**

### **Add and Set Carry/Add with Carry**

```
i32 __builtin_mips_addsc (i32, i32);  
i32 __builtin_mips_addwc (i32, i32);
```

Note that these two intrinsics can be used to add two 64-bit operands, each spread across two GPRs. The lower 32 bits are calculated first, then the carry from this addition is fed to the add of the upper 32 bit values.

```
#-----  
Ex:  
int i1 = 0;  
int i2 = 0xFFFFFFFF;  
int j1 = 1;  
int j2 = 1;  
int r1, r2;  
r2 = __builtin_mips_addsc (i2, j2); // r2 will be 0xFFFFFFFF+1 = 0 and C will be 1  
r1 = __builtin_mips_addwc (i1, j1); // r1 will be 0 + 1 + 1(C) = 2
```

### **Modular Subtraction on an Index Value**

```
i32 __builtin_mips_modsub (i32, i32);
```

Note that this intrinsic can be used to implement a circular buffer. The first parameter is the current index, that will be checked against zero. If the index is zero, the new index will be rolled back to the top of the buffer, assigned from the bit 23 to 8 of the second parameter. If the index is not zero, the new index will be decremented by the size of the element, assigned from the bit 7 to 0 of the second parameter.

```
#-----
```

```

Ex:
int index = 20;
int element = 0x1402;
while (1)
{
    index = __builtin_mips_modsub (index, element);
}
/* 'index' will be 20, 18, 16, ..., 4, 2, 0, 20, 18, 16, ... */

```

### **Bit Reverse a Half-word**

```

i32 __builtin_mips_bitrev (i32);
#-----
Ex:
int a = 0x1234; // 0001 0010 0011 0100
int r;
r = __builtin_mips_bitrev (a); // r will be 0x2c48 (0010 1100 0100 1000)

```

### **Insert Bit Field Variable**

```

i32 __builtin_mips_insv (i32, i32);

```

Note that using the same variable for the target and the first parameter could lead to better performance. This intrinsic inserts the value of the second parameter to the first parameter. The size to be extracted from the second parameter is specified in the SCOUNT field. The position to be inserted in the first parameter is specified in the POS field.

```

#-----
Ex:
int a = 0x12345678;
int r = 0xFFFFFFFF;
__builtin_mips_wrdspl ((16<<7)+4, 3); // set SCOUNT to 16, and set POS to 4
r = __builtin_mips_insv (r, a); // The lowest 16-bit value of a is inserted to r
// at bit 4. r will be 0xFFF5678F.

```

### **Load Unsigned Byte/Halfword/Word Indexed**

```

i32 __builtin_mips_lbx (void *, i32);
i32 __builtin_mips_lhx (void *, i32);
i32 __builtin_mips_lwx (void *, i32);

```

#### **NOTES:**

1. The first parameter is the base of the array, and the second parameter is the offset in byte.
2. The returned value is zero-extended for “\_\_builtin\_mips\_lbx”, and sign-extended for “\_\_builtin\_mips\_lhx” and “\_\_builtin\_mips\_lwx” to 32-bit integers.

```

#-----
Ex:
char array_a[100];
short array_b[100];
int array_c[100];
int offset = 20;
int r1, r2, r3;
r1 = __builtin_mips_lbx ((void *)array_a, offset);
r2 = __builtin_mips_lhx ((void *)array_b, offset);
r3 = __builtin_mips_lwx ((void *)array_c, offset);

```

### ***Signed Multiply and Add***

```
a64 __builtin_mips_madd (a64, i32, i32);
#-----
Ex:
a64 a = 1;
i32 b = 2;
i32 c = -3;
a = __builtin_mips_madd (a, b, c); // a will be 1 + 2 * (-3) = -5
```

### ***Unsigned Multiply and Add***

```
a64 __builtin_mips_maddu (a64, ui32, ui32);
#-----
a64 a = 1;
ui32 b = 2;
ui32 c = 3;
a = __builtin_mips_maddu (a, b, c); // a will be 1 + 2 * 3 = 7
```

### ***Signed Multiply and Subtract***

```
a64 __builtin_mips_msub (a64, i32, i32);
#-----
Ex:
a64 a = 1;
i32 b = 2;
i32 c = -3;
a = __builtin_mips_msub (a, b, c); // a will be 1 - 2 * (-3) = 7
```

### ***Unsigned Multiply and Subtract***

```
a64 __builtin_mips_msubu (a64, ui32, ui32);
#-----
Ex:
a64 a = 1;
ui32 b = 2;
ui32 c = 3;
a = __builtin_mips_msubu (a, b, c); // a will be 1 - 2 * 3 = -5
```

### ***Signed Multiply***

```
a64 __builtin_mips_mult (i32, i32);
#-----
a64 a;
i32 b = 2;
i32 c = -3;
a = __builtin_mips_mult (b, c); // a will be 2 * (-3) = -6
```

### ***Unsigned Multiply***

```
a64 __builtin_mips_multu (ui32, ui32);
#-----
Ex:
a64 a;
u32 b = 2;
u32 c = 3;
```

```
a = __builtin_mips_mulltu (b, c); // a will be 2 * 3 = 6
```

### **Left Shift and Append Bits**

```
i32 __builtin_mips_append (i32, i32, imm0_31); // DSPR2
#-----
Ex:
i32 a = 0x8765ABCD;
i32 b = 0x12345678;
i32 r;
r = __builtin_mips_append (a, b, 4); // r will be 0x765ABCD8
```

### **Byte Align Contents from Two Registers**

```
i32 __builtin_mips_balign (i32, i32, imm0_3); // DSPR2
#-----
Ex:
i32 a = 0x8765ABCD;
i32 b = 0x12345678;
i32 r;
r = __builtin_mips_balign (a, b, 3); // r will be 0xCD123456
```

### **Right Shift and Prepend Bits**

```
i32 __builtin_mips_prepend (i32, i32, imm0_31); // DSPR2
#-----
Ex:
i32 a = 0x8765ABCD;
i32 b = 0x12345678;
i32 r;
r = __builtin_mips_prepend (a, b, 4); // r will be 0x88765ABC
```

## **15.7.9.9 Using Intrinsics for 16-bit Integers**

### **Unsigned Add/Subtract with Optional Saturation**

```
v2i16 __builtin_mips_addu_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_addu_s_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_subu_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_subu_s_ph (v2i16, v2i16); // DSPR2
#-----
Ex:
v2i16 a = {0x0000, 0x8000};
v2i16 b = {0x8000, 0x8000};
v2i16 r1, r2, r3, r4;
r1 = __builtin_mips_addu_ph (a, b); // r1 will be {0x8000, 0x0000}
r2 = __builtin_mips_addu_s_ph (a, b); // r2 will be {0x8000, 0xFFFF}
r3 = __builtin_mips_subu_ph (a, b); // r3 will be {0x8000, 0x0000}
r4 = __builtin_mips_subu_s_ph (a, b); // r4 will be {0x0000, 0x0000}
```

### **Dot Product with Accumulate/Subtract**

```
a64 __builtin_mips_dpa_w_ph (a64, v2i16, v2i16); // DSPR2
a64 __builtin_mips_dps_w_ph (a64, v2i16, v2i16); // DSPR2
```

NOTES:

1. The result will be a 64-bit integer.
2. Using the same “a64” variable for both the target and the first parameter could result in better performance.

```
#-----
Ex:
v2i16 a = {0x0001, 0x8000};
v2i16 b = {0x0002, 0x8000};
a64 ac1, ac2;
ac1 = ac2 = 0;
ac1 = __builtin_mips_dpa_w_ph (ac1, a, b); // ac1 will be 0 + 1*2 +
                                           // 0x40000000 = 0x0000000040000002
ac2 = __builtin_mips_dps_w_ph (ac2, a, b); // ac2 will be 0 - 1*2 -
                                           // 0x40000000 = 0xFFFFFFFFFFFFFFFF
```

**Multiply with Optional Saturation**

```
v2i16 __builtin_mips_mul_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_mul_s_ph (v2i16, v2i16); // DSPR2
#-----
Ex:
v2i16 a = {0x7FFF, 0x8000};
v2i16 b = {0x7FFF, 0x8000};
v2i16 r1, r2;
r1 = __builtin_mips_mul_ph (a, b); // r1 will be {0x0001, 0x0000}
r2 = __builtin_mips_mul_s_ph (a, b); // r2 will be {0x7FFF, 0x7FFF}
```

**Multiply and Subtract and Accumulate**

```
a64 __builtin_mips_mulsa_w_ph (a64, v2i16, v2i16); // DSPR2
```

NOTES:

1. The result will be a 64-bit integer.
2. The processor endianness affects the format of the result.
3. Using the same “a64” variable for the target and the first parameter could lead to better performance.

```
#-----
Ex:
v2i16 a = {0x0001, 0x8000};
v2i16 b = {0x0002, 0x8000};
a64 ac1 = 0;
ac1 = __builtin_mips_mulsa_w_ph (ac1, a, b); // ac1 will be 0 + 1*2 -
                                           // 0x40000000 = 0xFFFFFFFFFC000002
```

**Shift Right Logical**

```
v2i16 __builtin_mips_shrl_ph (v2i16, imm0_15); // DSPR2
v2i16 __builtin_mips_shrl_ph (v2i16, i32); // DSPR2
#-----
Ex:
v2i16 a = {0x8000, 0x4000};
v2i16 r1, r2;
int shift_amount = 4;
r1 = __builtin_mips_shrl_ph (a, 4); // r1 will {0x0800, 0x0400};
r2 = __builtin_mips_shrl_ph (a, shift_amount); // r2 will {0x0800, 0x0400};
```

### **Cross Dot Product with Accumulate/Subtract**

```
a64 __builtin_mips_dpax_w_ph (a64, v2i16, v2i16); // DSPR2
a64 __builtin_mips_dpsx_w_ph (a64, v2i16, v2i16); // DSPR2
```

#### NOTES:

1. The result will be a 64-bit integer.
2. Using the same “a64” variable for both the target and the first parameter could result in better performance.

```
#-----
Ex:
v2i16 a = {0x0001, 0x0003};
v2i16 b = {0x0002, 0x0004};
a64 ac1, ac2;
ac1 = ac2 = 0;
ac1 = __builtin_mips_dpax_w_ph (ac1, a, b); // ac1 will be
// 0 + 1*4 + 2*3 = 0x000000000000000A
ac2 = __builtin_mips_dpsx_w_ph (ac1, a, b); // ac2 will be
// 0 - 1*4 - 2*3 = 0xFFFFFFFFFFFFFFF6
```

### **15.7.9.10 Using Intrinsics for Mixed Data Types: 16-bit and 32-bit Integers**

#### **Precision Reduce Two Integer Words to Halfwords After a Right Shift with Optional Rounding**

```
v2i16 __builtin_mips_preocr_sra_ph_w (i32, i32, imm0_31); // DSPR2
v2i16 __builtin_mips_preocr_sra_r_ph_w (i32, i32, imm0_31); // DSPR2
#-----
Ex:
i32 a = 0x80000000;
i32 b = 0x7FFFFFFF;
v2i16 r1, r2;
r1 = __builtin_mips_preocr_sra_ph_w (a, b, 4); // r1 will be {0x0000, 0xFFFF}
r2 = __builtin_mips_preocr_sra_r_ph_w (a, b, 4); // r2 will be {0x0000, 0x0000}
```

## **15.8 Code Optimizations for the MIPS DSP-R2**

This section provides information on writing efficient C programs using the MIPS32 DSP-R2 ASE. Note that the level of optimization is dependent on the specifics of how the particular GCC compiler works.

Programmers must select proper optimization levels to compile C code to suit their purposes. For example, for maximum speed: “-O3 -funroll-loops”. For good speed with moderate code sizes: “-O2”. For minimum code sizes: “-Os”. Note that, to allow the GCC compiler to efficiently schedule instructions based on the latency information, programmers must supply correct architecture and CPU options.

### **15.8.1 The Use of Intrinsics Versus ASM Macros**

The assembler has no knowledge of the pipeline and any code written using ASM macros will be treated as a single cycle latency instruction by the GCC compiler. This can lead to poor code scheduling and a lot of stalls in the resulting execution. On the other hand, the GCC compiler has knowledge of the pipeline latency of instructions and can schedule the DSP-R2 instructions correctly when programmers use intrinsics, that is “**\_\_builtin\_mips\_\***”. Hence it is important to try to avoid the use of ASM macros whenever possible.

## 15.8.2 Using Accumulators

To access only HI or LO of an accumulator, programmers are recommended to use a union type as follows.


```
typedef union
{
    long long a;    // One 64-bit accumulator
    int b[2];      // 32-bit HI and LO
} a64_union;
```

Note that CPU endianness affects how to access the accumulator as shown in [Figure 15.4](#). To access HI, b[0] is used in a Big Endian CPU, but b[1] is used in a Little Endian CPU. To access LO, b[1] is used in a Big Endian CPU, but b[0] is used in a Little Endian CPU.

```
Ex:
int test10 (long long a, v2q15 b, v2q15 c)
{
    a64_union temp;
    temp.a = __builtin_mips_dpaq_s_w_ph (a, b, c);
    return temp.b[0]; // Assume in a little-endian CPU we want to access LO.
}

# Generated Assembly
test10:
    mtlo    $4
    mthi    $5
    dpaq_s.w.ph    $ac0, $6, $7
    j       $31
    mflo    $2
```

**Figure 15.4 Accessing HI and LO of Accumulators**



## 15.8.3 Multiply “32-bit \* 32-bit = 64-bit”

To multiply 32 bits by 32 bits to obtain a 64-bit result, we must cast the 32-bit integer to a 64-bit integer (long long) and then perform the multiplication operation as follows.

```
Ex:
long long test11 (int a, int b)
{
    return (long long) a * b; // Same as (long long) a * (long long) b
                             // NOT the same as (long long) (a * b)
}

# Generated Assembly
test11:
```



```

mult    $4,$5
mflo    $2
j       $31
mfhi    $3

```

Combined with [Section 15.8.2](#) we can multiply 32-bit by 32-bit integers and get the highest 32-bit result from HI as follows.

```

Ex:
int test12 (int a, int b)
{
    a64_union temp;
    temp.a = (long long) a * b;
    return temp.b[1]; // Assume a little-endian CPU
}

# Generated Assembly
test12:
    mult    $4,$5
    j       $31
    mfhi    $2

```

### 15.8.4 Multiply and Add “32-bit \* 32-bit + 64-bit = 64-bit”

To perform multiplication and addition, we must cast the 32-bit integer to 64-bit (long long) and then perform multiplication and addition as follows.

```

Ex:
long long test13 (int a, int b, long long c)
{
    return c + (long long) a * b;
}

# Generated Assembly
test13:
    mtlo    $6
    mthi    $7
    madd    $4, $5
    mflo    $2
    j       $31
    mfhi    $3

```

### 15.8.5 Array Alignment and Data Layout

The GCC compiler provides a mechanism to specify the alignment of variables by using “`__attribute__ ((aligned (bytes)))`”. The alignment is important to loading or storing SIMD variables: “v4i8” and “v2q15”. If an array is aligned to a 4-byte boundary, that is, word-aligned, the GCC compiler can load or store four 8-bit data for v4i8 variables (or two 16-bit data for v2q15 variables) at a time using the load word class of instructions. The following example shows that when a char array A is aligned to a 4-byte boundary, we can cast this array to a v4i8 array and load four items to a v4i8 variable at a time by using the “lwx” instruction. However, if this char array A is not aligned to a 4-byte boundary, executing the following code will result in an address exception due to a mis-aligned load.

```

Ex: /* v4i8 Example */
char A[128] __attribute__ ((aligned (4)));
v4i8 test14 (int i)
{

```

```

v4i8 a;
v4i8 *myA = (v4i8 *)A;
a = myA[i];
return a;
}
# Generated Assembly
test14:
    lui    $2,%hi(A)
    sll    $4,$4,2
    addiu  $2,$2,%lo(A)
    j      $31
    lwx   $2,$2($4)

```

After SIMD data is loaded from memory into a register, it is best if the SIMD variables in the register are ready for use without requiring any rearrangement of the data. To avoid such data rearrangement which can reduce the benefit of parallelism, programmers must design their arrays with efficient data layout that is favorable for SIMD calculations.

## 15.8.6 GP-Relative Addressing

The GCC compiler provides an option “-G *num*” to put global or static data that is at most “*num*” bytes to the small data or BSS sections. This allows using only one instruction to access data via GP-relative addressing to improve the performance. Programmers can try to increase “*num*” to include more data into small data or BSS sections until these sections are full. Note that all ASEs should be compiled with the same “-G *num*”. The following example shows how the GCC compiler accesses a 16-byte array. When compiling the example with “-G 4”, calculating the base address of the array “C” needs two instructions: “lui \$3,%hi(C)” and “addiu \$3,\$3,%lo(C)”. But, when compiling with “-G 16” to put the whole array of “C” into the small data section, only one instruction “addiu \$3,\$28,%gp\_rel(C)” is required to get the base address of “C”.

```

Ex:
int C[4];
void test15 (int index, int value)
{
    C[index] = value;
}

# Generated Assembly when compiling with -G 4
test15:
    lui    $3,%hi(C)
    addiu  $3,$3,%lo(C)
    sll    $4,$4,2
    addu   $4,$4,$3
    j      $31
    sw     $5,0($4)
# -----

# Generated Assembly when compiling with -G 16
test15:
    addiu  $3,$28,%gp_rel(C)
    sll    $4,$4,2
    addu   $4,$4,$3
    j      $31
    sw     $5,0($4)

```

## 15.8.7 Fixed Registers and Register Variables

Register usage is defined by the Application Binary Interface (ABI). For example, the ABI defines that some registers are caller-saved, some are callee-saved, and a few registers are fixed (or called global) and not saved at all. When conforming to the ABI, functions are guaranteed to work with each other.

However, in very special cases where performance may be very critical, programmers may want to improve performance by avoiding the saving and restoring of registers and hence violating the ABI convention. This undertaking should be taken with caution and not normally recommended as general practice. The GCC compiler allows programmers to treat a register as fixed by using the command-line option: “-ffixed-*reg*” where *reg* must be the name of a register. When a register is fixed, the register allocation process does not touch the fixed register.

For example, the ABI defines that four accumulators (\$ac0 - \$ac3) are caller-saved registers, but programmers may want to dedicate one accumulator, \$ac1, for a special purpose. Note that because \$ac1 is a 64-bit register that consists of \$ac1hi and \$ac1lo, “-ffixed-*\$ac1hi* -ffixed-*\$ac1lo*” is specified in the command-line options to fix HI and LO of \$ac1.

The following example demonstrates that under the original ABI, the GCC compiler register allocator will allocate 64-bit variables to all accumulators. However, when \$ac1 is specified to be fixed, The GCC compiler only allocates 64-bit variables to \$ac0, \$ac2, and \$ac3.

```
Ex:
typedef long long a64;
typedef short v2q15 __attribute__((vector_size(4)));
void test16 (a64 a[4], v2q15 b[4], v2q15 c[4])
{
    a[0] = __builtin_mips_dpaq_s_w_ph (a[0], b[0], c[0]);
    a[1] = __builtin_mips_dpaq_s_w_ph (a[1], b[1], c[1]);
    a[2] = __builtin_mips_dpaq_s_w_ph (a[2], b[2], c[2]);
    a[3] = __builtin_mips_dpaq_s_w_ph (a[3], b[3], c[3]);
}

# Generated Assembly without using "-ffixed-$ac1hi --fixed-$ac1lo"
# Note that $ac0, $ac1, $ac2, and $ac3 are all used.
test16:
    lw     $15,4($4)
    lw     $14,0($4)
    lw     $13,12($4)
    lw     $10,8($4)
    lw     $9,20($4)
    lw     $8,16($4)
    lw     $3,28($4)
    lw     $7,24($4)
    lw     $2,0($6)
    lw     $12,12($5)
    lw     $11,12($6)
    mtlo   $15,$ac1
    mthi   $14,$ac1
    mtlo   $13,$ac2
    lw     $25,0($5)
    lw     $15,4($5)
    lw     $24,4($6)
    lw     $13,8($5)
    lw     $14,8($6)
    mthi   $10,$ac2
    mtlo   $9,$ac3
```

```

mthi    $8,$ac3
mtlo    $3
mthi    $7
dpaq_s.w.ph    $ac1,$25,$2
dpaq_s.w.ph    $ac2,$15,$24
dpaq_s.w.ph    $ac3,$13,$14
dpaq_s.w.ph    $ac0,$12,$11
mflo    $10
mfhi    $9
mflo    $8,$ac1
mfhi    $7,$ac1
mflo    $6,$ac2
mfhi    $5,$ac2
mflo    $3,$ac3
mfhi    $2,$ac3
sw      $10,28($4)
sw      $9,24($4)
sw      $8,4($4)
sw      $7,0($4)
sw      $6,12($4)
sw      $5,8($4)
sw      $3,20($4)
j        $31
sw      $2,16($4)
# -----

# Generated Assembly when using "-ffixed-\$ac1hi --fixed-\$ac1lo"
# Note that $ac0, $ac2, and $ac3 are used. But, $ac1 is not touched at all by the
# compiler.
test16:
lw      $3,4($4)
lw      $2,0($4)
lw      $25,12($4)
lw      $24,8($4)
lw      $9,20($4)
lw      $8,16($4)
lw      $15,12($5)
lw      $13,0($5)
lw      $11,0($6)
lw      $12,4($5)
lw      $7,4($6)
lw      $10,8($5)
lw      $5,8($6)
mtlo    $3,$ac2
mthi    $2,$ac2
lw      $3,28($4)
lw      $2,24($4)
mtlo    $25,$ac3
mthi    $24,$ac3
mtlo    $9
mthi    $8
dpaq_s.w.ph    $ac2,$13,$11
lw      $14,12($6)
dpaq_s.w.ph    $ac3,$12,$7
dpaq_s.w.ph    $ac0,$10,$5
mflo    $9
mfhi    $8
mtlo    $3

```

```

mthi    $2
dpaq_s.w.ph    $ac0,$15,$14
mflo    $25
mfhi    $24
mflo    $6,$ac2
mfhi    $5,$ac2
mflo    $3,$ac3
mfhi    $2,$ac3
sw      $25,28($4)
sw      $24,24($4)
sw      $6,4($4)
sw      $5,0($4)
sw      $3,12($4)
sw      $2,8($4)
sw      $9,20($4)
j       $31
sw      $8,16($4)

```

To use a fixed register, programmers must associate a register variable with the explicit name of the fixed register. For example, when \$a1 is fixed, we can declare “register a64 MYACC ASM (“\$ac1lo”)” in a Little Endian CPU, or “register a64 MYACC ASM (“\$ac1hi”)” in a Big Endian CPU. Then, the global variable “MYACC” is ready to be used across all functions via directly accessing \$a1.

The following example shows that when no global register variable is used, The GCC compiler needs to load or store a 64-bit global variable from or to memory.

```

Ex:
typedef long long a64;
typedef short v2q15 __attribute__((vector_size(4)));
a64 MYACC;
void test17 (v2q15 b, v2q15 c)
{
    MYACC = __builtin_mips_dpaq_s_w_ph (MYACC, b, c);
}

# Generated Assembly
test17:
    lw      $2,%gp_rel(MYACC)($28)
    lw      $3,%gp_rel(MYACC+4)($28)
    mtlo    $2
    mthi    $3
    dpaq_s.w.ph    $ac0,$4,$5
    mflo    $2
    mfhi    $3
    sw      $2,%gp_rel(MYACC)($28)
    j       $31
    sw      $3,%gp_rel(MYACC+4)($28)

```

However, when a register variable is used for a global variable, the overhead of storing and loading to and from memory is eliminated as follows, reducing the above 10 instructions to only 2.

```

Ex:
typedef long long a64;
typedef short v2q15 __attribute__((vector_size(4)));
register a64 MYACC asm ("$ac1lo"); /* Assume a little-endian CPU */
void test18 (v2q15 b, v2q15 c)
{

```

```

    MYACC = __builtin_mips_dpaq_s_w_ph (MYACC, b, c);
}

# Generated Assembly by
# "sde-gcc -mips32r2 -mdsp -O4 -S -ffixed-\$ac1hi --fixed-\$ac1lo -EL 18.c"
test18:
    j            $31
    dpaq_s.w.ph    $ac1,$4,$5

```

There are a few things to note when using the technique of fixing registers for global variables.

1. When fixing accumulators, because \$ac0 is the original HI and LO registers for multiplication and division instructions in MIPS32, \$ac0 cannot be fixed by using “-ffixed-hi -ffixed-lo”. The rest of the accumulators, that is \$ac1, \$ac2, and \$ac3 can be fixed.
2. When fixing \$ac1, \$ac2, or \$ac3, programmers must ensure that no third-party or library functions that can clobber \$ac1, \$ac2 or \$ac3 are called between accessing fixed accumulators. To practice safe programming methods, it is probably advisable to restrict the use of fixed accumulators inside an optimized kernel that consist of only internal functions.
3. The technique of fixing registers for use as global variables can be directly applied to callee-saved registers that are \$16 to \$23 (s0 to s7). Programmers do not need to change s0 to s7 to be fixed registers.

## 15.8.8 Conditional Moves

Typically conditional move instructions are used instead of branch instructions to avoid the penalty from branch delay slots and mis-predicted branches. For example, the GCC compiler can generate conditional move instructions for simple C code as follows.

```

Ex:
int test19 (int true_value, int false_value, int cond)
{
    if (cond)
        return true_value;
    else
        return false_value;
}

```

```

# Generated Assembly
test19:
    move    $2,$4
    j      $31
    movz   $2,$5,$6

```

# -----

```

Ex:
int test20 (int true_value, int false_value, int cond)
{
    return cond ? true_value : false_value;
}

```

```

# Generated Assembly
test20:
    move    $2,$4
    j      $31
    movz   $2,$5,$6

```

However, for complicated C code, the GCC compiler may not recognize the C patterns to generate conditional move instructions. Programmers can then use ASM macros to force the GCC compiler to use conditional move instructions. The following example shows how to use an ASM macro for a “movz” instruction. First, we need to assign a value “true\_value” (when the condition is true) to a resultant variable “result”. Then, we pass “result”, “false\_value” and “cond” to the ASM macro of “movz”. Note that the ASM macro uses “+d” for the output format, because the output register is also used as an input register.

```
Ex:
int test21 (int true_value, int false_value, int cond)
{
    int result = true_value;
    asm ("movz %0, %1, %2": "+d" (result): "d" (false_value), "d" (cond));
    return result;
}

# Generated Assembly
test21:
    #APP
        movz $4, $5, $6
    #NO_APP
        .set     noreorder
        .set     nomacro
        j       $31
        move    $2, $4
```

## 15.9 Programming Examples

This section describes the programming example for a 16-point FIR filter in three ways:

- FIR filter in traditional C code without SIMD variables and DSP-R2 intrinsics
- Hand-tuned assembly version
- FIR filter in efficient C code

### 15.9.1 The FIR Filter in Traditional C

The following C code implements a 16-point FIR filter without using SIMD variables and DSP-R2 intrinsics. The arrays of “coeffs” and “delay” store sixteen Q15 coefficients and sixteen Q15 delayed inputs.

```
Ex:
int i;
short x, y;
long long ac0 = 0;
for (i = 0; i < 16; i++)
{
    x = coeffs[i];
    y = delay[i];
    if ((unsigned short) x == 0x8000 && (unsigned short) y == 0x8000)
        ac0 += 0x7fffffff;
    else
        ac0 += ((x * y) << 1);
}
```

Inside a loop, a saturation condition needs to be detected when both values of “coeffs” and “delay” are 0x8000 (-1 in Q15). Moreover, to perform “Q15 x Q15” multiplication, a left shift is required after integer multiplication. This ver-

sion of the FIR filter takes 536 cycles to calculate one result. (The tools used for this experiment are SDE 6.03.00-rc3 and MIPSsim 4.6.23.) The traditional C code produces inefficient binary code, so DSP programmers would write it in assembly code.

## 15.9.2 The FIR Filter in Hand-Tuned Assembly

DSP programmers pack two coefficients to a register and pack two delayed inputs to a register, so a SIMD DSP-R2 instruction “dpaq\_s.w.ph” that performs saturation and “Q15 x Q15” multiplication can be applied efficiently. Instruction scheduling is performed by hand to avoid pipeline stalls. This FIR implementation can generate one result in 39 cycles which are much faster than the traditional C version in [Section 15.9.1](#). The hand-tuned assembly code for the FIR filter is as follows.

```
Ex:
mult   $0, $0
lw     $8, 0($5)
lw     $10, 0($6)
lw     $9, 4($5)
lw     $11, 4($6)
dpaq_s.w.ph $ac0, $8, $10
dpaq_s.w.ph $ac0, $9, $11
lw     $12, 8($5)
lw     $10, 8($6)
lw     $13, 12($5)
lw     $11, 12($6)
dpaq_s.w.ph $ac0, $12, $10
dpaq_s.w.ph $ac0, $13, $11
lw     $14, 16($5)
lw     $10, 16($6)
lw     $15, 20($5)
lw     $11, 20($6)
dpaq_s.w.ph $ac0, $14, $10
dpaq_s.w.ph $ac0, $15, $11
lw     $16, 24($5)
lw     $10, 24($6)
lw     $4, 28($5)
lw     $11, 28($6)
dpaq_s.w.ph $ac0, $16, $10
dpaq_s.w.ph $ac0, $4, $11
```

## 15.9.3 The FIR Filter in Efficient C

Although the hand-tuned assembly code yields good performance, it requires the programmer to manually do register allocation and code scheduling. A compromise is to write C code that uses SIMD variables and DSP-R2 intrinsics as shown.

```
Ex:
v2q15 *my_delay = (v2q15 *)delay;
v2q15 *my_coeffs = (v2q15 *)coeffs;
long long ac0 = 0;
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[0], my_coeffs[0]);
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[1], my_coeffs[1]);
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[2], my_coeffs[2]);
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[3], my_coeffs[3]);
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[4], my_coeffs[4]);
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[5], my_coeffs[5]);
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[6], my_coeffs[6]);
```



```
ac0 = __builtin_mips_dpaq_s_w_ph (ac0, my_delay[7], my_coeffs[7]);
```

This C code does not look as clean or as readable as the traditional C version in [Section 15.9.1](#), but it is efficient and calculates one result in 42 cycles which is only 7.69% slower than the hand-tuned assembly version in [Section 15.9.2](#). Compared to the hand-tuned assembly code, the efficient C code has three significant advantages as follows.

1. Register allocation is done by the compiler.
2. Code scheduling is done by the compiler.
3. Load and store of SIMD data is taken care of by the compiler.

Other DSP kernels can similarly benefit from C code.

## 15.10 MIPS32 DSP-R2 Intrinsics

This section lists the MIPS32 DSP-R2 intrinsics. Note that some parameters of intrinsics are immediate types. Programmers must pass a constant that is within the specific range in order to invoke these intrinsics.

### 15.10.1 Immediate Intrinsics

The immediate types are as follows:

```
imm0_3: the parameter must be a constant in the range 0 to 3.  
imm0_7: the parameter must be a constant in the range 0 to 7.  
imm0_15: the parameter must be a constant in the range 0 to 15.  
imm0_31: the parameter must be a constant in the range 0 to 31.  
imm0_63: the parameter must be a constant in the range 0 to 63.  
imm0_255: the parameter must be a constant in the range 0 to 255.  
imm_n512_511: the parameter must be a constant in the range -512 to 511.  
imm_n32_31: the parameter must be a constant in the range -32 to 31.
```

### 15.10.2 Intrinsics for DSPControl Register

```
void __builtin_mips_wrdsp (i32, imm0_63);  
i32 __builtin_mips_rddsp (imm0_63);  
i32 __builtin_mips_bposge32 ();
```

### 15.10.3 Intrinsics for Signed and Unsigned 8-bit Integers

```
v4i8 __builtin_mips_addu_qb (v4i8, v4i8);  
v4i8 __builtin_mips_addu_s_qb (v4i8, v4i8);  
v4i8 __builtin_mips_subu_qb (v4i8, v4i8);  
v4i8 __builtin_mips_subu_s_qb (v4i8, v4i8);  
i32 __builtin_mips_raddu_w_qb (v4i8);  
v4i8 __builtin_mips_shll_qb (v4i8, imm0_7);  
v4i8 __builtin_mips_shll_qb (v4i8, i32);  
v4i8 __builtin_mips_shrl_qb (v4i8, imm0_7);  
v4i8 __builtin_mips_shrl_qb (v4i8, i32);  
a64 __builtin_mips_dpau_h_qbl (a64, v4i8, v4i8);  
a64 __builtin_mips_dpau_h_qbr (a64, v4i8, v4i8);  
a64 __builtin_mips_dpsu_h_qbl (a64, v4i8, v4i8);  
a64 __builtin_mips_dpsu_h_qbr (a64, v4i8, v4i8);
```

```

v4i8 __builtin_mips_repl_qb (imm0_255);
v4i8 __builtin_mips_repl_qb (i32);
void __builtin_mips_cmpu_eq_qb (v4i8, v4i8);
void __builtin_mips_cmpu_lt_qb (v4i8, v4i8);
void __builtin_mips_cmpu_le_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgu_eq_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgu_lt_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgu_le_qb (v4i8, v4i8);
i32 __builtin_mips_cmpgdu_eq_qb (v4i8, v4i8); // DSPR2
i32 __builtin_mips_cmpgdu_lt_qb (v4i8, v4i8); // DSPR2
i32 __builtin_mips_cmpgdu_le_qb (v4i8, v4i8); // DSPR2
v4i8 __builtin_mips_pick_qb (v4i8, v4i8);
v4i8 __builtin_mips_absq_s_qb (v4i8); // DSPR2
v4i8 __builtin_mips_adduh_qb (v4i8, v4i8); // DSPR2
v4i8 __builtin_mips_adduh_r_qb (v4i8, v4i8); // DSPR2
v4i8 __builtin_mips_shra_qb (v4i8, imm0_7); // DSPR2
v4i8 __builtin_mips_shra_r_qb (v4i8, imm0_7); // DSPR2
v4i8 __builtin_mips_shra_qb (v4i8, i32); // DSPR2
v4i8 __builtin_mips_shra_r_qb (v4i8, i32); // DSPR2
v4i8 __builtin_mips_subuh_qb (v4i8, v4i8); // DSPR2
v4i8 __builtin_mips_subuh_r_qb (v4i8, v4i8); // DSPR2

```

## 15.10.4 Intrinsics for Q15

```

v2q15 __builtin_mips_addq_ph (v2q15, v2q15);
v2q15 __builtin_mips_addq_s_ph (v2q15, v2q15);
v2q15 __builtin_mips_subq_ph (v2q15, v2q15);
v2q15 __builtin_mips_subq_s_ph (v2q15, v2q15);
v2q15 __builtin_mips_absq_s_ph (v2q15);
v2q15 __builtin_mips_shll_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shll_ph (v2q15, i32);
v2q15 __builtin_mips_shll_s_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shll_s_ph (v2q15, i32);
v2q15 __builtin_mips_shra_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shra_ph (v2q15, i32);
v2q15 __builtin_mips_shra_r_ph (v2q15, imm0_15);
v2q15 __builtin_mips_shra_r_ph (v2q15, i32);
v2q15 __builtin_mips_mulq_rs_ph (v2q15, v2q15);
a64 __builtin_mips_dpaq_s_w_ph (a64, v2q15, v2q15);
a64 __builtin_mips_dpsq_s_w_ph (a64, v2q15, v2q15);
a64 __builtin_mips_mulsaq_s_w_ph (a64, v2q15, v2q15);
a64 __builtin_mips_maq_s_w_phl (a64, v2q15, v2q15);
a64 __builtin_mips_maq_s_w_phr (a64, v2q15, v2q15);
a64 __builtin_mips_maq_sa_w_phl (a64, v2q15, v2q15);
a64 __builtin_mips_maq_sa_w_phr (a64, v2q15, v2q15);
q31 __builtin_mips_muleq_s_w_phl (v2q15, v2q15);
q31 __builtin_mips_muleq_s_w_phr (v2q15, v2q15);
v2q15 __builtin_mips_repl_ph (imm_n512_511);
v2q15 __builtin_mips_repl_ph (i32);
void __builtin_mips_cmp_eq_ph (v2q15, v2q15);
void __builtin_mips_cmp_lt_ph (v2q15, v2q15);
void __builtin_mips_cmp_le_ph (v2q15, v2q15);
v2q15 __builtin_mips_pick_ph (v2q15, v2q15);
v2q15 __builtin_mips_packr1_ph (v2q15, v2q15);
v2q15 __builtin_mips_mulq_s_ph (v2q15, v2q15); // DSPR2
v2q15 __builtin_mips_addqh_ph (v2q15, v2q15); // DSPR2
v2q15 __builtin_mips_addqh_r_ph (v2q15, v2q15); // DSPR2

```

```

v2q15 __builtin_mips_subqh_ph (v2q15, v2q15); // DSPR2
v2q15 __builtin_mips_subqh_r_ph (v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpaqx_s_w_ph (a64, v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpaqx_sa_w_ph (a64, v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpsqx_s_w_ph (a64, v2q15, v2q15); // DSPR2
a64 __builtin_mips_dpsqx_sa_w_ph (a64, v2q15, v2q15); // DSPR2

```

### 15.10.5 Intrinsics for Q31

```

q31 __builtin_mips_addq_s_w (q31, q31);
q31 __builtin_mips_subq_s_w (q31, q31);
q31 __builtin_mips_absq_s_w (q31);
q31 __builtin_mips_shll_s_w (q31, imm0_31);
q31 __builtin_mips_shll_s_w (q31, i32);
q31 __builtin_mips_shra_r_w (q31, imm0_31);
q31 __builtin_mips_shra_r_w (q31, i32);
a64 __builtin_mips_dpaq_sa_l_w (a64, q31, q31);
a64 __builtin_mips_dpsq_sa_l_w (a64, q31, q31);
q31 __builtin_mips_mulq_rs_w (q31, q31); // DSPR2
q31 __builtin_mips_mulq_s_w (q31, q31); // DSPR2
q31 __builtin_mips_addqh_w (q31, q31); // DSPR2
q31 __builtin_mips_addqh_r_w (q31, q31); // DSPR2
q31 __builtin_mips_subqh_w (q31, q31); // DSPR2
q31 __builtin_mips_subqh_r_w (q31, q31); // DSPR2

```

### 15.10.6 Intrinsics for Mixed Data Types: 8-bit Integers and Q15/16-bit Integers

```

v4i8 __builtin_mips_preocrq_qb_ph (v2q15, v2q15);
v4i8 __builtin_mips_preocrqu_s_qb_ph (v2q15, v2q15);
v4i8 __builtin_mips_preocr_qb_ph (v2i16, v2i16); // DSPR2
v2q15 __builtin_mips_precequ_ph_qbl (v4i8);
v2q15 __builtin_mips_precequ_ph_qbr (v4i8);
v2q15 __builtin_mips_precequ_ph_qbla (v4i8);
v2q15 __builtin_mips_precequ_ph_qbra (v4i8);
v2q15 __builtin_mips_preceu_ph_qbl (v4i8);
v2q15 __builtin_mips_preceu_ph_qbr (v4i8);
v2q15 __builtin_mips_preceu_ph_qbla (v4i8);
v2q15 __builtin_mips_preceu_ph_qbra (v4i8);
v2q15 __builtin_mips_muleu_s_ph_qbl (v4i8, v2q15);
v2q15 __builtin_mips_muleu_s_ph_qbr (v4i8, v2q15);

```

### 15.10.7 Intrinsics for Mixed Data Types: Q15 and Q31

```

v2q15 __builtin_mips_preocrq_ph_w (q31, q31);
v2q15 __builtin_mips_preocrq_rs_ph_w (q31, q31);
q31 __builtin_mips_preceq_w_phl (v2q15);
q31 __builtin_mips_preceq_w_phr (v2q15);

```

### 15.10.8 Intrinsics for 64-bit Accumulators

```

i32 __builtin_mips_extr_w (a64, imm0_31);
i32 __builtin_mips_extr_w (a64, i32);
i32 __builtin_mips_extr_r_w (a64, imm0_31);
i32 __builtin_mips_extr_r_w (a64, i32);
i32 __builtin_mips_extr_rs_w (a64, imm0_31);

```

```

i32 __builtin_mips_extr_rs_w (a64, i32);
i32 __builtin_mips_extr_s_h (a64, imm0_31);
i32 __builtin_mips_extr_s_h (a64, i32);
i32 __builtin_mips_extp (a64, imm0_31);
i32 __builtin_mips_extp (a64, i32);
i32 __builtin_mips_extpdp (a64, imm0_31);
i32 __builtin_mips_extpdp (a64, i32);
a64 __builtin_mips_shilo (a64, imm_n32_31);
a64 __builtin_mips_shilo (a64, i32);
a64 __builtin_mips_mthlip (a64, i32);

```

### 15.10.9 Intrinsic for 32-bit Integers

```

i32 __builtin_mips_addsc (i32, i32);
i32 __builtin_mips_addwc (i32, i32);
i32 __builtin_mips_modsub (i32, i32);
i32 __builtin_mips_bitrev (i32);
i32 __builtin_mips_insv (i32, i32);
i32 __builtin_mips_lbx (void *, i32);
i32 __builtin_mips_lhx (void *, i32);
i32 __builtin_mips_lwx (void *, i32);
i32 __builtin_mips_append (i32, i32, imm0_31); // DSPR2
i32 __builtin_mips_balign (i32, i32, imm0_3); // DSPR2
i32 __builtin_mips_prepend (i32, i32, imm0_31); // DSPR2
a64 __builtin_mips_madd (a64, i32, i32);
a64 __builtin_mips_maddu (a64, ui32, ui32);
a64 __builtin_mips_msub (a64, i32, i32);
a64 __builtin_mips_msubu (a64, ui32, ui32);
a64 __builtin_mips_mult (i32, i32);
a64 __builtin_mips_multu (ui32, ui32);

```

### 15.10.10 Intrinsic for 16-bit Integers

```

v2i16 __builtin_mips_addu_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_addu_s_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_subu_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_subu_s_ph (v2i16, v2i16); // DSPR2
a64 __builtin_mips_dpa_w_ph (a64, v2i16, v2i16); // DSPR2
a64 __builtin_mips_dps_w_ph (a64, v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_mul_ph (v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_mul_s_ph (v2i16, v2i16); // DSPR2
a64 __builtin_mips_mulsa_w_ph (a64, v2i16, v2i16); // DSPR2
v2i16 __builtin_mips_shrl_ph (v2i16, imm0_15); // DSPR2
v2i16 __builtin_mips_shrl_ph (v2i16, i32); // DSPR2
a64 __builtin_mips_dpax_w_ph (a64, v2i16, v2i16); // DSPR2
a64 __builtin_mips_dpax_w_ph (a64, v2i16, v2i16); // DSPR2

```

### 15.10.11 Intrinsic for Mixed Data Types: 16-bit and 32-bit Integers

```

v2i16 __builtin_mips_preocr_sra_ph_w (i32, i32, imm0_31); // DSPR2
v2i16 __builtin_mips_preocr_sra_r_ph_w (i32, i32, imm0_31); // DSPR2

```

## 15.11 DSP-R2 ASE Instruction Groups

The following tables list the DSP-R2 instructions per function group. The input and output data type for each instruction is included, as well as the intended application. Refer to [Section 15.12, "Listing of DSP-R2 ASE Instruction Groups"](#) for an alphabetical listing of DSP-R2 instructions and associated links to the corresponding instruction. Refer to [Section "Repeat rate is measured as number of independent instructions that can be sent in 1 cycle."](#) for a definition and encoding of each individual DSP-R2 instruction.

[Table 15.4](#) through [Table 15.11](#) in this section list all the instructions in the DSP-R2 ASE. In each table below, the column entitled "Writes GPR / ac / DSPControl", indicates the explicit write performed by each instruction. This column indicates the writing of a field in the *DSPControl* register other than the *ouflag* field (which is written by a large number of instructions).

**Table 15.4 List of Instructions in MIPS® DSP-R2 ASE in Arithmetic Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
ADDQ.PH rd,rs,rt ADDQ_S.PH rd,rs,rt	Pair Q15	Pair Q15	GPR	VoIP SoftM	Element-wise addition of two vectors of Q15 fractional values, with optional saturation.
ADDQ_S.W rd,rs,rt	Q31	Q31	GPR	Audio	Add two Q31 fractional values with saturation.
ADDU.QB rd,rs,rt ADDU_S.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise addition of vectors of four unsigned byte values. Results may be optionally saturated to 255.
ADDUH.QB rd,rs,rt ADDUH_R.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise addition of vectors of four unsigned byte values, halving each result by right-shifting by one bit position. Results may be optionally rounded up in the least-significant bit.
ADDU.PH rd,rs,rt ADDU_S.PH rd,rs,rt	Pair Unsigned Halfword	Pair Unsigned Halfword	GPR	Video	Element-wise addition of vectors of two unsigned halfword values, with optional saturation on overflow.
ADDQH.PH rd,rs,rt ADDQH_R.PH rd,rs,rt	Pair Signed Halfword	Pair Signed Halfword	GPR	Misc	Element-wise addition of vectors of two signed halfword values, halving each result with right-shifting by one bit position. Results may be optionally rounded up in the least-significant bit.
ADDQH.W rd,rs,rt ADDQH_R.W rd,rs,rt	Signed Word	Signed Word	GPR	Misc	Add two signed word values, halving the result with right-shifting by one bit position. Result may be optionally rounded up in the least-significant bit.
SUBQ.PH rd,rs,rt SUBQ_S.PH rd,rs,rt	Pair Q15	Pair Q15	GPR	VoIP	Element-wise subtraction of two vectors of Q15 fractional values, with optional saturation.
SUBQ_S.W rd,rs,rt	Q31	Q31	GPR	Audio	Subtraction with Q31 fractional values, with saturation.

**Table 15.4 List of Instructions in MIPS® DSP-R2 ASE in Arithmetic Sub-class (*continued*)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
SUBU.QB rd,rs,rt SUBU_S.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise subtraction of unsigned byte values, with optional unsigned saturation.
SUBUH.QB rd,rs,rt SUBUH_R.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise subtraction of unsigned byte values, shifting the results right one bit position (halving). The results may be optionally rounded up by adding 1 to each result at the most-significant discarded bit position before shifting.
SUBU.PH rd,rs,rt SUBU_S.PH rd,rs,rt	Pair Unsigned Halfword	Pair Unsigned Halfword	GPR	Video	Element-wise subtraction of vectors of two unsigned halfword values, with optional saturation on overflow.
SUBQH.PH rd,rs,rt SUBQH_R.PH rd,rs,rt	Pair Signed Halfword	Pair Signed Halfword	GPR	Misc	Element-wise subtraction of vectors of two signed halfword values, halving each result with right-shifting by one bit position. Results may be optionally rounded up in the least-significant bit.
SUBQH.W rd,rs,rt SUBQH_R.W rd,rs,rt	Signed Word	Signed Word	GPR	Misc	Subtract two signed word values, halving the result with right-shifting by one bit position. Result may be optionally rounded up in the least-significant bit.
ADDSC rd,rs,rt	Signed Word	Signed Word	GPR & <i>DSPControl</i>	Audio	Add two signed words and set the carry bit in the <i>DSPControl</i> register.
ADDWC rd,rs,rt	Signed Word	Signed Word	GPR	Audio	Add two signed words with the carry bit from the <i>DSPControl</i> register.
MODSUB rd,rs,rt	Signed Word	Signed Word	GPR	Misc	Modulo addressing support: update a byte index into a circular buffer by subtracting a specified decrement (in bytes) from the index, resetting the index to a specified value if the subtraction results in underflow.
RADDU.W.QB rd,rs	Quad Unsigned Byte	Unsigned Word	GPR	Misc	Reduce (add together) the 4 unsigned byte values in <i>rs</i> , zero-extending the sum to 32 bits before writing to the destination register. For example, if all 4 input values are 0x80 (decimal 128), then the result in <i>rd</i> is 0x200 (decimal 512).
ABSQ_S.QB rd,rt	Quad Q7	Quad Q7	GPR	Misc	Find the absolute value of each of four Q7 fractional byte elements in the source register, saturating values of -1.0 to the maximum positive Q7 fractional value.
ABSQ_S.PH rd,rt	Pair Q15	Pair Q15	GPR	Misc	Find the absolute value of each of two Q15 fractional halfword elements in the source register, saturating values of -1.0 to the maximum positive Q15 fractional value.

**Table 15.4 List of Instructions in MIPS® DSP-R2 ASE in Arithmetic Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
ABSQ_S.W rd,rt	Q31	Q31	GPR	Misc	Find the absolute value of the Q31 fractional element in the source register, saturating the value -1.0 to the maximum positive Q31 fractional value.
PRECR.QB.PH rd,rs,rt	Two Pair Integer Halfwords	Four Integer Bytes	GPR	Misc	Reduce the precision of four signed integer halfword input values by discarding the eight most-significant bits from each to create four signed integer byte output values. The two halfword values from register <i>rs</i> are used to create the two left-most byte results, allowing an endian-agnostic implementation.
PRECRQ.QB.PH rd,rs,rt	2 Pair Q15	Quad Byte	GPR	Misc	Reduce the precision of four Q15 fractional input values by truncation to create four Q7 fractional output values. The two Q15 values from register <i>rs</i> are written to the two left-most byte results, allowing an endian-agnostic implementation.
PRECR_SRA.PH.W rt,rs,sa PRECR_SRA_R.PH.W rt,rs,sa	Two Integer Words	Pair Integer Halfword	GPR	Misc	Reduce the precision of two integer word values to create a pair of integer halfword values. Each word value is first shifted right arithmetically by <i>sa</i> bit positions, and optionally rounded up by adding 1 at the most-significant discard bit position. The 16 least-significant bits of each word are then written to the corresponding halfword elements of destination register <i>rt</i> .
PRECRQ.PH.W rd,rs,rt PRECRQ_RS.PH.W rd,rs,rt	2 Q31	Pair halfword	GPR	Misc	Reduce the precision of two Q31 fractional input values by truncation to create two Q15 fractional output values. The Q15 value obtained from register <i>rs</i> creates the left-most result, allowing an endian-agnostic implementation. Results may be optionally rounded up and saturated before being written to the destination.
PRECRQU_S.QB.PH rd,rs,rt	2 Pair Q15	Quad Unsigned Byte	GPR	Misc	Reduce the precision of four Q15 fractional values by saturating and truncating to create four unsigned byte values.
PRECEQ.W.PHL rd,rt PRECEQ.W.PHR rd,rt	Q15	Q31	GPR	Misc	Expand the precision of a Q15 fractional value to create a Q31 fractional value by adding 16 least-significant bits to the input value.
PRECEQU.PH.QBL rd,rt PRECEQU.PH.QBR rd,rt PRECEQU.PH.QBLA rd,rt PRECEQU.PH.QBRA rd,rt	Unsigned Byte	Q15	GPR	Video	Expand the precision of two unsigned byte values by prepending a sign bit and adding seven least-significant bits to each to create two Q15 fractional values.

**Table 15.4 List of Instructions in MIPS® DSP-R2 ASE in Arithmetic Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
PRECEU.PH.QBL rd,rt PRECEU.PH.QBR rd,rt PRECEU.PH.QBLA rd,rt PRECEU.PH.QBRA rd,rt	Unsigned Byte	Unsigned half-word	GPR	Video	Expand the precision of two unsigned byte values by adding eight least-significant bits to each to create two unsigned halfword values.

**Table 15.5 List of Instructions in MIPS® DSP ASE in GPR-Based Shift Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
SHLL.QB rd, rt, sa SHLLV.QB rd, rt, rs	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Misc	Element-wise left shift of eight signed bytes. Zeros are inserted into the bits emptied by the shift. The shift amount is specified by the three least-significant bits of <i>sa</i> or <i>rs</i> .
SHLL.PH rd, rt, sa SHLLV.PH rd, rt, rs SHLL_S.PH rd, rt, sa SHLLV_S.PH rd, rt, rs	Pair Signed halfword	Pair Signed halfword	GPR	Misc	Element-wise left shift of two signed halfwords, with optional saturation on overflow. Zeros are inserted into the bits emptied by the shift. The shift amount is specified by the four least-significant bits of <i>sa</i> or <i>rs</i> .
SHLL_S.W rd, rt, sa SHLLV_S.W rd, rt, rs	Signed Word	Signed Word	GPR	Misc	Left shift of a signed word, with saturation on overflow. Zeros are inserted into the bits emptied by the shift. The shift amount is specified by the five least-significant bits of <i>sa</i> or <i>rs</i> . Use the MIPS32 instructions SLL or SLLV for non-saturating shift operations.
SHRL.QB rd, rt, sa SHRLV.QB rd, rt, rs	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise logical right shift of four byte values. Zeros are inserted into the bits emptied by the shift. The shift amount is specified by the three least-significant bits of <i>sa</i> or <i>rs</i> .
SHRL.PH rd, rt, sa SHRLV.PH rd, rt, rs	Pair Half-words	Pair Half-words	GPR	Video	Element-wise logical right shift of two halfword values. Zeros are inserted into the bits emptied by the shift. The shift amount is specified by the four least-significant bits of <i>rs</i> or the <i>sa</i> argument.
SHRA.QB rd,rt,sa SHRA_R.QB rd,rt,sa SHRAV.QB rd,rt,rs SHRAV_R.QB rd,rt,rs	Quad Byte	Quad Byte	GPR	Misc	Element-wise arithmetic (sign preserving) right shift of four byte values. Optional rounding may be performed, adding 1 at the most-significant discard bit position. The shift amount is specified by the three least-significant bits of <i>rs</i> or by the argument <i>sa</i> .



**Table 15.5 List of Instructions in MIPS® DSP ASE in GPR-Based Shift Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
SHRA.PH rd, rt, sa SHRAV.PH rd, rt, rs SHRA_R.PH rd, rt, sa SHRAV_R.PH rd, rt, rs	Pair Signed halfword	Pair Signed halfword	GPR	Misc	Element-wise arithmetic (sign preserving) right shift of two halfword values. Optionally, rounding may be performed, adding 1 at the most-significant discard bit position. The shift amount is specified by the four least-significant bits of <i>rs</i> or by the argument <i>sa</i> .
SHRA_R.W rd, rt, sa SHRAV_R.W rd, rt, rs	Signed Word	Signed Word	GPR	Video	Arithmetic (sign preserving) right shift of a word value. Optionally, rounding may be performed, adding 1 at the most-significant discard bit position. The shift amount is specified by the five least-significant bits of <i>rs</i> or the argument <i>sa</i> .

**Table 15.6 List of Instructions in MIPS® DSP-R2 ASE in Multiply Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
MULEU_S.PH.QBL rd,rs,rt MULEU_S.PH.QBR rd,rs,rt	Pair Unsigned Byte, Pair Unsigned Halfword,	Pair Unsigned Halfword	GPR	Still Image	Element-wise multiplication of two unsigned byte values from register <i>rs</i> with two unsigned halfword values from register <i>rt</i> . Each 24-bit product is truncated to 16 bits, with saturation if the product exceeds 0xFFFF, and written to the corresponding element in the destination register.
MULQ_RS.PH rd,rs,rt	Pair Q15	Pair Q15	GPR	Misc	Element-wise multiplication of two Q15 fractional values to create two Q15 fractional results, with rounding and saturation. After multiplication, each 32-bit product is rounded up by adding 0x00008000, then truncated to create a Q15 fractional value that is written to the destination register. If both multiplicands are -1.0, the result is saturated to the maximum positive Q15 fractional value. To stay compliant with the base architecture, this instruction leaves the base <i>HI-LO</i> pair <b>UNPREDICTABLE</b> after the operation. The other DSP-R2 ASE accumulators <i>ac1-ac3</i> are untouched.

**Table 15.6 List of Instructions in MIPS® DSP-R2 ASE in Multiply Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
MULEQ_S.W.PHL rd,rs,rt MULEQ_S.W.PHR rd,rs,rt	Pair Q15	Q31	GPR	VoIP	Multiplication of two Q15 fractional values, shifting the product left by 1 bit to create a Q31 fractional result. If both multiplicands are -1.0 the result is saturated to the maximum positive Q31 value. To stay compliant with the base architecture, this instruction leaves the base <i>HI-LO</i> pair <b>UNPREDICTABLE</b> after the operation. The other DSP-R2 ASE accumulators <i>ac1-ac3</i> must be untouched.
DPAU.H.QBL DPAU.H.QBR	Pair Bytes	Halfword	Acc	Image	Dot-product accumulation. Two pairs of corresponding unsigned byte elements from source registers <i>rt</i> and <i>rs</i> are separately multiplied, and the two 16-bit products are then summed together. The summed products are then added to the accumulator.
DPSU.H.QBL DPSU.H.QBR	Pair Bytes	Halfword	Acc	Image	Dot-product subtraction. Two pairs of corresponding unsigned byte elements from source registers <i>rt</i> and <i>rs</i> are separately multiplied, and the two 16-bit products are then summed together. The summed products are then subtracted from the accumulator.
DPA.W.PH ac,rs,rt	Pair Signed Halfword	Pair Signed Halfword	ac	VoIP / SoftM	Dot-product accumulation. The two pairs of corresponding signed integer halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate integer word products. The products are then summed and accumulated into the specified accumulator.
DPAX.W.PH ac,rs,rt	Pair Signed Halfword	Doubleword	ac	VoIP	Dot-product with crossed operands and accumulation. The two crossed pairs of signed integer halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate integer word products. The products are then summed and accumulated into the specified accumulator.

**Table 15.6 List of Instructions in MIPS® DSP-R2 ASE in Multiply Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
DPAQ_S.W.PH ac,rs,rt	Pair Q15	Q32.31	ac	VoIP / SoftM	Dot-product accumulation. Two pairs of corresponding Q15 fractional values from source registers <i>rt</i> and <i>rs</i> are separately multiplied and left-shifted 1 bit to create two Q31 fractional products. For each product, if both multiplicands are equal to -1.0 the product is clamped to the maximum positive Q31 fractional value. The products are then summed, and the sum is then sign extended to the width of the accumulator and accumulated into the specified accumulator. This instruction may be used to compute the imaginary component of a 16-bit complex multiplication operation after first swapping the operands to place them in the correct order.
DPAQX_S.W.PH ac,rs,rt	Pair Signed Halfword	Q32.31	ac	VoIP	Dot-product with saturating fractional multiplication and using crossed operands, with a final accumulation. The two crossed pairs of signed fractional halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate fractional word products. The products are then summed and accumulated into the specified accumulator.
DPAQX_SA.W.PH ac,rs,rt	Pair Signed Halfword	Q32.31	ac	VoIP	Dot-product with saturating fractional multiplication and using crossed operands, with a final saturating accumulation. The two crossed pairs of signed fractional halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate fractional word products. The products are then summed and accumulated with saturation into the specified accumulator.
DPS.W.PH ac,rs,rt	Pair Signed Halfword	Doubleword	ac	VoIP / SoftM	Dot-product subtraction. The two pairs of corresponding signed integer halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate integer word products. The products are then summed and subtracted from the specified accumulator.

**Table 15.6 List of Instructions in MIPS® DSP-R2 ASE in Multiply Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
DPSX.W.PH ac,rs,rt	Pair Signed Halfword	Q32.31	ac	VoIP	Dot-product with crossed operands and subtraction. The two crossed pairs of signed integer halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate integer word products. The products are then summed and subtracted into the specified accumulator.
DPSQ_S.W.PH ac,rs,rt	Pair Q15	Q32.31	ac	VoIP / SoftM	Dot-product subtraction. Two pairs of corresponding Q15 fractional values from source registers <i>rt</i> and <i>rs</i> are separately multiplied and left-shifted 1 bit to create two Q31 fractional products. For each product, if both multiplicands are equal to -1.0 the product is clamped to the maximum positive Q31 fractional value. The products are then summed, and the sum is then sign extended to the width of the accumulator and subtracted from the specified accumulator. This instruction may be used to compute the imaginary component of a 16-bit complex multiplication operation after first swapping the operands to place them in the correct order.
DPSQX_S.W.PH ac,rs,rt	Pair Signed Halfword	Q32.31	ac	VoIP	Dot-product with saturating fractional multiplication and using crossed operands, with a final subtraction. The two crossed pairs of signed fractional halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate fractional word products. The products are then summed and subtracted from the specified accumulator.
DPSQX_SA.W.PH ac,rs,rt	Pair Signed Halfword	Q32.31	ac	VoIP	Dot-product with saturating fractional multiplication and using crossed operands, with a final saturating subtraction. The two crossed pairs of signed fractional halfword values from source registers <i>rt</i> and <i>rs</i> are separately multiplied to create two separate fractional word products. The products are then summed and subtracted with saturation into the specified accumulator.

**Table 15.6 List of Instructions in MIPS® DSP-R2 ASE in Multiply Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
MULSAQ_S.W.PH ac,rs,rt	Pair Q15	Q32.31	ac	SoftM	Complex multiplication step. Performs element-wise fractional multiplication of the two Q15 fractional values from registers <i>rt</i> and <i>rs</i> , subtracting one product from the other to create a Q31 fractional result that is added to accumulator <i>ac</i> . The intermediate products are saturated to the maximum positive Q31 fractional value if both multiplicands are equal to -1.0.
DPAQ_SA.L.W ac,rs,rt	Q31	Q63	ac	Audio	Fractional multiplication of two Q31 fractional values to produce a Q63 fractional product. If both multiplicands are -1.0 the product is saturated to the maximum positive Q63 fractional value. The product is then added to accumulator <i>ac</i> . If the addition results in overflow or underflow, the accumulator is saturated to the maximum positive or minimum negative value.
DPSQ_SA.L.W ac,rs,rt	Q31	Q63	ac	Audio	Fractional multiplication of two Q31 fractional values to produce a Q63 fractional product. If both multiplicands are -1.0 the product is saturated to the maximum positive Q63 fractional value. The product is then subtracted from accumulator <i>ac</i> . If the addition results in overflow or underflow, the accumulator is saturated to the maximum positive or minimum negative value.
MAQ_S.W.PHL ac,rs,rt MAQ_S.W.PHR ac,rs,rt	Q15	Q32.31	ac	SoftM	Fractional multiply-accumulate. The product of two Q15 fractional values is sign extended to the width of the accumulator and added to accumulator <i>ac</i> . The intermediate product is saturated to the maximum positive Q31 fractional value if both multiplicands are equal to -1.0.
MAQ_SA.W.PHL ac,rs,rt MAQ_SA.W.PHR ac,rs,rt	Q15	Q31	ac	speech	Fractional multiply-accumulate with saturation after accumulation. The product of two Q15 fractional values is sign extended to the width of the accumulator and added to accumulator <i>ac</i> . The intermediate product is saturated to the maximum positive Q31 fractional value if both multiplicands are equal to -1.0. If the accumulation results in overflow or underflow, the accumulator value is saturated to the maximum positive or minimum negative Q31 fractional value.

**Table 15.6 List of Instructions in MIPS® DSP-R2 ASE in Multiply Sub-class (continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
MUL.PH rd,rs,rt MUL_S.PH rd,rs,rt	Pair Signed Halfword	Pair Signed Halfword	GPR	speech	Element-wise multiplication of two vectors of signed integer halfwords, writing the 16 least-significant bits of each 32-bit product to the corresponding element of the destination register. Optional saturation clamps each 16-bit result to the maximum positive or minimum negative value if the product cannot be accurately represented in 16 bits.
MULQ_S.PH rd,rs,rt	Pair Q15	Pair Q15	GPR	speech	Element-wise multiplication of two vectors of Q15 fractional halfwords, writing the 16 most-significant bits of each Q31-format product to the corresponding element of the destination register. Each result is saturated to the maximum positive Q15 value if both multiplicands were equal to -1.0 (0x8000 hexadecimal).
MULQ_S.W rd,rs,rt	Q31	Q31	GPR	speech	Fractional multiplication of two Q31 format words to create a Q63 format result that is truncated by discarding the 32 least-significant bits before being written to the destination register. The result is saturated to the maximum positive Q31 value if both multiplicands were equal to -1.0 (0x80000000 hexadecimal).
MULQ_RS.W rd,rs,rt	Q31	Q31	GPR	speech	Multiplication of two Q31 fractional words to create a Q63-format intermediate product that is rounded up by adding a 1 at bit position 31. The 32 most-significant bits of the rounded result are then written to the destination register. If both multiplicands were equal to -1.0 (0x80000000 hexadecimal), rounding is not performed and the result is clamped to the maximum positive Q31 value before being written to the destination.
MULSA.W.PH ac,rs,rt	Pair Signed Halfword	Doubleword	ac	speech	Element-wise multiplication of two vectors of signed integer halfwords to create two 32-bit word intermediate results. The right intermediate result is subtracted from the left intermediate result, and the resulting sum is accumulated into the specified accumulator.
MADD, MADDU, MSUB, MSUBU, MULT, MULTU	Word	DoubleWord	ac	Misc	Allows these instructions to target accumulators <i>ac1</i> , <i>ac2</i> , and <i>ac3</i> (in addition to the original <i>ac0</i> destination).

**Table 15.7 List of Instructions in MIPS® DSP-R2 ASE in Bit Manipulation Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
BITREV rd,rt	Unsigned Word	Unsigned Word	GPR	Audio / FFT	Reverse the order of the 16 least-significant bits of register <i>rt</i> , writing the result to register <i>rd</i> . The 16 most-significant bits are set to zero.
INSV rt,rs	Unsigned Word	Unsigned Word	GPR	Misc	Like the Release 2 INS instruction, except that the 5 bits for <i>pos</i> and <i>size</i> values are obtained from the <i>DSPControl</i> register. <i>size</i> = <i>scount</i> [14:10], and <i>pos</i> = <i>pos</i> [20:16].
REPL.QB rd,imm REPLV.QB rd,rt	Byte	Quad Byte	GPR	Video / Misc	Replicate a signed byte value into the four byte elements of register <i>rd</i> . The byte value is given by the 8 least-significant bits of the specified 10-bit immediate constant or by the 8 least-significant bits of register <i>rt</i> .
REPL.PH rd,imm REPLV.PH rd,rt	Signed halfword	Pair Signed halfword	GPR	Misc	Replicate a signed halfword value into the two halfword elements of register <i>rd</i> . The halfword value is given by the 16 least-significant bits of register <i>rt</i> , or by the value of the 10-bit immediate constant, sign-extended to 16 bits.

**Table 15.8 List of Instructions in MIPS® DSP-R2 ASE in Compare-Pick Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
CMPU.EQ.QB rs,rt CMPU.LT.QB rs,rt CMPU.LE.QB rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	DSPControl	Video	Element-wise unsigned comparison of the four unsigned byte elements of <i>rs</i> and <i>rt</i> , recording the boolean comparison results to the four right-most bits in the <i>ccond</i> field of the <i>DSPControl</i> register.
CMPGDU.EQ.QB rd,rs,rt CMPGDU.LT.QB rd,rs,rt CMPGDU.LE.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR DSPControl	Video	Element-wise unsigned comparison of the four right-most unsigned byte elements of <i>rs</i> and <i>rt</i> , recording the boolean comparison results to the four least-significant bits of register <i>rd</i> and to the four right-most bits in the <i>ccond</i> field of the <i>DSPControl</i> register.
CMPGU.EQ.QB rd,rs,rt CMPGU.LT.QB rd,rs,rt CMPGU.LE.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise unsigned comparison of the four right-most unsigned byte elements of <i>rs</i> and <i>rt</i> , recording the boolean comparison results to the four least-significant bits of register <i>rd</i> .
CMP.EQ.PH rs,rt CMP.LT.PH rs,rt CMP.LE.PH rs,rt	Pair Signed halfword	Pair Signed halfword	DSPControl	Misc	Element-wise signed comparison of the two halfword elements of <i>rs</i> and <i>rt</i> , recording the boolean comparison results to the two right-most bits in the <i>ccond</i> field of the <i>DSPControl</i> register.

**Table 15.8 List of Instructions in MIPS® DSP-R2 ASE in Compare-Pick Sub-class(continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
PICK.QB rd,rs,rt	Quad Unsigned Byte	Quad Unsigned Byte	GPR	Video	Element-wise selection of unsigned bytes from the four bytes of registers <i>rs</i> and <i>rt</i> into the corresponding elements of register <i>rd</i> , based on the value of the four right-most bits of the <i>ccond</i> field in the <i>DSPControl</i> register. If the corresponding <i>ccond</i> bit is 1, the byte value is copied from register <i>rs</i> , otherwise it is copied from <i>rt</i> .
PICK.PH rd,rs,rt	Pair Signed halfword	Pair Signed halfword	GPR	Misc	Element-wise selection of signed halfwords from the two halfwords in registers <i>rs</i> and <i>rt</i> into the corresponding elements of register <i>rd</i> , based on the value of the two right-most bits of the <i>ccond</i> field in the <i>DSPControl</i> register. If the corresponding <i>ccond</i> bit is 1, the halfword value is copied from register <i>rs</i> , otherwise it is copied from <i>rt</i> .
APPEND rt,rs,sa	Two Words	Word	GPR	Misc	Shifts the 32-bit word in register <i>rt</i> left by <i>sa</i> bits, inserting the <i>sa</i> least-significant bits from register <i>rs</i> into the bit positions emptied by the shift. The 32-bit result is then written to register <i>rt</i> .
PREPEND rt,rs,sa	Two Words	Word	GPR	Misc	Shifts the 32-bit word in register <i>rt</i> right by <i>sa</i> bits, inserting the <i>sa</i> least-significant bits from register <i>rs</i> into the bit positions emptied by the shift. The 32-bit result is then written to register <i>rt</i> .
BALIGN rt,rs,bp	Two Words	Word	GPR	Misc	Packs <i>bp</i> bytes from register <i>rt</i> and $(4-bp)$ bytes from register <i>rs</i> into a 32-bit word and writes it to register <i>rt</i> .
PACKRL.PH rd,rs,rt	Pair Signed Halfwords	Pair Signed Halfword	GPR	Misc	Pack two halfwords taken from registers <i>rs</i> and <i>rt</i> into destination register <i>rd</i> .



**Table 15.9 List of Instructions in Accumulator and DSPControl Access Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
EXTR.W rt,ac,shift EXTR_R.W rt,ac,shift EXTR_RS.W rt,ac,shift	Q63	Q31	GPR	Misc	Extract a Q31 fractional value from the 32 least-significant bits of 64-bit accumulator <i>ac</i> . The accumulator value may be shifted right logically by <i>shift</i> bits prior to the extraction, and the extracted value may be optionally rounded or rounded and saturated before being written to register <i>rt</i> . The <i>shift</i> argument value ranges from 0 to 31. The optional rounding step adds 1 at the most-significant bit position discarded by the shift. The optional saturation clamps the extracted value to the maximum positive Q31 value if the rounding step results in overflow.
EXTR_S.H rt,ac,shift	Q63	Q15	GPR	Misc	Extract a Q15 fractional value from the 16 least-significant bits of 64-bit accumulator <i>ac</i> . The accumulator value may be shifted right logically by <i>shift</i> bits prior to the extraction, and the extracted value is saturated before being written to register <i>rt</i> . The <i>shift</i> argument value ranges from 0 to 31. The saturation clamps the extracted value to the maximum positive or minimum negative Q15 value if the shifted accumulator value cannot be represented accurately as a Q15 format value.
EXTRV_S.H rt,ac,rs	Q63	Q15	GPR	Misc	Extract a Q15 fractional value from the 16 least-significant bits of 64-bit accumulator <i>ac</i> . The accumulator value may be shifted right logically by <i>shift</i> bits prior to the extraction, and the extracted value is saturated before being written to register <i>rt</i> . The <i>shift</i> argument ranges from 0 to 31 and is given by the five least-significant bits of register <i>rs</i> . The saturation clamps the extracted value to the maximum positive or minimum negative Q15 value if the shifted accumulator value cannot be represented accurately as a Q15 format value.

**Table 15.9 List of Instructions in Accumulator and DSPControl Access Sub-class(continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
EXTRV.W <i>rt,ac,rs</i> EXTRV_R.W <i>rt,ac,rs</i> EXTRV_RS.W <i>rt,ac,rs</i>	Q63	Q31	GPR	Misc	Extract a Q31 fractional value from the 32 least-significant bits of 64-bit accumulator <i>ac</i> . The accumulator value may be shifted right logically by <i>shift</i> bits prior to the extraction, and the extracted value may be optionally rounded or rounded and saturated before being written to register <i>rt</i> . The <i>shift</i> argument value is provided by the five least-significant bits of <i>rs</i> and ranges from 0 to 31. The optional rounding step adds 1 at the most-significant bit position discarded by the shift. The optional saturation clamps the extracted value to the maximum positive Q31 value if the rounding step results in overflow.
EXTP <i>rt,ac,size</i> EXTPV <i>rt,ac,rs</i> EXTPDP <i>rt,ac,size</i> EXTPDPV <i>rt,ac,rs</i>	Unsigned DWord	Unsigned Word	GPR / DSPControl	Audio / Video	Extract a set of <i>size</i> +1 contiguous bits from accumulator <i>ac</i> , right-justifying and sign-extending the result to 32 bits before writing the result to register <i>rt</i> . The position of the left-most bit to extract is given by the value of the <i>pos</i> field in the <i>DSPControl</i> register. The number of bits (less one) to extract is provided either by the <i>size</i> immediate operand or by the five least-significant bits of <i>rs</i> . The EXTPDP and EXTPDPV instructions also decrement the <i>pos</i> field by <i>size</i> +1 to facilitate sequential bit field extraction operations.
SHILO <i>ac,shift</i> SHILOV <i>ac,rs</i>	Unsigned DWord	Unsigned DWord	ac	Misc	Shift accumulator <i>ac</i> left or right by the specified number of bits, writing the shifted value back to the accumulator. The signed shift argument is specified either by the immediate operand <i>shift</i> or by the six least-significant bits of register <i>rs</i> . A negative shift argument results in a right shift of up to 32 bits, and a positive shift argument results in a left shift of up to 31 bits.
MTHLIP <i>rs, ac</i>	Unsigned Word	Unsigned Word	ac / DSPControl	Audio / Video	Copy the <i>LO</i> register of the specified accumulator to the <i>HI</i> register, copy <i>rs</i> to <i>LO</i> , and increment the <i>pos</i> field in <i>DSPControl</i> by 32.
MFHI/MFLO/MTHI/MTLO	Unsigned Word	Unsigned Word	GPR/ac	Misc	Copy an unsigned word to or from the specified accumulator <i>HI</i> or <i>LO</i> register to the specified GPR.
WRDSP <i>rt,mask</i>	Unsigned Word	Unsigned Word	DSPControl	Misc	Overwrite specific fields in the <i>DSPControl</i> register using the corresponding bits from the specified GPR. Bits in the <i>mask</i> argument correspond to specific fields in <i>DSPControl</i> ; a value of 1 causes the corresponding <i>DSPControl</i> field to be overwritten using the corresponding bits in <i>rt</i> , otherwise the field is unchanged.

**Table 15.9 List of Instructions in Accumulator and DSPControl Access Sub-class(continued)**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
RDDSP rt,mask	Unsigned Word	Unsigned Word	GPR	Misc	Copy the values of specific fields in the <i>DSPControl</i> register to the specified GPR. Bits in the <i>mask</i> argument correspond to specific fields in <i>DSPControl</i> ; a value of 1 causes the corresponding <i>DSPControl</i> field to be copied to the corresponding bits in <i>rt</i> , otherwise the bits in <i>rt</i> are unchanged.

**Table 15.10 List of Instructions in MIPS® DSP-R2 ASE in Indexed-Load Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
LBUX rd,index(base)	-	Unsigned byte	GPR	Misc	Index byte load from address base+(index). Loads the byte in the low-order bits of the destination register and zero-extends the result.
LHX rd,index(base)	-	Signed halfword	GPR	Misc	Index halfword load from address base+(index). Loads the halfword in the low-order bits of the register and sign-extends the result.
LWX rd, index(base)	-	Signed Word	GPR	Misc	Indexed word load from address base+(index).

**Table 15.11 List of Instructions in MIPS® DSP-R2 ASE in Branch Sub-class**

Instruction Mnemonics	Input Data Type	Output Data Type	Writes GPR / ac / DSPControl	App	Description
BPOSGE32 offset	-	-	-	Audio / Video	Branch if the <i>pos</i> value is greater than or equal to integer 32.

## 15.12 Listing of DSP-R2 ASE Instruction Groups

Table 15.12 shows an alphabetical listing of the DSP-R2 instruction set, along with the associated instruction group, the page number location of the actual instruction. The actual instruction can be viewed by clicking on either the instruction or the page number reference in the table below. For the definition of each instruction, refer to Table 15.4 through Table 15.11 above.

**Table 15.12 Alphabetical Listing of DSP-R2 Instructions**

Instruction Name	Instruction Group
ABSQ_S.PH	Arithmetic

**Table 15.12 Alphabetical Listing of DSP-R2 Instructions (*continued*)**

<b>Instruction Name</b>	<b>Instruction Group</b>
ABSQ_S.W	Arithmetic
ADDQ_S.W	Arithmetic
ADDQH[_R].PH	Arithmetic
ADDQH[_R].W	Arithmetic
ADDSC	Arithmetic
ADDU[_S].PH	Arithmetic
ADDU[_S].QB	Arithmetic
ADDUH[_R].QB	Arithmetic
ADDWC	Arithmetic
APPEND	Compare-Pick
BALIGN	Compare-Pick
BPOSGE32	Branch
BITREV	Bit Manipulation
CMP.cond.PH	Compare-Pick
CMPGDU.cond.QB	Compare-Pick
CMPGU.cond.QB	Compare-Pick
CMPU.cond.QB	Compare-Pick
DPA.W.PH	Multiply
DPAQ_S.W.PH	Multiply
DPAQ_SA.L.W	Multiply
DPAQX_S.W.PH	Multiply
DPAQX_SA.W.PH	Multiply
DPAU.H.QBL	Multiply
DPAU.H.QBT	Multiply
DPAX.W.PH	Multiply
DPS.W.PH	Multiply
DPSQ_S.W.PH	Multiply
DPSQ_SA.L.W	Multiply
DPSQX_S.W.PH	Multiply
DPSQX_SA.W.PH	Multiply
DPSU.H.QBL	Multiply
DPSU.H.QBR	Multiply
DPSX.W.PH	Multiply
EXTPV	Accumulator/DSPControl Access
EXTPDP	Accumulator/DSPControl Access
EXTPDPV	Accumulator/DSPControl Access
EXTPV	Accumulator/DSPControl Access

**Table 15.12 Alphabetical Listing of DSP-R2 Instructions (continued)**

<b>Instruction Name</b>	<b>Instruction Group</b>
EXTR[_RS].W	Accumulator/DSPControl Access
EXTR_S.H	Accumulator/DSPControl Access
EXTRV[_RS].W	Accumulator/DSPControl Access
EXTRV_S.H	Accumulator/DSPControl Access
INSV	Bit Manipulation
LBUX	Indexed Load
LHX	Indexed Load
LWX	Indexed Load
MADD	Multiply
MADDU	Multiply
MAQ_S[A].W.PHL	Multiply
MAQ_S[A].W.PHR	Multiply
MFHI	Accumulator/DSPControl Access
MFLO	Accumulator/DSPControl Access
MODSUB	Arithmetic
MSUB	Multiply
MSUBU	Multiply
MTHI	Accumulator/DSPControl Access
MTHILIP	Accumulator/DSPControl Access
MTLO	Accumulator/DSPControl Access
MUL[_S].PH	Arithmetic
MULEQ_S.W.PHL	Arithmetic
MULEQ_S.W.PHR	Arithmetic
MULEU_S.PH.QBL	Arithmetic
MULEU_S.PH.QBR	Arithmetic
MULQ_RS.W	Arithmetic
MULQ_S.PH	Arithmetic
MULQ_RS.PH	Arithmetic
MULQ_S.W	Arithmetic
MULSA.W.PH	Arithmetic
MULSAQ_S.W.PH	Arithmetic
MULT	Arithmetic
MULTU	Arithmetic
PACKRL.PH	Compare-Pick
PICK.PH	Compare-Pick
PICK.QB	Compare-Pick
PRECEQ.W.PHL	Arithmetic

**Table 15.12 Alphabetical Listing of DSP-R2 Instructions (*continued*)**

<b>Instruction Name</b>	<b>Instruction Group</b>
PRECEQ.W.PHR	Arithmetic
PRECEQU.PH.QBL	Arithmetic
PRECEQU.PH.QBLA	Arithmetic
PRECEQU.PH.QBR	Arithmetic
PRECEQU.PH.QBRA	Arithmetic
PRECEU.PH.QBL	Arithmetic
PRECEU.PH.QBLA	Arithmetic
PRECEU.PH.QBR	Arithmetic
PRECEU.PH.QBRA	Arithmetic
PRECR.QB.PH	Arithmetic
PRECRQ.PH.W	Arithmetic
PRECRQ_RS.PH.W	Arithmetic
PRECRQU_S.QB.PH	Arithmetic
PRECR_SRA[_R].PH.W	Arithmetic
PREPEND	Compare-Pick
RADDU.W.QB	Arithmetic
RDDSP	Accumulator/DSPControl Access
REPL.PH	Bit Manipulation
REPL.QB	Bit Manipulation
REPLV.PH	Bit Manipulation
REPL.QB	Bit Manipulation
SHILO	Accumulator/DSPControl Access
SHILOV	Accumulator/DSPControl Access
SHLL[_S].PH	GPR-Based Shift
SHLL.QB	GPR-Based Shift
SHLLV.QB	GPR-Based Shift
SHLL_S.W	GPR-Based Shift
SHLLV[_S].PH	GPR-Based Shift
SHLLV_S.W	GPR-Based Shift
SHRA[_R].PH	GPR-Based Shift
SHRAV[_R].PH	GPR-Based Shift
SHRA[_R].QB	GPR-Based Shift
SHRAV[_R].QB	GPR-Based Shift
SHRAV[_R].W	GPR-Based Shift
SHRL.PH	GPR-Based Shift
SHRLV.PH	GPR-Based Shift
SHRL.QB	GPR-Based Shift

**Table 15.12 Alphabetical Listing of DSP-R2 Instructions (continued)**

<b>Instruction Name</b>	<b>Instruction Group</b>
SHRLV.QB	GPR-Based Shift
SUBQ[_S].PH	Arithmetic
SUBQ_S.W	Arithmetic
SUBQH[_R].PH	Arithmetic
SUBQH[_R].W	Arithmetic
SUBU[_S].PH	Arithmetic
SUBU[_S].QB	Arithmetic
SUBUH[_R].QB	Arithmetic
WRDSP	Accumulator/DSPControl Access

## 15.13 DSP Instruction Latencies and Repeat Rates

Latency is defined with respect to instruction pair, but for ease of documenting they are defined for the instruction. If the behavior per instruction differs from that of an instruction pair, this difference is mentioned in the Notes column. If the instruction does not load any general purpose register (GPR) then it is shown as not applicable (n/a).

Repeat rate is measured as number of independent instructions that can be sent in 1 cycle.

**Table 15.13 interAptiv DSP Instruction Latencies and Repeat Rates**

Instruction	Latency	Repeat Rate	Notes
ADDQ{ _S }.PH, ADDQ_S.W,	2	1	
ADDU{ _S }.PH, ADDU{ _S }.QB	2	1	
ADDUH{ _R }.QB,	2	1	
ADDQH{ _R }.PH, ADDQH{ _R }.W	2	1	
ADDSC, ADDWC	2	1	
SUBQ{ _S }.PH, SUB_S.W	2	1	
SUBU{ _S }.PH, SUBU{ _S }.QB	2	1	
SUBUH{ _R }.QB	2	1	
SUBQH{ _R }.PH, SUBQH{ _R }.W	2	1	
MODSUB, RADDU.W.QB	2	1	
ABSQ_S.QB, ABSQ_S.PH, ABSQ_S.W	2	1	
PRECR.QB.PH	2	1	
PRECRQ.QB.PH	2	1	
PRECR_SRA{ _R }.PH.W	2	1	
PRECRQ{ _RS }.PH.W	2	1	
PRECRQU_S.QB.PH	2	1	
PRECEQ.W.PHL, PRECEQ.W.PHR	2	1	
PRECEQU.PH.QBL{ A }, PRECEQU.PH.QBR{ A }	2	1	
PRECEU.PH.QBL{ A }, PRECEU.PH.QBR{ A }	2	1	
SHLL.QB, SHLLV.QB	2	1	
SHLL{ _S }.PH, SHLLV{ _S }.PH	2	1	
SHLL_S.W, SHLLV_S.W	2	1	
SHRL.QB, SHRLV.QB	2	1	
SHRL.PH, SHRLV.PH	2	1	
SHRA{ _R }.QB, SHRAV{ _R }.QB	2	1	
SHRA{ _R }.PH, SHRAV{ _R }.PH	2	1	
SHRA_R.W, SHRAV_R.W	2	1	
MULEU_S.PH.QBL, MULEU_S.PH.QBR	6	1	
MULQ_RS.PH	6	1	
MULEQ_S.W.PHL, MULEQ_S.W.PHR	6	1	



**Table 15.13 interAptiv DSP Instruction Latencies and Repeat Rates (continued)**

Instruction	Latency	Repeat Rate	Notes
DPAU.H.QBL, DPAU.H.QBR	6/1	1	DPA to MADD/DPA is 1 while MFHI/MFLO is 6.
DPSU.H.QBL, DPSU.H.QBR	6/1	1	
DPA.W.PH, DPAX.W.PH	6/1	1	
DPAQ_S.W.PH, DPAQX_S.W.PH, DPAQX_SA.W.PH	6/1	1	
DPS.W.PH, DPSX.W.PH	6/1	1	
DPSQ_S.W.PH, DPSQX_S.W.PH, DPSQX_SA.W.PH	6/1	1	
DPAQ_SA.L.W, DPSQ_SA.L.W	6/1	1	
MAQ_S{A}.W.PHL, MAQ_S{A}.W.PHR	6/1	1	
MADD, MADDU, MSUB, MSUBU, MULT, MULTU	6/1	1	
MULSAQ_S.W.PH	6	1	
MUL{S}.PH	6	1	
MULQ_S.PH, MULQ_S.W, MULQ_RS.W	6	1	
MULSA.W.PH	6	1	
BITREV	2	1	
INSV	2	1	
REPL{V}.QB, REPL{V}.PH	2	1	
CMPU.EQ.QB, CMPU.LT.QB, CMPU.LE.QB	2	1	
CMPGDU.EQ.QB, CMPGDU.LT.QB, CMPGDU.LE.QB	2	1	
CMPGU.EQ.QB, CMPGU.LT.QB, CMPGU.LE.QB	2	1	
CMPEQ.PH, CMP.LT.PH, CMP.LE.PH	2	1	
PICK.QB, PICK.PH	2	1	
APPEND, PREPEND	2	1	
BALIGN	2	1	
PACKRL.PH	2	1	
EXTR.W, EXTR_R.W, EXTR_RS.W	6	1	
EXTR_S.H, EXTRV_S.H	6	1	
EXTRV{R_RS}.W	6	1	
EXTP, EXTPV, EXTPDP, EXTPDPV	2	1	
SHILO, SHILOV	5	5	
MTHLIP	5/13	5	5 cycles for acc register, while 13 cycles for DSPCTL.POS register.
MFHI	2	1	
MFLO	2	1	
MTHI	5	5	
MTLO	5	5	
WRDSP	n/a	1	
RDDSP	2	1	
LBUX, LHX, LWX	4	1	Assuming L1 data cache hit.
BPOSGE32	n/a	1	



## EJTAG Debug Support

The EJTAG block provides a system debug facility for the device. The EJTAG functions are not normally controlled by the end user, but rather are controlled by a debugger. This chapter is meant to be read in conjunction with the MIPS EJTAG Specification that was included as part of the release.

An EJTAG debug block is present in all cores available from MIPS Technologies, Inc. It contains support for things like hardware and software breakpoints, hardware single-step, and a JTAG based debug TAP for debug probe connection.

This chapter is used for debug of the interAptiv core. For more information on the debugging of the Multiprocessing System, including the CM2 and CPC, refer to the next chapter entitled “Multi-CPU Debug”

This chapter contains the following sections:

- [Section 16.1 “Overview”](#)
- [Section 16.2 “Trace Funnel and Trace Types”](#)
- [Section 16.3 “Detecting Debug Mode”](#)
- [Section 16.4 “Ways of Entering Debug Mode”](#)
- [Section 16.5 “Exiting Debug Mode”](#)
- [Section 16.6 “EJTAG and PDTrace Revisions”](#)
- [Section 16.7 “Connection Options”](#)
- [Section 16.8 “Hardware Breakpoints”](#)
- [Section 16.9 “Debug Vector Addressing”](#)
- [Section 16.10 “Test Access Port \(TAP\)”](#)
- [Section 16.11 “PDTrace”](#)
- [Section 16.12 “PDtrace Cycle-by-Cycle Behavior”](#)
- [Section 16.13 “PC Sampling”](#)
- [Section 16.14 “EJTAG Registers”](#)
- [Section 16.15 “Fast Debug Channel”](#)
- [Section 16.16 “TCB Trigger Logic”](#)

## 16.1 Overview

The EJTAG debug logic in the interAptiv core is compliant with EJTAG Specification 5.0 and includes:

1. Standard core debug features
2. Optional hardware breakpoints
3. Standard Test Access Port (TAP) for a dedicated connection to a debug host
4. Optional PDtrace capability for program counter/data address/data value trace to On-chip memory or to Trace probe

EJTAG debug resources are often controlled via high level debugger commands. The following is a brief overview of some EJTAG features.

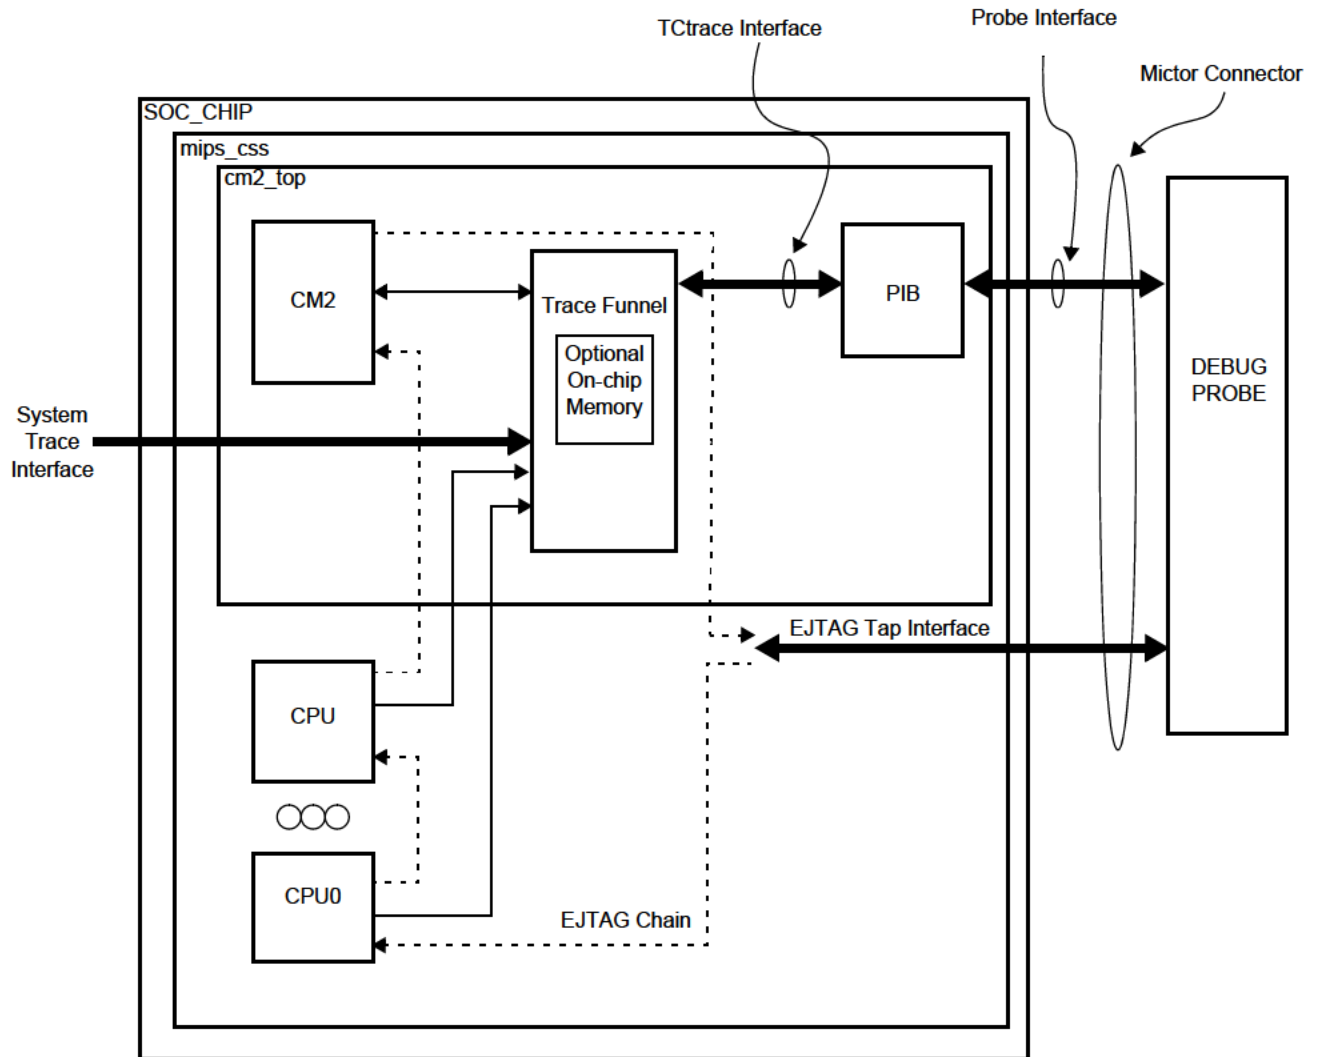
- **PCSAMPLE:** A feature allowing for non-intrusive reading of recently completed instruction addresses. The PCSAMPLE TAP instruction selects the TAP data register “PCSAMPLE” which contains an execution address and a flag indicating whether or not a new instruction has completed since the last read of the PCSAMPLE TAP data register.
- **EJTAG TAP:** The optional JTAG TAP associated with an EJTAG debug block used for communications with an EJTAG probe and debugger.
- **ECR (EJTAG Control Register):** This register is used mostly by probe developers and can only be accessed via a probe.
- **DCR (Debug Control Register):** This register is located in the drseg memory segment and can only be accessed in Debug mode.
- **DINT (Debug Interrupt):** an interrupt which causes a debug exception and entry into debug mode.
- **DRSEG (Debug Register Segment):** A memory overlay, present only while executing in debug mode, that allows access to registers controlling various EJTAG debug features.
- **DMSEG (Debug Memory Segment):** A memory overlay, present only while in debug mode and ECR.ProbEn is set, that an EJTAG probe emulates by satisfying processor accesses (fetches, loads, and stores.) The emulation is carried out via TAP data registers CONTROL, ADDRESS, and DATA.
- **Single-Step:** A debug setting that results in a debug exception after execution of a single non-debug mode instruction has completed.
- **Hardware Breakpoint:** A hardware resource capable of detecting execution or data access at virtual addresses.
- **Software Breakpoint:** The instruction “sdbbp” which causes a debug exception on execution. Debuggers will temporarily replace an instruction of your program with this instruction on setting a breakpoint in writeable memory.

## 16.2 Trace Funnel and Trace Types

The interAptiv Multiprocessing System implements a trace funnel that is used to communicate with the debug probe via the probe interface block. The trace funnel can accept trace information from either the CM2, the core, or the MIPS system trace.

The trace funnel and its connections are shown in [Figure 16.1](#). Refer to [Section 16.2.1 “Trace Types”](#) for more information on the types of traces shown.

**Figure 16.1 Trace Connections in the MIPS Debug Architecture**



## 16.2.1 Trace Types

The interAptiv Multiprocessing System supports the following trace types:

1. Core Trace
2. CM2 Trace
3. System Trace

**Core Trace** — Core trace allows CPU signals to be traced and routed to the trace funnel for processing. The functionality of core trace and the registers used to control it are described throughout this chapter.

**CM2 Trace** — The CM2 has its own trace and also manages the trace funnel. The functionality of CM2 trace and the registers used to control it are described in the CM2 chapter. Refer to 8 of this manual for more information.

**MIPS System Trace** — The MIPS System trace is a new feature to the interAptiv core and allows the SoC designer to place signals from their non-probe SoC logic directly into the trace funnel for PDTrace to capture. The logic and registers that controls System Trace are handled by the CM2. Refer to Chapter 8 of the interAptiv Hardware User’s Manual for more information on MIPS System Trace. For additional information, refer to Section 7.6.2 of the interAptiv Hardware User’s Manual.

## 16.2.2 EJTAG TAP Interface

Every TAP register access (also referred to as a “scan”) is a read-before-write operation. A TAP register access captures (reads) a register value from the target and then that value is serially shifted out to the tool as a new value is simultaneously shifted in. After all of the bits of the register have been shifted the input value is updated (written.)

There are two main paths through an EJTAG TAP state machine. One provides access to the single, 5-bit instruction register and the other provides access to the currently selected data register(s). Every TAP instruction access should result in the 5 bit binary value “00001” being read. Most EJTAG TAP instructions’ sole purpose is to select which data register is accessed during a data scan. EJTAG TAP instructions not intended to select specific TAP data registers will select the BYPASS data register.

In a multi-device target system, the term “scan chain” is used to describe the serial (daisy-chained) set of TAPS which are read/written in a single scan.

## 16.2.3 EJTAGBOOT vs NORMALBOOT

The EJTAGBOOT TAP instruction modifies the reset value of the *ECR.ProbTrap*, *ECR.ProbEn*, and *ECR.EjtagBrk*, thereby changing device reset behavior. Subsequent warm resets result in a debug exception after release from reset. Any EJTAG TAP reset will clear the EJTAGBOOT indication as will sending a NORMALBOOT TAP instruction.

## 16.3 Detecting Debug Mode

The DM bit of the CP0 Debug register (CP0 Register 23, Select 0) indicates if the processor is operating in debug mode. If this bit is set, the processor is operating in debug mode. This bit is set on any debug exception and is cleared by executing a *DERET* instruction. Refer to Chapter 2, CP0 Registers, for more information on the *Debug* register.

This bit is available to both probe and non-probe related configurations and can be read at any time. The user does not need to be in Debug mode in order to read this bit. This bit, along with the associated fields in this register, can be used by software to determine the conditions under which Debug mode was entered.

## 16.4 Ways of Entering Debug Mode

There are five ways to enter Debug mode. Each of these ways can be entered from either software, or from a debug probe. All of these ways cause the *DM* bit in the *CP0 Debug* register to be set.

1. EJTAG Debug Single Step
2. EJTAG Debug Interrupt. Caused by the assertion of the external EJ\_DINT input, or by setting the EJTAGBrk bit in the ECR register.
3. EJTAG debug hardware data breakpoint match
4. EJTAG debug hardware instruction breakpoint match
5. EJTAG Breakpoint (execution of SDBBP instruction)

### 16.4.1 EJTAG Debug Single Step

To enter Debug single step mode, the core must implement the single step mode. This can be determined by reading the *NoSST* bit (9) of the *CP0 Debug* register. If this bit is zero, the debug single step feature is implemented in the core. In the interAptiv core, this bit is always zero to indicate that the single step feature is implemented by the core.

Single step mode can be enabled or disabled by writing to the *SST* bit (8) of the *CP0 Debug* register. If the *SST* bit is set, the single step function is available once the core enters debug mode using any of the ways listed above. For implementation that include a probe, the common way is to generate the EJTAG DINT signal, which causes a debug interrupt to the core. For non-probe implementations, software can set the EJTAGBRK bit. Both of these methods are described in the following subsection.

### 16.4.2 EJTAG Debug Interrupt

The EJTAG DINT signal is an implementation dependent feature that is determined at build time. The *DINTsup* bit (24) in the *Implementation* register indicates whether the DINT signal is supported. This bit is written by the *EJ\_DINTsup* signal at reset depending on whether this option is selected at build time. This is a common way for probe or logic analyzer implementations to enter debug mode. Refer to [Section 16.14.4.5 “Implementation Register”](#) for more information.

Software can enter debug mode by setting the *EJTAGbrk* bit (12) or the *EJTAG Control* register. Setting this bit to 1 causes a debug exception to the processor, unless the CPU was in debug mode or another debug exception occurred. When the debug exception occurs, the processor core clock is restarted if the CPU was in low power mode. This bit is cleared by hardware when the debug exception is taken. Refer to [Section 16.14.4.6 “EJTAG Control Register”](#) for more information.

### 16.4.3 EJTAG Hardware Data Breakpoint Match

Data breakpoints occur on load/store transactions. Breakpoints are set on virtual address and ASID values, similar to the Instruction breakpoint. Data breakpoints can be set on a load, a store or both. Data breakpoints can also be set based on the value of the load/store operation. Finally, masks can be applied to both the virtual address and the

load/store value. Refer to [Section 16.8 “Hardware Breakpoints”](#) for more information and a list of registers used to set up a data breakpoint.

#### 16.4.4 EJTAG Hardware Instruction Breakpoint Match

Instruction breaks occur on instruction fetch operations and the break is set on the virtual address used by the instruction fetch unit. Instruction breaks can also be made on the ASID value used by the MMU. Finally, a mask can be applied to the virtual address to set breakpoints on a range of instructions. Instruction breakpoints compare the virtual address of the executed instructions (PC) and the ASID with the registers for each instruction breakpoint including masking of address and ASID. When an instruction breakpoint matches, a trigger is generated and a debug exception is optionally signalled. An internal bit in the instruction breakpoint registers is set to indicate that the match occurred.

Refer to [Section 16.8 “Hardware Breakpoints”](#) for more information and a list of register used to set up an instruction breakpoint.

#### 16.4.5 EJTAG Software Breakpoint

Software can execute a software debug breakpoint using the SDBBP instruction. When this instruction is executed, the debugger temporarily replaces the program instruction with the SDBBP instruction when setting a breakpoint in memory.

### 16.5 Exiting Debug Mode

As described above, there are five basic ways to enter debug mode. On in debug mode, the mode can only be exited in one of three ways:

- Execution of a Debug Exception Return (DERET) instruction.
- Reset the core
- Power cycle the core

During normal operation, exceptions are taken by the core and processed. Once the exception processing is complete, software executes an Exception Return (ERET) instruction. When in debug mode, software executes a Debug Exception Return (DERET) instruction. This causes the core to exit debug mode and return to previous mode as determined by the programmer (normal, kernel, supervisor, etc.).

Note that for a DERET instruction to be executed, the core must be in a state where it is fetching instructions. If for any reason the instruction stream has been halted and cannot resume, then the DERET instruction cannot be executed. In this case, the only other options are resetting the core, or cycling the power to the interAptiv core.

### 16.6 EJTAG and PDTrace Revisions

This chapter is intended to be used in conjunction with the EJTAG specification (MIPS document number MD00047) and the MIPS PDTrace specification (MIPS document number MD00439). These documents contain information for multiple types of MIPS cores, so the EJTAG and PDTrace versions of the core in question must be known in order to use these documents.

- ***EJTAG version with probe***: When using the MIPS Debug facility with a debug probe, the EJTAG version used in the interAptiv core can be determined by reading the EJTAGver field in bits 31:29 of the *Implementation* register. This is a TAP controller register that is only accessible through an EJTAG probe. The interAptiv core implements EJTAG revision 5.0. Refer to [Section 16.14.4.5 “Implementation Register”](#) for more information.



Note that the probe can read either the *Implementation* register of the CP0 *Debug* register described below to determine the EJTAG revision number.

- **EJTAG version without probe:** When using the MIPS Debug facility without a debug probe, the EJTAG version used in the interAptiv core can be determined by reading the *EJTAGver* field in bits 17:15 of the CP0 *Debug* register located at CP0 register 23, select 0. The interAptiv core implements EJTAG revision 5.0. Refer to Chapter 2 of this manual for more information on the CP0 *Debug* register. Note that the kernel can only read the CP0 *Debug* register to determine the EJTAG version and does not have access to the EJTAG *Implementation* register described above.
- **PDTrace version with probe:** When using the MIPS Debug facility with a debug probe, the PDTrace version used in the interAptiv core can be determined by reading the REV field in bits 3:0 of the *Trace Buffer Configuration* (TCBCONFIG) register located in the EJTAG TAP controller. Refer to the [Section 16.14.10.7 “TCBCONFIG Register \(Reg 0\)”](#) for more information on this register. The current revision is 3.0 as noted by the default value. Note that this register can only be read when an EJTAG probe is connected to the device.
- **PDTrace version without probe:** When using the MIPS Debug facility without a debug probe, the PDTrace version used in the interAptiv core can be determined by reading the REV field in bits 3:0 of the *Trace Buffer Configuration* (TCBCONFIG) register located in the EJTAG TAP controller.

However, since a probe is not attached in this case, the core must be in Debug mode in order to read this register. Debug mode can be entered in any of the ways described in [Section 16.4 “Ways of Entering Debug Mode”](#). Refer to the [Section 16.14.10.7 “TCBCONFIG Register \(Reg 0\)”](#) for more information on this register.

It should be noted that the *Device Identification* register located in [Section 16.14.4.4 on page 710](#) contains version and part number information. This register is only accessible when an EJTAG probe is attached, but does not provide EJTAG or PDTrace revision information. This register is used to by the manufacturer for their own device identification purposes and should not be used in an attempt to determine the EJTAG or PDTrace revisions.

## 16.7 Connection Options

The EJTAG debug port of the interAptiv core can be accessed either via a TAP (five JTAG pins), or the EJTAG debug block through the CP0 Debug register, the DCR, and drseg space. If the TAP is used, no ROM monitor is required and there is no interference with the customer's code. If there is no TAP, then the user must write their own ROM monitor.

There are two ways to connect to access the EJTAG debug facility:

- Software via the General Control Registers (GCR)
- Debug probe via the EJTAG Test Access Port (TAP)

The DCR (Debug Control Register) can be used to access the EJTAG debug port via software. This register is located in the drseg memory segment and can only be accessed in Debug mode. This register can be accessed by anyone that enters Debug mode and does not require that a probe be attached.

Access via software would mostly be performed during normal operation. As described in [Section 16.4 “Ways of Entering Debug Mode”](#) above, the CP0 Debug register (CP0 Register 23, Select 0) indicates whether or not the device is in Debug mode and the cause as to how it got there. Bit 30 of this register indicates if the core has entered Debug mode. If the core is not in Debug mode, the other bits have no meaning. If the core is in Debug mode, the other bits are used to provide additional information about how the device got into Debug mode. For example, setting a software breakpoint allows the core to enter Debug mode.

The ECR (EJTAG Control Register) is used mostly by probe developers and can only be accessed via a probe. Refer to [Section 16.14.4.6 “EJTAG Control Register”](#) for more information.

## 16.8 Hardware Breakpoints

Hardware breakpoints provide for the comparison by hardware of executed instructions and data load/store transactions. It is possible to set instruction breakpoints on addresses even in ROM area. Data breakpoints can be set to cause a debug exception on a specific data transaction. Instruction and data hardware breakpoints are alike for many aspects, and are thus described in parallel in the following. The term hardware is not applied to breakpoint, unless required to distinguish it from software breakpoint.

There are two types of simple hardware breakpoints implemented in the interAptiv core; Instruction breakpoints and Data breakpoints.

A core may be configured with the following breakpoint options:

- Zero, two, or four instruction breakpoints
- Zero, one, or two data breakpoints

### 16.8.1 Instruction Breakpoints

Instruction breaks occur on instruction fetch operations and the break is set on the virtual address used by the instruction fetch unit. Instruction breaks can also be made on the ASID value used by the TLB-based MMU. Finally, a mask can be applied to the virtual address to set breakpoints on a range of instructions.

Instruction breakpoints compare the virtual address of the executed instructions (PC) and the ASID with the registers for each instruction breakpoint including masking of address and ASID. When an instruction breakpoint matches, a trigger is generated and a debug exception is optionally signalled. An internal bit in the instruction breakpoint registers is set to indicate that the match occurred.

### 16.8.2 Data Breakpoints

Data breakpoints occur on load/store transactions. Breakpoints are set on virtual address and ASID values, similar to the Instruction breakpoint. Data breakpoints can be set on a load, a store or both. Data breakpoints can also be set based on the value of the load/store operation. Finally, masks can be applied to both the virtual address and the load/store value.

Data breakpoints compare the transaction type (TYPE), which may be load or store, the virtual address of the transaction (ADDR), the ASID, accessed bytes (BYTLANE) and data value (DATA), with the registers for each data breakpoint including masking or qualification on the transaction properties. When a data breakpoint matches, a trigger is generated and a debug exception is optionally signalled. An internal bit in the data breakpoint registers is set to indicate that the match occurred.

### 16.8.3 Instruction Breakpoint Registers Overview

The register with implementation indication and status for instruction breakpoints in general is shown in [Table 16.1](#).

**Table 16.1 Overview of Status Register for Instruction Breakpoints**

Register Mnemonic	Register Name and Description
<i>IBS</i>	Instruction Breakpoint Status

Up to four instruction breakpoints are available and are numbered 0 to 3 for registers and breakpoints, and the number is indicated by  $n$ . The registers for each breakpoint are shown in [Table 16.2](#).

**Table 16.2 Overview of Registers for Each Instruction Breakpoint**

Register Mnemonic	Register Name and Description
<i>IBAn</i>	Instruction Breakpoint Address $n$
<i>IBMn</i>	Instruction Breakpoint Address Mask $n$
<i>IBASIDn</i>	Instruction Breakpoint ASID $n$
<i>IBCn</i>	Instruction Breakpoint Control $n$

## 16.8.4 Data Breakpoint Registers Overview

The register with implementation indication and status for data breakpoints in general is shown in [Table 16.3](#).

**Table 16.3 Overview of Status Register for Data Breakpoints**

Register Mnemonic	Register Name and Description
<i>DBS</i>	Data Breakpoint Status

Up to two data breakpoints are available and are numbered 0 and 1 for registers and breakpoints, and the number is indicated by  $n$ . The registers for each breakpoint are shown in [Table 16.4](#).

**Table 16.4 Overview of Registers for Each Data Breakpoint**

Register Mnemonic	Register Name and Description
<i>DBAn</i>	Data Breakpoint Address $n$
<i>DBMn</i>	Data Breakpoint Address Mask $n$
<i>DBASIDn</i>	Data Breakpoint ASID $n$
<i>DBCn</i>	Data Breakpoint Control $n$
<i>DBVn</i>	Data Breakpoint Value $n$

## 16.8.5 Conditions for Matching Breakpoints

A number of conditions must be fulfilled in order for a breakpoint to match on an executed instruction or a data transaction, as described in this section. Breakpoints only match for instructions executed in non-debug mode, never on instructions executed in debug mode.

The match of an enabled breakpoint always generates a trigger indication and can also generate a debug exception. The *BE* and/or *TE* bits in the *IBCn* or *DBCn* registers are used to enable the breakpoints.

Debug software should not configure breakpoints to compare on an ASID value unless a TLB is present in the implementation.

### 16.8.5.1 Conditions for Matching Instruction Breakpoints

When an instruction breakpoint is enabled, that breakpoint is evaluated for the address of every executed instruction in non-debug mode, including execution of instructions at an address causing an address error on an instruction fetch. The breakpoint is not evaluated on instructions from a speculative fetch or execution, nor for addresses which are unaligned with an executed instruction.

A breakpoint match depends on the virtual address of the executed instruction (PC), which can be masked at the bit level. The match can also include an optional compare of the ASID and TC values. The registers for each instruction breakpoint contain the values and mask used in the compare, and the equation that determines the match is shown below in C-like notation.

```

IB_match =
    ( ! IBCnTCuse || ( TC == IBCTC ) ) &&
    ( ! IBCnASIDuse || ( ASID == IBASIDnASID ) ) &&
    ( <all 1's> == ( IBMnIBM | ~ ( PC ^ IBAnIBA ) ) &&
    ( ( IBMnISAM | ~ ( ISAMode ^ IBAnISA ) ) )

```

The match indication for instruction breakpoints is always precise, i.e., indicated on the instruction causing the `IB_match` to be true.

### 16.8.5.2 Conditions for Matching Data Breakpoints

When a data breakpoint is enabled, that breakpoint is evaluated for every data transaction due to a load/store instruction executed in non-debug mode, including coprocessor loads/stores and transactions causing an address error on data access. The breakpoint is not evaluated due to a PREF instruction or other transactions which are not part of explicit load/store transactions in the execution flow, nor for addresses which are not the explicit load/store source or destination address.

A breakpoint match depends on the transaction type (TYPE) as load or store, the address, and optionally the data value of a transaction. Match also includes an optional compare of the TC value. The registers for each data breakpoint contain the values and mask used in the compare, and the equation that determines the match is shown below in C-like notation.

The overall match equation is the `DB_match`.

```

DB_match =
    ( !DBCnTCuse || ( TC == DBCnTC ) ) &&
    ( ( ( TYPE == load ) && ! DBCnNoLB ) ||
    ( ( TYPE == store ) && ! DBCnNoSB ) ) &&
    DB_addr_match && ( DB_no_value_compare || DB_value_match )

```

The match on the address part, `DB_addr_match`, depends on the virtual address of the transaction (ADDR), the ASID value, and the accessed bytes (BYTELANE) where `BYTELANE[0]` is 1 only if the byte at bits [7:0] on the bus is accessed, and `BYTELANE[1]` is 1 only if the byte at bits [15:8] is accessed, etc. The `DB_addr_match` is shown below.

```

DB_addr_match =
    ( ! DBCnASIDuse || ( ASID == DBASIDnASID ) ) &&
    ( <all 1's> == ( DBMnDBM | ~ ( ADDR ^ DBAnDBA ) ) ) &&
    ( <all 0's> != ( ~ BAI & BYTELANE ) )

```

The size of `DBCnBAI` and `BYTELANE` is 8 bits. They are 8 bits to allow for data value matching on doubleword floating point loads and stores. For non-doubleword loads and stores, only the lower 4 bits will be used.

Data value compare is included in the match condition for the data breakpoint depending on the bytes (`BYTELANE` as described above) accessed by the transaction, and the contents of breakpoint registers. The `DB_no_value_compare` is shown below.

```

DB_no_value_compare =
    ( <all 1's> == ( DBCnBLM | DBCnBAI | ~ BYTELANE ) )

```

The size of  $DBCn_{BLM}$ ,  $DBCn_{BAI}$  and BYTELANE is 8 bits.

In case a data value compare is required,  $DB\_no\_value\_compare$  is false, then the data value from the data bus (DATA) is compared and masked with the registers for the data breakpoint. The endianness is not considered in these match equations for value, as the compare uses the data bus value directly, thus debug software is responsible for setup of the breakpoint corresponding with endianness.

```
DB_value_match =
  ( ( DATA[7:0] == DBVnDBV[7:0] ) || !BYTELANE[0] || DBCnBLM[0] || DBCnBAI[0] ) &&
  ( ( DATA[15:8] == DBVnDBV[15:8] ) || !BYTELANE[1] || DBCnBLM[1] || DBCnBAI[1] ) &&
  ( ( DATA[23:16] == DBVnDBV[23:16] ) || !BYTELANE[2] || DBCnBLM[2] || DBCnBAI[2] ) &&
  ( ( DATA[31:24] == DBVnDBV[31:24] ) || !BYTELANE[3] || DBCnBLM[3] || DBCnBAI[3] ) &&
  ( ( DATA[39:32] == DBVnDBV[39:32] ) || !BYTELANE[4] || DBCnBLM[4] || DBCnBAI[4] ) &&
  ( ( DATA[47:40] == DBVnDBV[47:40] ) || !BYTELANE[5] || DBCnBLM[5] || DBCnBAI[5] ) &&
  ( ( DATA[55:48] == DBVnDBV[55:48] ) || !BYTELANE[6] || DBCnBLM[6] || DBCnBAI[6] ) &&
  ( ( DATA[63:56] == DBVnDBV[63:56] ) || !BYTELANE[7] || DBCnBLM[7] || DBCnBAI[7] ) )
```

The match for a data breakpoint without value compare is always precise, since the match expression is fully evaluated at the time the load/store instruction is executed. A true  $DB\_match$  can thereby be indicated on the very same instruction causing the  $DB\_match$  to be true. The match for data breakpoints with value compare is always imprecise.

## 16.8.6 Debug Exceptions from Breakpoints

Instruction and data breakpoints may be set up to generate a debug exception when the match condition is true, as described below.

### 16.8.6.1 Debug Exception by Instruction Breakpoint

If the breakpoint is enabled by the  $BE$  bit in the  $IBCn$  register, then a debug instruction break exception occurs if the  $IB\_match$  equation is true. The corresponding  $BS[n]$  bit in the  $IBS$  register is set when the breakpoint generates the debug exception.

The debug instruction break exception is always precise, so the  $DEPC$  register and the  $DBD$  bit in the  $Debug$  register point to the instruction that caused the  $IB\_match$  equation to be true.

The instruction receiving the debug exception does not update any registers due to the instruction, nor does any load or store by that instruction occur. Thus a debug exception from a data breakpoint cannot occur for instructions receiving a debug instruction break exception.

The debug handler usually returns to the instruction causing the debug instruction break exception, whereby the instruction is executed. Debug software is responsible for disabling the breakpoint when returning to the instruction; otherwise the debug instruction break exception reoccurs.

### 16.8.6.2 Debug Exception by Data Breakpoint

If the breakpoint is enabled by  $BE$  bit in the  $DBCn$  register, then a debug exception occurs when the  $DB\_match$  condition is true. The corresponding  $BS[n]$  bit in the  $DBS$  register is set when the breakpoint generates the debug exception. A matching data breakpoint generates either a precise or imprecise debug exception.

### ***Debug Data Break Load/Store Exception as a Precise Debug Exception***

A precise debug data break exception occurs when a data breakpoint without value compare indicates a match. In this case the *DEPC* register and *DBD* bit in the *Debug* register points to the instruction that caused the *DB\_match* equation to be true.

The instruction causing the debug data break exception does not update any registers due to the instruction, and the following applies to the load or store transaction causing the debug exception:

- A store transaction is not allowed to complete the store to the memory system.
- A load transaction with no data value compare, i.e. where the *DB\_no\_value\_compare* is true for the match, is not allowed to complete the load.

The result of this is that the load or store instruction causing the debug data break exception appears as not executed.

If both data breakpoints without and with data value compare would match the same transaction and generate a debug exception, then the rules shown in [Table 16.5](#) apply with respect to updating the *BS[n]* bits.

**Table 16.5 Rules for Update of BS Bits on Data Breakpoint Exceptions**

Instruction	Breakpoints that Match		Update of BS Bits for Matching Data Breakpoints	
	Without Value Compare	With Value Compare	Without Value Compare	With Value Compare
Load/Store	One or more	None	BS bits set for all	(No matching break-points)
Load	One or more	One or more	BS bits set for all	Unchanged BS bits since load of data value does not occur so match of the breakpoint cannot be determined
Load	None	One or more	(No matching break-points)	BS bits set for all
Store	One or more	One or more	BS bits set for all	BS bits set for all
Store	None	One or more	(No matching break-points)	BS bits set for all

Any *BS[n]* bit set prior to the match and debug exception are kept set, since *BS[n]* bits are only cleared by debug software.

The debug handler usually returns to the instruction causing the debug data break exception, whereby the instruction is re-executed. Debug software is responsible for disabling breakpoints when returning to the instruction, otherwise the debug data break exception will reoccur.

### ***Debug Data Break Load/Store Exception as a Imprecise Debug Exception***

An Debug Data Break Load/Store Imprecise exception occurs when a data breakpoint indicates an imprecise match. Imprecise matches are generated when data value compare is used. In this case, the *DEPC* register and *DBD* bit in the *Debug* register point to an instruction later in the execution flow rather than at the load/store instruction that caused the *DB\_match* equation to be true.

The load/store instruction causing the Debug Data Break Load/Store Imprecise exception always updates the destination register and completes the access to the external memory system. Therefore this load/store instruction is not re-executed on return from the debug handler, because the *DEPC* register and *DBD* bit do not point to that instruction.

Several imprecise data breakpoints can be pending at a given time, if the bus system supports multiple outstanding data accesses. The breakpoints are evaluated as the accesses finalize, and a Debug Data Break Load/Store Imprecise exception is generated only for the first one that matches. Both the first and succeeding matches cause corresponding *BS* bits and *DDBLImpr/DDBSImpr* to be set, but no debug exception is generated for succeeding matches, because the processor is already in Debug Mode. Similarly, if a debug exception had already occurred at the time of the first match (for example, due to a precise debug exception), then all matches cause the corresponding *BS* bits and *DDBLImpr/DDBSImpr* to be set, but no debug exception is generated because the processor is already in Debug Mode.

The SYNC instruction, followed by appropriate spacing must be executed before the *BS* bits and *DDBLImpr/DDBSImpr* bits are accessed for read or write. This delay ensures that these bits are fully updated.

Any *BS* bit set prior to the match and debug exception remains set, because only debug software can clear the *BS* bits.

### 16.8.7 Breakpoint used as Triggerpoint

When an enabled instruction or data breakpoint matches, the corresponding bit in the *IBS.BS* or *DBS.BS* field is set. These fields are externalized on the *SI\_lbs* and *SI\_Dbs* core outputs, respectively. These outputs are intended to be used to trigger external devices such as logic analyzers. Furthermore, breakpoint matches can also be used to start or stop PDtrace. See [Section 16.11.8 “Enabling PDtrace”](#) for details.

If the breakpoints are to be used only as trigger events, the signalling of the debug exception can be suppressed by clearing the *IBCn/DBCn.BE* field and setting the *IBCn/DBCn.TE* field.

## 16.9 Debug Vector Addressing

The debug vector address size is managed by the *Debug Vector Address* register as described in [Section 16.14.1.2 “DebugVectorAddr Register”](#). The *Debug Vector Address* register is a read/write register containing the base address of the debug exception vectors in bits 31:7, and a WG bit that determines whether the bits 31:30 of this field are a fixed value, or are programmable.

Bits 31:12 of the *DebugVectorAddress* register are concatenated with zeros to form the base of the debug exception vector. The exception vector base address comes from the fixed defaults for any EJTAG Debug exception. The reset state of bits 31:12 of the *DebugVectorAddress* register initialize the exception base register to 0xFC00\_0480.

The size of the *DebugVectorAddr* field depends on the state of the WG bit. At reset, the WG bit is cleared by default. In this case, the *DebugVectorAddr* field is comprised of bits 29:7. Bits 31:30 of the *DebugVectorAddr* Register are not writeable and are forced to a value of 2'b10 by hardware so that the debug exception handler will be executed from the *kseg0/kseg1* segments.

When the WG bit is set, bits 31:30 of the *DebugVectorAddr* field become writeable and are used to relocate the *DebugVectorAddr* field to other segments after they have been setup using the *SegCtl0* through *SegCtl2* registers. Note that if the WG bit is set by software (allowing bits 31:30 to become part of the *DebugVectorAddr* field) and then cleared, bits 31:30 can no longer be written by software and the state of these bits remains unchanged for any writes after WG was cleared. Therefore, it is the responsibility of software to write a value of 2'b10 to bits 31:30 of the *DebugVectorAddr* register prior to clearing the WG bit if it wants to ensure that future debug exceptions will be executed from the *kseg0* or *kseg1* segments.



Note that the WG bit is different from the CV bit in the SegCtl0 register located in [Section 2.2.3.1, "Segmentation Control 0 — SegCtl0 \(CP0 Register 5, Select 2\)"](#). Although their functions are similar, the CV bit applies only to cache error exceptions, whereas the WG bit applies to all exceptions.

If the value of the exception base register is to be changed, this must be done with *StatusBEV* equal to 1. The operation of the processor is **UNDEFINED** if the exception base field is written with a different value when *StatusBEV* is 0.

[Table 16.11](#) shows the different debug exception vector locations that are possible.

**Table 16.6 Debug Exception Vectors**

ECR <sub>ProbTrap</sub>	DCR <sub>RdVec</sub>	Config5 <sub>K</sub>	SI_UseExceptionBase	Cache Error?	Debug Exception Vector
1	x	x	x	x	0xFF20_0200
0	1	0	x	0	2'b10    DebugVectorAddr[29:0]
0	1	1	x	0	DebugVectorAddr[31:0]
0	1	0	x	1	3'b101    DebugVectorAddr[28:0]
0	1	1	x	1	DebugVectorAddr[31:0]
0	0	0	1	0	2'b10    SI_ExceptionBase[29:12]    0x480
0	0	1	1	0	SI_ExceptionBase[31:12]    0x480
0	0	0	1	1	3'b101    SI_ExceptionBase[28:12]    0x480
0	0	1	1	1	SI_ExceptionBase[31:12]    0x480
0	0	x	0	x	0xBFC0_0480

As shown in the table above, if the *ECR<sub>ProbeTrap</sub>* bit (14) is set in the EJTAG Control register, then all other bits or signals that determine the location of the debug vector address have no meaning and the location of the debug exception vector default to 0xFF20\_0200. Note that the *ECR<sub>ProbeEn</sub>* bit (15) must be set in order for this bit to have meaning.

## 16.10 Test Access Port (TAP)

The TAP is used only when a probe is connected to the interAptiv core.

The following main features are supported by the TAP module:

- 5-pin industry standard JTAG Test Access Port (*TCK*, *TMS*, *TDI*, *TDO*, *TRST\_N*) interface which is compatible with IEEE Std. 1149.1.
- Target chip and EJTAG feature identification available through the Test Access Port (TAP) controller.
- The processor can access external memory on the EJTAG Probe serially through the EJTAG pins. This is achieved through Processor Access (PA), and is used to eliminate the use of the system memory for debug routines.
- Support for both ROM based debugger and debugging both through TAP.



## 16.10.1 EJTAG Internal and External Interfaces

The external interface of the EJTAG module consists of the 5 signals defined by the IEEE standard.

**Table 16.7 EJTAG Interface Pins**

Pin	Type	Description
<i>TCK</i>	I	Test Clock Input Input clock used to shift data into or out of the Instruction or data registers. The <i>TCK</i> clock is independent of the processor clock, so the EJTAG probe can drive <i>TCK</i> independently of the processor clock frequency. The core signal for this is called <i>EJ_TCK</i>
<i>TMS</i>	I	Test Mode Select Input The <i>TMS</i> input signal is decoded by the TAP controller to control test operation. <i>TMS</i> is sampled on the rising edge of <i>TCK</i> . The core signal for this is called <i>EJ_TMS</i>
<i>TDI</i>	I	Test Data Input Serial input data ( <i>TDI</i> ) is shifted into the Instruction register or data registers on the rising edge of the <i>TCK</i> clock, depending on the TAP controller state. The core signal for this is called <i>EJ_TDI</i>
<i>TDO</i>	O	Test Data Output Serial output data is shifted from the Instruction or data register to the <i>TDO</i> pin on the falling edge of the <i>TCK</i> clock. When no data is shifted out, the <i>TDO</i> is 3-stated. The core signal for this is called <i>EJ_TDO</i> with output enable controlled by <i>EJ_TDOzstate</i> .
<i>TRST_N</i>	I	Test Reset Input (Optional pin) The <i>TRST_N</i> pin is an active-low signal for asynchronous reset of the TAP controller and instruction in the TAP module, independent of the processor logic. The processor is not reset by the assertion of <i>TRST_N</i> . The core signal for this is called <i>EJ_TRST_N</i> This signal is optional, but power-on reset must apply a low pulse on this signal at power-on and then leave it high, in case the signal is not available as a pin on the chip. If available on the chip, then it must be low on the board when the EJTAG debug features are unused by the probe.

## 16.10.2 Test Access Port Operation

The TAP controller is controlled by the Test Clock (*TCK*) and Test Mode Select (*TMS*) inputs. These two inputs determine whether an the Instruction register scan or data register scan is performed. The TAP consists of a small controller, driven by the *TCK* input, which responds to the *TMS* input as shown in the state diagram in [Figure 16.2](#). The TAP uses both clock edges of *TCK*. *TMS* and *TDI* are sampled on the rising edge of *TCK*, while *TDO* changes on the falling edge of *TCK*.

At power-up the TAP is forced into the *Test-Logic-Reset* by low value on *TRST\_N*. The TAP instruction register is thereby reset to IDCODE. No other parts of the EJTAG hardware are reset through the *Test-Logic-Reset* state.


When test access is required, a protocol is applied via the *TMS* and *TCK* inputs, causing the TAP to exit the *Test-Logic-Reset* state and move through the appropriate states. From the *Run-Test/Idle* state, an Instruction register scan or a data register scan can be issued to transition the TAP through the appropriate states shown in [Figure 16.2](#).

The states of the data and instruction register scan blocks are mirror images of each other adding symmetry to the protocol sequences. The first action that occurs when either block is entered is a capture operation. For the data registers, the *Capture-DR* state is used to capture (or parallel load) the data into the selected serial data path. In the Instruction register, the *Capture-IR* state is used to capture status information into the Instruction register.

From the *Capture* states, the TAP transitions to either the *Shift* or *Exit1* states. Normally the *Shift* state follows the *Capture* state so that test data or status information can be shifted out for inspection and new data shifted in. Following the *Shift* state, the TAP either returns to the *Run-Test/Idle* state via the *Exit1* and *Update* states or enters the *Pause* state via *Exit1*. The reason for entering the *Pause* state is to temporarily suspend the shifting of data through either the Data or Instruction Register while a required operation, such as refilling a host memory buffer, is performed. From the *Pause* state shifting can resume by re-entering the *Shift* state via the *Exit2* state or terminate by entering the *Run-Test/Idle* state via the *Exit2* and *Update* states.

Upon entering the data or Instruction register scan blocks, shadow latches in the selected scan path are forced to hold their present state during the *Capture* and *Shift* operations. The data being shifted into the selected scan path is not output through the shadow latch until the TAP enters the *Update-DR* or *Update-IR* state. The *Update* state causes the shadow latches to update (or parallel load) with the new data that has been shifted into the selected scan path.

**Figure 16.2 TAP Controller State Diagram**



### 16.10.2.1 Test-Logic-Reset State

In the *Test-Logic-Reset* state the boundary scan test logic is disabled. The test logic enters the *Test-Logic-Reset* state when the *TMS* input is held HIGH for at least five rising edges of *TCK*. The *BYPASS* instruction is forced into the instruction register output latches during this state. The controller remains in the *Test-Logic-Reset* state as long as *TMS* is HIGH.

### 16.10.2.2 Run-Test/Idle State

The controller enters the *Run-Test/Idle* state between scan operations. The controller remains in this state as long as *TMS* is held LOW. The instruction register and all test data registers retain their previous state. The instruction cannot change when the TAP controller is in this state.

When *TMS* is sampled HIGH on the rising edge of *TCK*, the controller transitions to the *Select\_DR* state.

#### 16.10.2.3 Select\_DR\_Scan State

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, then the controller transitions to the *Capture\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Select\_IR* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.4 Select\_IR\_Scan State

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW on the rising edge of *TCK*, the controller transitions to the *Capture\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Test-Reset-Logic* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.5 Capture\_DR State

In this state the boundary scan register captures the value of the register addressed by the Instruction register, and the value is then shifted out in the *Shift\_DR*. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Shift\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_DR* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.6 Shift\_DR State

In this state the test data register connected between *TDI* and *TDO* as a result of the current instruction shifts data one stage toward its serial output on the rising edge of *TCK*. If *TMS* is sampled LOW on the rising edge of *TCK*, the controller remains in the *Shift\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_DR* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.7 Exit1\_DR State

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Pause\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Update\_DR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.8 Pause\_DR State

The *Pause\_DR* state allows the controller to temporarily halt the shifting of data through the test data register in the serial path between *TDI* and *TDO*. All test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW on the rising edge of *TCK*, the controller remains in the *Pause\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Exit2\_DR* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.9 Exit2\_DR State

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Shift\_DR* state to allow another serial shift of data. A HIGH on *TMS* causes the controller to transition to the *Update\_DR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.10 Update\_DR State

When the TAP controller is in this state the value shifted in during the *Shift\_DR* state takes effect on the rising edge of the *TCK* for the register indicated by the Instruction register.

If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Run-Test/Idle* state. A HIGH on *TMS* causes the controller to transition to the *Select\_DR\_Scan* state. The instruction cannot change while the TAP controller is in this state and all shift register stages in the test data registers selected by the current instruction retain their previous state.

#### 16.10.2.11 Capture\_IR State

In this state the shift register contained in the Instruction register loads a fixed pattern (00001<sub>2</sub>) on the rising edge of *TCK*. The data registers selected by the current instruction retain their previous state.

If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Shift\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_IR* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.12 Shift\_IR State

In this state the instruction register is connected between *TDI* and *TDO* and shifts data one stage toward its serial output on the rising edge of *TCK*. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller remains in the *Shift\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_IR* state.

#### 16.10.2.13 Exit1\_IR State

This is a temporary controller state in which all registers retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Pause\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Update\_IR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state and the instruction register retains its previous state.

#### 16.10.2.14 Pause\_IR State

The *Pause\_IR* state allows the controller to temporarily halt the shifting of data through the instruction register in the serial path between *TDI* and *TDO*. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller remains in the *Pause\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Exit2\_IR* state. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.15 Exit2\_IR State

This is a temporary controller state in which the instruction register retains its previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, then the controller transitions to the *Shift\_IR* state to allow another serial shift of data. A HIGH on *TMS* causes the controller to transition to the *Update\_IR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state.

#### 16.10.2.16 Update\_IR State

The instruction shifted into the instruction register takes effect on the rising edge of *TCK*.

If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Run-Test/Idle* state. A HIGH on *TMS* causes the controller to transition to the *Select\_DR\_Scan* state.

### 16.10.3 Test Access Port (TAP) Instructions

The TAP Instruction register allows instructions to be serially input into the device when TAP controller is in the *Shift-IR* state. Instructions are decoded and define the serial test data register path that is used to shift data between *TDI* and *TDO* during data register scanning.

The Instruction register is a 5-bit register. In the current EJTAG implementation only some instructions have been decoded; the unused instructions default to the BYPASS instruction.

**Table 16.8 Implemented EJTAG Instructions**

Value	Instruction	Function
0x01	IDCODE	Select Chip Identification data register
0x03	IMPCODE	Select Implementation register
0x08	ADDRESS	Select Address register
0x09	DATA	Select Data register
0x0A	CONTROL	Select EJTAG Control register
0x0B	ALL	Select the Address, Data and EJTAG Control registers
0x0C	EJTAGBOOT	Set EjtagBrk, ProbEn and ProbTrap to 1 as reset value
0x0D	NORMALBOOT	Set EjtagBrk, ProbEn and ProbTrap to 0 as reset value
0x0E	FASTDATA	Selects the Data and Fastdata registers
0x10	TCBCONTROLA	Selects the <i>TCBTCONTROLA</i> register in the Trace Control Block
0x11	TCBCONTROLB	Selects the <i>TCBTCONTROLB</i> register in the Trace Control Block
0x12	TCBDATA	Selects the <i>TCBDATA</i> register in the Trace Control Block
0x13	TCBCONTROLC	Selects the <i>TCBTCONTROLC</i> register in the Trace Control Block
0x14	PCSAMPLE	Selects the <i>PCSAMPLE</i> register
0x15	TCBCONTROLD	Selects the <i>TCBTCONTROLD</i> register in the Trace Control Block
0x16	TCBCONTROLE	Selects the <i>TCBTCONTROLE</i> register in the Trace Control Block
0x17	FDC	Select Fast Debug Channel
0x1F	BYPASS	Bypass mode

#### 16.10.3.1 BYPASS Instruction

The required BYPASS instruction allows the processor to remain in a functional mode and selects the Bypass register to be connected between *TDI* and *TDO*. The BYPASS instruction allows serial data to be transferred through the processor from *TDI* to *TDO* without affecting its operation. The bit code of this instruction is defined to be all ones by the IEEE 1149.1 standard. Any unused instruction is defaulted to the BYPASS instruction.

#### 16.10.3.2 IDCODE Instruction

The IDCODE instruction allows the processor to remain in its functional mode and selects the Device Identification (ID) register to be connected between *TDI* and *TDO*. The Device ID register is a 32-bit shift register containing information regarding the IC manufacturer, device type, and version code. Accessing the Identification Register does not interfere with the operation of the processor. Also, access to the Identification Register is immediately available, via a TAP data scan operation, after power-up when the TAP has been reset with on-chip power-on or through the optional *TRST\_N* pin.

### 16.10.3.3 IMPCODE Instruction

This instruction selects the Implementation register for output, which is always 32 bits.

### 16.10.3.4 ADDRESS Instruction

This instruction is used to select the Address register to be connected between *TDI* and *TDO*. The EJTAG Probe shifts 32 bits through the *TDI* pin into the Address register and shifts out the captured address via the *TDO* pin.

### 16.10.3.5 DATA Instruction

This instruction is used to select the Data register to be connected between *TDI* and *TDO*. The EJTAG Probe shifts 32 bits of *TDI* data into the Data register and shifts out the captured data via the *TDO* pin.


### 16.10.3.6 CONTROL Instruction

This instruction is used to select the EJTAG Control register to be connected between *TDI* and *TDO*. The EJTAG Probe shifts 32 bits of *TDI* data into the EJTAG Control register and shifts out the EJTAG Control register bits via *TDO*.

### 16.10.3.7 ALL Instruction

This instruction is used to select the concatenation of the Address and Data register, and the EJTAG Control register (ECR) between *TDI* and *TDO*. It can be used in particular to minimize the overhead in switching the instruction in the instruction register. The first bit shifted out is bit 0 of the ECR.

**Figure 16.3 Concatenation of the EJTAG Address, Data and Control Registers**



### 16.10.3.8 EJTAGBOOT Instruction

EJTAGBOOT provides a means to enter debug mode just after a reset, without fetching or executing any instructions from the normal memory area. This can be used for download of code to a system which has no code in ROM.

When the EJTAGBOOT instruction is given and the Update-IR state is left, the EJTAGBOOT indication will become active. When EJTAGBOOT is active, a core reset will set the ProbTrap, ProbEn and EjtagBrk bits in the EJTAG Control register to 1. This will cause a debug exception that is serviced by the probe immediately after reset is deasserted.

This EJTAGBOOT indication is effective until a NORMALBOOT instruction is given, *TRST\_N* is asserted or a rising edge of *TCK* occurs when the TAP controller is in Test-Logic-Reset state.

Each VPE has its own TAP controller and thus EJTAGBOOT can be set independently per VPE. Even though a VPE is not activated at core reset, EJTAGBOOT on that VPE will still cause a debug exception immediately after reset.

The Bypass register is selected when the EJTAGBOOT instruction is given.



### 16.10.3.9 NORMALBOOT Instruction


When the NORMALBOOT instruction is given and the Update-IR state is left, then the EJTAGBOOT indication will be cleared. When NORMALBOOT is active (EJTAGBOOT is not active), a core reset will set the ProbTrap, ProbEn and EhtagBrk bits in the EJTAG Control register to 0.

The Bypass register is selected when the NORMALBOOT instruction is given.

### 16.10.3.10 FASTDATA Instruction

This selects the Data and the Fastdata registers at once, as shown in [Figure 16.4](#).

**Figure 16.4 TDI to TDO Path When in Shift-DR State and FASTDATA Instruction is Selected**



The FASTDATA access is used for efficient block transfers between dmseg (on the probe) and target memory (on the processor). An “upload” is defined as a sequence of processor loads from target memory and stores to dmseg. A “download” is a sequence of processor loads from dmseg and stores to target memory. The “Fastdata area” specifies the legal range of dmseg addresses (0xFF20.0000 - 0xFF20.000F) that can be used for uploads and downloads. The Data + Fastdata registers (selected with the FASTDATA instruction) allow efficient completion of pending Fastdata area accesses.

During Fastdata uploads and downloads, the processor will stall on accesses to the Fastdata area. The PrAcc (processor access pending bit) will be 1 indicating the probe is required to complete the access. Both upload and download accesses are attempted by shifting in a zero *SPrAcc* value (to request access completion) and shifting out *SPrAcc* to see if the attempt will be successful (i.e., there was an access pending and a legal Fastdata area address was used). Downloads will also shift in the data to be used to satisfy the load from dmseg’s Fastdata area, while uploads will shift out the data being stored to dmseg’s Fastdata area.

As noted above, two conditions must be true for the Fastdata access to succeed. These are:

- *PrAcc* must be 1, i.e., there must be a pending processor access.
- The Fastdata operation must use a valid Fastdata area address in dmseg (0xFF20.0000 to 0xFF20.000F).

[Table 16.9](#) shows the values of the *PrAcc* and *SPrAcc* bits and the results of a Fastdata access.

**Table 16.9 Operation of the FASTDATA Access**

Probe Operation	Address Match Check	PrAcc in the Control Register	LSB (SPrAcc) Shifted In	Action in the Data Register	PrAcc Changes to	Lsb Shifted Out	Data Shifted Out
Download using FAST-DATA	Fails	x	x	none	unchanged	0	invalid
	Passes	1	1	none	unchanged	1	invalid
		1	0	write data	0 (SPrAcc)	1	valid (previous) data
		0	x	none	unchanged	0	invalid

**Table 16.9 Operation of the FASTDATA Access (continued)**

Probe Operation	Address Match Check	PrAcc in the Control Register	LSB (SPrAcc) Shifted In	Action in the Data Register	PrAcc Changes to	Lsb Shifted Out	Data Shifted Out
Upload using FASTDATA	Fails	x	x	none	unchanged	0	invalid
	Passes	1	1	none	unchanged	1	invalid
		1	0	read data	0 (SPrAcc)	1	valid data
		0	x	none	unchanged	0	invalid

There is no restriction on the contents of the Data register. It is expected that the transfer size is negotiated between the download/upload transfer code and the probe software. Note that the most efficient transfer size is a 32-bit word.

The Rocc bit of the Control register is not used for the FASTDATA operation.

#### 16.10.3.11 TCBCONTROLA Instruction

This instruction is used to select the TCBCONTROLA register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

#### 16.10.3.12 TCBCONTROLB Instruction

This instruction is used to select the TCBCONTROLB register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

#### 16.10.3.13 TCBCONTROLC Instruction

This instruction is used to select the TCBCONTROLC register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

#### 16.10.3.14 TCBDATA Instruction

This instruction is used to select the TCBDATA register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register. It should be noted that the TCBDATA register is only an access register to other TCB registers. The width of the TCBDATA register is dependent on the specific TCB register.

#### 16.10.3.15 PCSAMPLE Instruction

This instruction is used to select the PCSAMPLE register to be connected between *TDI* and *TDO*. This register is always implemented.

#### 16.10.3.16 TCBCONTROLD Instruction

This instruction is used to select the TCBCONTROLD register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.



### 16.10.3.17 TCBCONTROLE Instruction

This instruction is used to select the TCBCONTROLE register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

### 16.10.3.18 FDC Instruction

This instruction is used to select the Fast Debug Channel register to be connected between *TDI* and *TDO*. This register is always implemented

## 16.10.4 TAP Processor Accesses

The TAP modules support handling of fetches, loads and stores from the CPU through the dmseg segment, whereby the TAP module can operate like a *slave unit* connected to the on-chip bus. The core can then execute code taken from the EJTAG Probe and it can access data (via a load or store) which is located on the EJTAG Probe. This occurs in a serial way through the EJTAG interface: the core can thus execute instructions e.g. debug monitor code, without occupying the memory.

Accessing the dmseg segment (EJTAG memory) can only occur when the processor accesses an address in the range from 0xFF20.0000 to 0xFF2F.FFFF, the ProbEn bit is set, and the processor is in debug mode (DM=1). In addition the LSNM bit in the CP0 Debug register controls transactions to/from the dmseg.

When a debug exception is taken, while the ProbTrap bit is set, the processor will start fetching instructions from address 0xFF20.0200.

A pending processor access can only finish if the probe writes 0 to PrAcc or by a reset.

### 16.10.4.1 Fetch/Load and Store From/To the EJTAG Probe Through dmseg

1. The internal hardware latches the requested address into the Address register (in case of the Debug exception: 0xFF20.0200).
2. The internal hardware sets the following bits in the EJTAG Control register:  
PrAcc = 1 (selects Processor Access operation)  
PRnW = 0 (selects processor read operation)  
Psz[1:0] = value depending on the transfer size
3. The EJTAG Probe selects the EJTAG Control register, shifts out this control register's data and tests the PrAcc status bit (Processor Access): when the PrAcc bit is found 1, it means that the requested address is available and can be shifted out.
4. The EJTAG Probe checks the PRnW bit to determine the required access.
5. The EJTAG Probe selects the Address register and shifts out the requested address.
6. The EJTAG Probe selects the Data register and shifts in the instruction corresponding to this address.
7. The EJTAG Probe selects the EJTAG Control register and shifts a PrAcc = 0 bit into this register to indicate to the processor that the instruction is available.
8. The instruction becomes available in the instruction register and the processor starts executing.

9. The processor increments the program counter and outputs an instruction read request for the next instruction. This starts the whole sequence again.

Using the same protocol, the processor can also execute a load instruction to access the EJTAG Probe's memory. For this to happen, the processor must execute a load instruction (e.g. a LW, LH, LB) with the target address in the appropriate range.

Almost the same protocol is used to execute a store instruction to the EJTAG Probe's memory through dmseg. The store address must be in the range: 0xFF20.0000 to 0xFF2F.FFFF, the ProbEn bit must be set and the processor has to be in debug mode (DM=1). The sequence of actions is found below:

1. The internal hardware latches the requested address into the Address register
2. The internal hardware latches the data to be written into the Data register.
3. The internal hardware sets the following bits in the EJTAG Control register:  
PrAcc = 1 (selects Processor Access operation)  
PRnW = 1 (selects processor write operation)  
Psz[1:0] = value depending on the transfer size
4. The EJTAG Probe selects the EJTAG Control register, shifts out this control register's data and tests the PrAcc status bit (Processor Access): when the PrAcc bit is found 1, it means that the requested address is available and can be shifted out.
5. The EJTAG Probe checks the PRnW bit to determine the required access.
6. The EJTAG Probe selects the Address register and shifts out the requested address.
7. The EJTAG Probe selects the Data register and shifts out the data that was written.
8. The EJTAG Probe selects the EJTAG Control register and shifts a PrAcc = 0 bit into this register to indicate to the processor that the write access is finished.
9. The EJTAG Probe writes the data to the appropriate address in its memory.
10. The processor detects that PrAcc bit = 0, which means that it is ready to handle a new access.

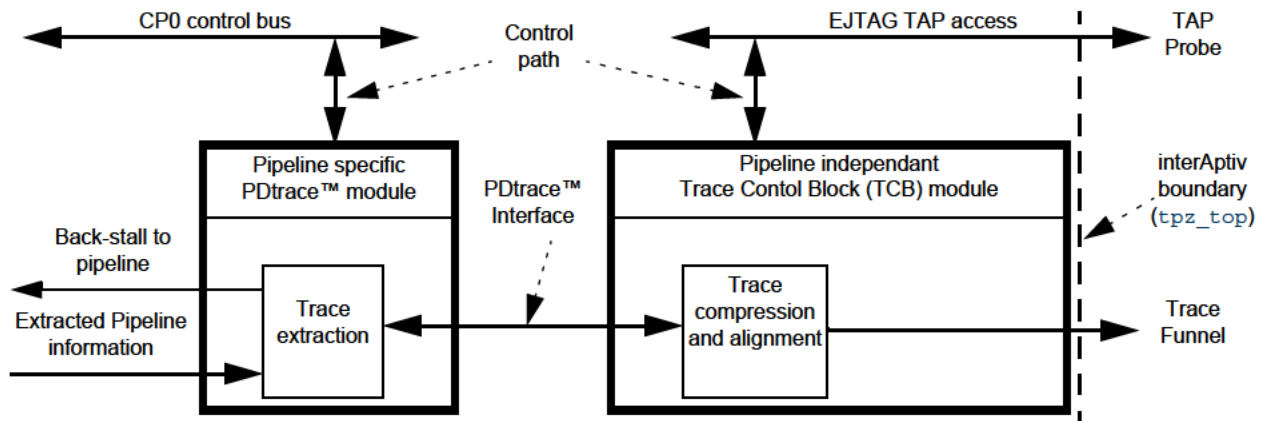
The above examples imply that no reset occurs during the operations, and that Rocc is cleared.

## 16.11 PDTrace

PDTrace enables the ability to trace program flow, load/store addresses and load/store data. Several run-time options exist for the level of information which is traced, including tracing only when in specific processor modes (e.g., User-Mode or KernelMode). PDtrace is an optional block in the interAptiv core. If PDtrace is not implemented, the rest of this chapter is irrelevant. If PDtrace is implemented, the *CP0 Config3<sub>TL</sub>* bit is set.

There are two primary blocks involved in the PDtrace solution. The pipeline specific part of PDtrace is called the PDtrace module. It extracts the trace information from the processor pipeline, and presents it to a pipeline-independent module called the Trace Control Block (TCB). While working closely together, the two parts of PDtrace are controlled separately by software. [Figure 16.5](#) shows an overview of the PDtrace modules within the core.

**Figure 16.5 MIPS® Trace Modules in the interAptiv™ Core**



To some extent, the two modules both provide similar trace control features, but the access to these features is quite different. The PDtrace controls can only be reached through access to CP0 registers. In general, the PDtrace control registers select what information is captured for tracing. The TCB controls can be reached through EJTAG TAP access or through load/store access to registers mapped in drseg space. The TCB registers control what is traced through the PDtrace™ Interface.

Before describing the PDtrace implemented in the interAptiv core, some common terminology and basic features are explained. The remaining sections of this chapter will then provide a more thorough explanation.

### 16.11.1 Processor Modes

Tracing can be enabled or disabled based on various processor modes. This section precisely describes these modes. The terminology is then used elsewhere in the document.

```

DebugMode ← (DebugDM = 1)
ExceptionMode ← (not DebugMode) and ((StatusEXL = 1) or (StatusERL = 1))
KernelMode ← (not (DebugMode or ExceptionMode)) and (StatusKSU = 2#00)
SupervisorMode ← (not (DebugMode or ExceptionMode)) and (StatusKSU = 2#01)
UserMode ← (not (DebugMode or ExceptionMode)) and (StatusKSU = 2#10)
    
```

### 16.11.2 Software Versus Hardware Control

In some of the specifications and in this text, the terms “software control” and “hardware control” are used to refer to the method for how trace is controlled. Software control is when the CP0 register *TraceControl* is used to select the modes to trace, etc. Hardware control is when the EJTAG register *TCBCONTROLA* in the TCB, via the PDtrace interface, is used to select the trace modes. The *TraceControl<sub>TS</sub>* bit determines whether software or hardware control is active.

### 16.11.3 Trace Information

The main object of trace is to show the exact program flow from a specific program execution or just a small window of the execution. In PDtrace this is done by providing the minimal cycle-by-cycle information necessary on the PDtrace™ interface for trace regeneration software to reproduce the trace. The following is a summary of the type of information traced:

- Only instructions which complete at the end of the pipeline are traced, and indicated with a completion-flag. The PC is implicitly pointing to the next instruction.
- Load instructions are indicated with a load-flag.
- Store instructions are indicated with a store-flag<sup>1</sup>.
- Taken branches are indicated with a branch-taken-flag on the target instruction.
- New PC information for a branch is only traced if the branch target is unpredictable from the static program image.
- When branch targets are unpredictable, only the delta value from current PC is traced, if it is dynamically determined to reduce the number of bits necessary to indicate the new PC. Otherwise the full PC value is traced.
- When a completing instruction is executed in a different processor mode from the previous one, the new processor mode is traced.
- The first instruction is always traced as a branch target, with processor mode and full PC.
- Periodic synchronization instructions are identified with a sync-flag, and traced with the processor mode and full PC.

All the instruction flags above are combined into one 3-bit value, to minimize the bit information to trace. The possible processor modes are explained in [Section 16.11.1 “Processor Modes”](#).

The target address is statically predictable for all branch and all jump-immediate instructions. If the branch is taken, then the branch-taken-flag will indicate this. All jump-register instructions and ERET/DERET are instructions which have an unpredictable target address. These will have full/delta PC values included in the trace information. Also treated as unpredictable are PC changes which occur due to exceptions, such as an interrupt, reset, etc.

Trace regeneration software is required to know the static program image in memory, in order to reproduce the dynamic flow with the above information. Only the virtual value of the PC is used. Physical memory location will typically differ.

It is possible to turn on PC delta/full information for all branches, but this should not normally be necessary. As a safety check for trace regeneration software, a periodic synchronization with a full PC is sent. The period of this synchronization is cycle based and programmable.

#### 16.11.4 Load/Store Address and Data Trace Information

In addition to PC flow, it is possible to get information on the load/store addresses, as well as the data read/written. When enabled, the following information is optionally added to the trace.

- When load-address tracing is on, the full load address of the first load instruction is traced (indicated by the load-flag). For subsequent loads, a dynamically-determined delta to the previous load address is traced to compress the information which must be sent.
- When store-address tracing is on, the full store address of the first store instruction is traced (indicated by the store-flag). For subsequent stores, a dynamically-determined delta to the previous store address is traced.
- When load-data tracing is on, the full load data read by each load instruction is traced (indicated by the load-flag). Only actual read bytes are traced.

---

1. A SC (Store Conditional) instruction is not flagged as a store instruction if the load-locked bit prevented the actual store.

- When store-data tracing is on, the full store data written by each store instruction is traced (indicated by the store-flag). Only written bytes are traced.

After each synchronization instruction, the first load address and the first store address following this are both traced with the full address if load/store address tracing is enabled.

### 16.11.5 Programmable Processor Trace Mode Options

To enable tracing, a global Trace On signal must be set. When trace is on, it is possible to enable tracing in any combination of the processor modes described in [Section 16.11.1 “Processor Modes”](#). In addition to this, trace can be turned on globally for all processes, or only for specific processes by tracing only specific masked values of the ASID found in *EntryHi*<sub>ASID</sub>.

Additionally, an EJTAG Simple Break trigger point can override the processor mode and ASID selection and turn them all on. Another trigger point can disable this override again.

### 16.11.6 Programmable Trace Information Options

The processor mode changes are always traced:

- On the first instruction.
- On any synchronization instruction.
- When the mode changes and either the previous or the current processor mode is selected for trace.

The amount of extra information traced is programmable to include:

- PC information only.
- PC and cross product of load/store address/data
- Performance counter values, if the optional performance counter trace is enabled.

If the full internal state of the processor is known prior to trace start, PC and load data are the only information needed to recreate all register values on an instruction by instruction basis.

#### 16.11.6.1 User Data Trace

Two special CP0 registers, *UserTraceData1* and *UserTraceData2*, can generate a data trace. When either of these registers is written, and the global Trace On is set, then the 32-bit data written is put in the trace as special User Data information. Since writing these registers is performed via an MTC0 operation, only one register is updated in any given cycle. Thus in the same cycle, only one of the UserTraceData registers is traced. However in back to back cycles, the tracing of the two registers can alternate, and is handled correctly.

*Remark:* The User Data is sent even if the processor is operating in an un-traced processor mode.

### 16.11.7 Enable Trace to Probe On-Chip Memory

When trace is On, based on the options listed in [Section 16.11.5 “Programmable Processor Trace Mode Options”](#), the trace information is continuously sent on the PDtrace™ interface to the TCB. The TCB must be enabled to transmit the trace information to the Trace funnel by having the *TCBCONTROLB*<sub>EN</sub> bit set. It is possible to enable and disable the TCB in a number of ways:

- Set/clear the *TCBCONTROLB<sub>EN</sub>* bit via an EJTAG TAP operation.
- Initialize a TCB trigger to set/clear the *TCBCONTROLB<sub>EN</sub>* bit.
- Use the drseg mapping of *TCBCONTROLB* to clear *TCBCONTROLB<sub>EN</sub>* via a store to drseg space.

## 16.11.8 Enabling PDtrace

As there are several ways to enable tracing, it can be quite confusing to figure out how to turn tracing on and off. This section should help clarify the enabling of trace.

### 16.11.8.1 Trace Trigger from EJTAG Hardware Instruction/Data Breakpoints

If hardware instruction/data simple breakpoints are implemented in the interAptiv core, then these breakpoints can be used as triggers to start/stop trace. When used for this, the breakpoints need not also generate a debug exception, but are capable of only generating an internal trigger to the trace logic. This is done by only setting the TE bit and not the BE bit in the Breakpoint Control register. Please see [Section 16.14.2.5 “Instruction Breakpoint Control n \(IBn\) Register”](#) and [Section 16.14.3.5 “Data Breakpoint Control n \(DBCn\) Register”](#) for details on breakpoint control.

In connection with the breakpoints, the Trace BreakPoint Control (*TraceBPC*) register is used to define the trace action when a trigger happens. When a breakpoint is enabled as a trigger (TE = 1), it can be selected to be either a start or a stop trigger to the trace logic.

### 16.11.8.2 Turning On PDtrace™ Trace

Trace enabling and disabling from software is similar to the hardware method, with the exception that the bits in the control register are used instead of the input enable signals from the TCB. The *TraceControl<sub>TS</sub>* bit controls whether hardware (via the TCB), or software (via the *TraceControl* register) controls tracing functionality.

Trace is turned on when the following expression evaluates true:

```
(
  (
    (TraceControlTS and TraceControlOn) or
    ((not TraceControlTS) and TCBCONTROLAOn)
  )
  and
  (MatchEnable or TriggerEnable)
)
```

where,

```
MatchEnable ←
(
  TraceControlTS
  and
  ((TraceControl2TCV and (TraceControl2TCNUM equal TCIDofCompletedInst)) or
  ((not TraceControl2TCV) and TraceControl2CPUIDv and
  (TraceControl2CPUID equal CPUIDofCompletedInst )) or
  (TraceControl2TCV nor TraceControl2CPUIDv ))
  and
  (
    TraceControlG or
    (((TraceControlASID xor EntryHiASID) and (not TraceControlASID_M)) = 0)
  )
)
```

```

and
(
  (TraceControlU and UserMode)      or
  (TraceControlS and SupervisorMode) or
  (TraceControlK and KernelMode)    or
  (TraceControlE and ExceptionMode) or
  (TraceControlD and DebugMode)
)
)
or
(
  (not TraceControlTS)
and
  ((TCBCONTROLCTCV and (TCBCONTROLCTNUM equal TCIDofCompletedInst)) or
  ((not TCBCONTROLCTCV) and TCBCONTROLCCPUIDV and
  (TCBCONTROLCCPUID equal CPUIDofCompletedIns )) or
  (TCBCONTROLCTCV nor TCBCONTROLCCPUIDV ))
and
  (TCBCONTROLAG or (TCBCONTROLAASID = EntryHiASID))
and
  (
    (TCBCONTROLAU and UserMode)      or
    (TCBCONTROLAS and SupervisorMode) or
    (TCBCONTROLAK and KernelMode)    or
    (TCBCONTROLAE and ExceptionMode) or
    (TCBCONTROLADM and DebugMode)
  )
)
)

```

and where,

```

TriggerEnable ←
(
  DBCiTE      and
  DBSBS[i]    and
  TraceBPCDE and
  (TraceBPCDBPOn[i] = 1)
)
or
(
  IBCiTE      and
  IBSBS[i]    and
  TraceBPCIE and
  (TraceBPCIBPOn[i] = 1)
)

```

As seen in the expression above, trace can be turned on only if the master switch *TraceControl<sub>On</sub>* or *TCBCONTROL<sub>A</sub><sub>On</sub>* is first asserted.

Once this is asserted, there are three ways to turn on tracing. The first way, the *MatchEnable* expression, uses the input enable signals from the TCB or the bits in the *TraceControl* register. This tracing is done over general program areas. For example, all of the user-level code for a particular process (if ASID is specified), and so on.

The second way to turn on tracing, the *TriggerEnable* expression, is from the processor side using the EJTAG hardware breakpoint triggers. If EJTAG is implemented, and hardware breakpoints can be set, then using this method enables finer grain tracing control. It is possible to send a trigger signal that turns on tracing at a particular instruction.

For example, it would be possible to trace a single procedure in a program by triggering on trace at the first instruction, and triggering off trace at the last instruction.

The third way to enable tracing is in Filtered Data Trace Mode. When this mode is enabled, data load and store addresses are compared to the hardware data breakpoint address, if the addresses match, the data value associated with that match along with the address are traced out.

The easiest way to unconditionally turn on trace is to assert either hardware or software tracing and the corresponding trace on signal with other enables. For example, with  $TraceControl_{TS} = 0$ , i.e., hardware controlled tracing, assert  $TCBCONTROLA_{On}$ ,  $TCBCONTROLA_G$  and all the other signals in the second part of expression *MatchEnable*. To only trace when a particular process with a known ASID is executing, assert  $TCBCONTROLA_{On}$ , the correct  $TCBCONTROLA_{ASID}$  value, and all of  $TCBCONTROLA_U$ ,  $TCBCONTROLA_K$ ,  $TCBCONTROLA_E$ , and  $TCBCONTROLA_{DM}$ . (If it is known that the particular process is a user-level process, then it would be sufficient to only assert  $TCBCONTROLA_U$  for example). When using the EJTAG hardware triggers to turn trace on and off, it is best if  $TCBCONTROLA_{On}$  is asserted and all the other processor mode selection bits in  $TCBCONTROLA$  are turned off. This would be the least confusing way to control tracing with the trigger signals. Tracing can be controlled via software with the *TraceControl* register in a similar manner.

### 16.11.8.3 Turning Off PDtrace™ Trace

Trace is turned off when the following expression evaluates true:

```
(
  (TraceControlTS and (not TraceControlOn)) or
  ((not TraceControlTS) and (not TCBCONTROLAOn))
)
or
(
  (not MatchEnable)      and
  (not TriggerEnable)    and
  (not FilterDataTraceActive) and
  TriggerDisable
)
```

where,

```
TriggerDisable ←
(
  DBCiTE      and
  DBSBS[i]    and
  TraceBPCDE  and
  (TraceBPCDBPON[i] = 0)
)
or
(
  IBCiTE      and
  IBSBS[i]    and
  TraceBPCIE  and
  (TraceBPCIBPON[i] = 0)
)
```

Tracing can be unconditionally turned off by de-asserting the  $TraceControl_{On}$  bit or the  $TCBCONTROLA_{On}$  signal. When either of these are asserted, tracing can be turned off if all of the enables are de-asserted, irrespective of the  $TraceControl_G$  bit ( $TCBCONTROLA_G$ ) and  $TraceControl_{ASID}$  ( $TCBCONTROLA_{ASID}$ ) values. EJTAG hardware break-



points can be used to trigger trace off as well. Note that if simultaneous triggers are generated, and even one of them turns on tracing, then even if all of the others attempt to trigger trace off, then tracing will still be turned on. This condition is reflected in presence of the “(not TriggerEnable)” term in the expression above.

## 16.12 PDtrace Cycle-by-Cycle Behavior

A key reason for using trace, and not single stepping to debug a software problem, is often to get a picture of the real-time behavior. However the trace logic itself can, when enabled, affect the exact cycle-by-cycle behavior,

### 16.12.1 FIFO Logic in PDtrace and TCB Modules

Both the PDtrace module and the TCB module contain a fifo. This might seem like extra overhead, but there are good reasons for this. The vast majority of the information compression happens in the PDtrace module. Any data information, like PC and load/store address values (delta or full), load/store data and processor mode changes, are sent on two 32-bit data busses to the TCB on the internal PDtrace™ interface. When an instruction requires more than 2x32 bits of information to be traced properly, the PDtrace fifo will buffer the information, and send it on subsequent clock cycles.

In the TCB, the on-chip trace memory is defined as a 128-bit wide synchronous memory running at core-clock speed. In this case the FIFO is not needed. For off-chip trace through the Trace Probe, the FIFO comes into play, because only a limited number of pins (16) exist. Also the speed of the Trace Probe interface can be different (either faster or slower) from that of the interAptiv core. So for off-chip tracing, a specific TCB TW FIFO is needed.

### 16.12.2 Handling of FIFO Overflow in the PDtrace Module

Depending on the amount of trace information selected for trace, and the frequency with which the 2x32-bit data interface is needed, it is possible for the PDtrace FIFO to overflow from time to time. There are two ways to handle this case:

1. Allow the overflow to happen, and thereby lose some information from the trace data.
2. Prevent the overflow by back-stalling the core until the FIFO has enough empty slots to accept new trace data.

The PDtrace fifo option is controlled by either the *TraceControl<sub>IO</sub>* or the *TCBCONTROL<sub>AIO</sub>* bit, depending on the setting of *TraceControl<sub>TS</sub>* bit.

The first option is free of any cycle-by-cycle change whether trace is turned on or not. This is achieved at the cost of potentially losing trace information. After an overflow, the fifo is completely emptied, and the next instruction is traced as if it was the start of the trace (processor mode and full PC are traced). This guarantees that only the un-traced fifo information is lost.

The second option guarantees that all the trace information is traced to the TCB. In some cases this is then achieved by back-stalling the core pipeline, giving the PDtrace fifo time to empty enough room in the fifo to accept new trace information from a new instruction. This option can obviously change the real-time behavior of the core when tracing is turned on.

If PC trace information is the only thing enabled (in *TraceControl<sub>2MODE</sub>* or *TCBCONTROL<sub>LCMODE</sub>*, depending on the setting of *TraceControl<sub>TS</sub>*), and Trace of all branches is turned off (via *TraceControl<sub>TB</sub>* or *TCBCONTROL<sub>ATB</sub>*, depending on the setting of *TraceControl<sub>TS</sub>*), then the fifo is unlikely to overflow very often, if at all. This is of course very dependent on the code executed, and the frequency of exception handler jumps, but with this setting there is very little information overhead.

### 16.12.3 Handling of FIFO Overflow in the TCB

The TCB also holds a fifo, used to buffer the TW's which are sent off-chip through the Trace Probe. The data width of the probe is 16 pins. The speed of these data pins can range from the core-clock speed to 1/10th of the core clock speed (the trace probe clock always runs at a double data rate multiple to the core-clock). See [Section 16.12.3.1 “Probe Width and Clock-ratio Settings”](#) for a description of probe width and clock-ratio options. The combination between the probe width and the data speed allows for different data rates through the trace probe. The high extreme is not likely to be supported in any implementation, but the low one might be.

The data rate is an important figure when the likelihood of a TCB fifo overflow is considered. The TCB will at maximum produce two 64-bit trace words per core-clock cycle. This is true for any selection of trace mode in *TraceControl2*<sub>MODE</sub> or *TCBCONTROL*<sub>MODE</sub>. The PDtrace module will guarantee the limited amount of data. If the TCB data rate cannot be matched by the off-chip probe width and data speed, then the TCB fifo can possibly overflow. Similar to the PDtrace module FIFO, this can be handled in two ways:

1. Allow the overflow to happen, and thereby lose some information from the trace data.
2. Prevent the overflow by asserting a stall-signal back to the core (*PDI\_StallSending*). This will in turn stall the core pipeline.

As a practical matter, the amount of data to the TCB can be minimized by only tracing PC information and excluding any cycle accurate information. This is explained in [Section 16.12.2 “Handling of FIFO Overflow in the PDtrace Module”](#) and below in [Section 16.12.4 “Adding Cycle Accurate Information to the Trace”](#). With this setting, a data rate of 8-bits per core-clock cycle is usually sufficient. No guarantees can be given here, however, as heavy interrupt activity can increase the number of unpredictable jumps considerably.

#### 16.12.3.1 Probe Width and Clock-ratio Settings

Note: the registers called out in this section are located in the Coherence Manager TAP described in Chapter 15, Multi-CPU Debug. All of these fields are reserved in the interAptiv core TAP registers.

The actual number of data pins (16) is defined by the CM TAP *TCBCONFIG*<sub>PW</sub> field. Furthermore, the frequency of the Trace Probe can be different from the core-clock frequency. The trace clock (*TR\_CLK*) is a double data rate clock. This means that the data pins (*TR\_DATA*) change their value on both edges of the trace clock. When the trace clock is running at clock ratio of 1:2 (one half) of core clock, the data output registers are running a core-clock frequency. The clock ratio is set in the CM TAP *TCBCONTROL*<sub>CR</sub> field. The legal range for the clock ratio is defined in CM TAP *TCBCONFIG*<sub>CRMax</sub> and CM TAP *TCBCONFIG*<sub>CRMin</sub> (both values inclusive). If the CM TAP *TCBCONTROL*<sub>CR</sub> bit is set to an unsupported value, the result is UNPREDICABLE.

### 16.12.4 Adding Cycle Accurate Information to the Trace

Depending on the trace regeneration software, it is possible to obtain the exact cycle time relationship between each instruction in the trace. This information is added to the trace, when the *TCBCONTROL*<sub>CA</sub> bit is set. The overhead on the trace information is a little more than one extra bit per core-clock cycle.

This setting only affects the TCB module and not the PDtrace module. The extra bit therefore only affects the likelihood of the TCB FIFO overflowing.

## 16.13 PC Sampling

The PC sampling feature enables sampling of the PC value periodically. This information can be used for statistical profiling of the program akin to gprof. This information is also very useful for detecting hot-spots in the code.

In PC sampling, the PC is sampled periodically and sent to the TAP register. Note that although the PC sampling function can be used both with and without a probe, if a probe is not connected, the sampled information cannot be read out since the TAP registers can only be read when a probe is connected. Therefore, MIPS recommends using the PC sampling capability only when a probe is connected.

The presence or absence of the PC Sampling feature is available in the Debug Control register as bit 9 (PCS). The sampled PC values are written into a TAP register. The old value in the TAP register is overwritten by a new value even if this register has not been read out by the debug probe. The sample rate is specified in a manner similar to the PDtrace synchronization period, with three bits. These bits in the Debug Control register are 8:6 and called PCSR (PC Sample Rate). These three bits take the value  $2^5$  to  $2^{12}$  similar to SyncPeriod. Note that the processor samples PC even when it is asleep, that is, in a WAIT state. This permits an analysis of the amount of time spent by a processor in WAIT state which may be used for example to revert to a low power mode during the non-execution phase of a real-time application.

The sampled values include a new data bit, the PC, the ASID of the sampled PC as well as the Enhanced Virtual Address (EVA) K/U bit. Figure 16.6 shows the format of the sampled values in the TAP register PCsample. The new data bit is used by the probe to determine if the PCsample register data just read out is new or already been read and must be discarded.

**Figure 16.6 TAP Register PCsample Format**

49	42	41	40	33	32	1	0
TC		K/U	ASID		PC		New

The sampled PC value is the PC of the graduating instruction in the current cycle. If the processor is stalled when the PC sample counter overflows, then the sampled PC is the PC of the next graduating instruction. The processor continues to sample the PC value even when it is in Debug mode.

### 16.13.1 PC Sampling in Wait State

When the processor is in a WAIT state to save power for example, an external agent might want to know how long it stays in the WAIT state. But counting cycles to update the PC sample value is a waste of power. Hence, when in a WAIT state, the processor must simply switch the New bit to 1 every time it is set to 0 by the probe hardware. Hence, the external agent or probe reading the PC value will detect a WAIT instruction for as long as the processor remains in the WAIT state. When the processor leaves the WAIT state, then counting is resumed as before.

## 16.14 EJTAG Registers

The following subsections describe the EJTAG register interface.

### 16.14.1 General Purpose Control and Status

The following registers provide general control and status information for EJTAG.

### 16.14.1.1 Debug Control Register

The Debug Control Register (*DCR*) register controls and provides information about debug issues and is always provided with the interAptiv core. The register is memory-mapped in drseg at offset 0x0.

The DataBrk and InstBrk bits indicate if hardware breakpoints are included in the implementation, and debug software is expected to read hardware breakpoint registers for additional information.

Hardware and software interrupts are maskable for non-debug mode with the INTE bit, which works in addition to the other mechanisms for interrupt masking and enabling. NMI is maskable in non-debug mode with the NMIE bit, and a pending NMI is indicated through the NMIP bit.

The SRE bit allows implementation dependent masking of some sources for reset. The interAptiv core does not distinguish between soft and hard reset, but typically only soft reset sources in the system would be maskable and hard sources such as the reset switch would not be. The soft reset masking should only be applied to a soft reset source if that source can be efficiently masked in the system, thus resulting in no reset at all. If that is not possible, then that soft reset source should not be masked, since a partial soft reset may cause the system to fail or hang. There is no automatic indication of whether the SRE is effective, so the user must consult system documentation.

The PE bit reflects the ProbEn bit from the EJTAG Control register (*ECR*), whereby the probe can indicate to the debug software that the probe will service dmseg accesses. The reset value in the table below takes effect on any CPU reset.

**Figure 16.7 Debug Control Register Format**

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0	ENM	0	PCIM	PCno ASID	DASQ	DASe	DAS	0				FDC Impl	Data Brk	Inst Brk	
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IVM	DVM	0		RD Vec	CBT	PCS	PCR		PCSe	IntE	NMIE	NMI pend	SRstE	Prob En	

**Table 16.10 Debug Control Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
0	31:30	Must be written as zeros; return zeros on reads.	0	0						
ENM	29	Endianness in which the processor is running in kernel and Debug Mode: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Little endian</td> </tr> <tr> <td>1</td> <td>Big endian</td> </tr> </tbody> </table>	Encoding	Meaning	0	Little endian	1	Big endian	R	Preset
Encoding	Meaning									
0	Little endian									
1	Big endian									
0	28:27	Must be written as zeros; return zeros on reads.	0	0						

**Table 16.10 Debug Control Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
PCIM	26	<p>Configure PC Sampling to capture all executed addresses or only those that miss the instruction cache. This feature is not supported and this bit will read as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>All PC's captured</td> </tr> <tr> <td>1</td> <td>Capture only PC's that miss the cache.</td> </tr> </tbody> </table>	Encoding	Meaning	0	All PC's captured	1	Capture only PC's that miss the cache.	R	0
Encoding	Meaning									
0	All PC's captured									
1	Capture only PC's that miss the cache.									
PCnoASID	25	<p>Controls whether the PCSAMPLE scan chain includes or omits the ASID field. ASID is always included so this bit will read as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>ASID included in PCSAMPLE scan</td> </tr> <tr> <td>1</td> <td>ASID omitted from PCSAMPLE scan</td> </tr> </tbody> </table>	Encoding	Meaning	0	ASID included in PCSAMPLE scan	1	ASID omitted from PCSAMPLE scan	R	0
Encoding	Meaning									
0	ASID included in PCSAMPLE scan									
1	ASID omitted from PCSAMPLE scan									
DASQ	24	<p>Qualifies Data Address Sampling using a data breakpoint. Data address sampling is not supported so this bit will read as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>All data addresses are sampled</td> </tr> <tr> <td>1</td> <td>Sample matches of data breakpoint 0</td> </tr> </tbody> </table>	Encoding	Meaning	0	All data addresses are sampled	1	Sample matches of data breakpoint 0	R	0
Encoding	Meaning									
0	All data addresses are sampled									
1	Sample matches of data breakpoint 0									
DASe	23	<p>Enables Data Address Sampling. Data address sampling is not supported so this bit will read as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Data Address sampling disabled.</td> </tr> <tr> <td>1</td> <td>Data Address sampling enabled.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Data Address sampling disabled.	1	Data Address sampling enabled.	R	0
Encoding	Meaning									
0	Data Address sampling disabled.									
1	Data Address sampling enabled.									
DAS	22	<p>Indicates if the Data Address Sampling feature is implemented. Data address sampling is not supported so this bit will read as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No DA Sampling implemented</td> </tr> <tr> <td>1</td> <td>DA Sampling implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No DA Sampling implemented	1	DA Sampling implemented	R	0
Encoding	Meaning									
0	No DA Sampling implemented									
1	DA Sampling implemented									
0	21:19	Must be written as zeros; return zeros on reads.	0	0						

**Table 16.10 Debug Control Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
FDCImpl	18	Indicates if the fast debug channel is implemented <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No fast debug channel implemented</td> </tr> <tr> <td>1</td> <td>Fast debug channel implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No fast debug channel implemented	1	Fast debug channel implemented	R	1
Encoding	Meaning									
0	No fast debug channel implemented									
1	Fast debug channel implemented									
DataBrk	17	Indicates if data hardware breakpoint is implemented: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No data hardware breakpoint implemented</td> </tr> <tr> <td>1</td> <td>Data hardware breakpoint implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No data hardware breakpoint implemented	1	Data hardware breakpoint implemented	R	Preset
Encoding	Meaning									
0	No data hardware breakpoint implemented									
1	Data hardware breakpoint implemented									
InstBrk	16	Indicates if instruction hardware breakpoint is implemented: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No instruction hardware breakpoint implemented</td> </tr> <tr> <td>1</td> <td>Instruction hardware breakpoint implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No instruction hardware breakpoint implemented	1	Instruction hardware breakpoint implemented	R	Preset
Encoding	Meaning									
0	No instruction hardware breakpoint implemented									
1	Instruction hardware breakpoint implemented									
IVM	15	Indicates if inverted data value match on data hardware breakpoints is implemented: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No inverted data value match on data hardware breakpoints implemented</td> </tr> <tr> <td>1</td> <td>Inverted data value match on data hardware breakpoints implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No inverted data value match on data hardware breakpoints implemented	1	Inverted data value match on data hardware breakpoints implemented	R	0
Encoding	Meaning									
0	No inverted data value match on data hardware breakpoints implemented									
1	Inverted data value match on data hardware breakpoints implemented									
DVM	14	Indicates if a data value store on a data value breakpoint match is implemented: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No data value store on a data value breakpoint match implemented</td> </tr> <tr> <td>1</td> <td>Data value store on a data value breakpoint match implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No data value store on a data value breakpoint match implemented	1	Data value store on a data value breakpoint match implemented	R	0
Encoding	Meaning									
0	No data value store on a data value breakpoint match implemented									
1	Data value store on a data value breakpoint match implemented									
0	13:12	Must be written as zeros; return zeros on reads.	0	0						
RDVec	11	Enables relocation of the debug exception vector. The value in the DebugVectorAddr register is used for EJTAG exceptions when ProbTrap = 0, and RDVec = 1.	R/W	0						

**Table 16.10 Debug Control Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
CBT	10	Indicates if complex breakpoint block is implemented: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No complex breakpoint block implemented</td> </tr> <tr> <td>1</td> <td>Complex breakpoint block implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No complex breakpoint block implemented	1	Complex breakpoint block implemented	R	0
Encoding	Meaning									
0	No complex breakpoint block implemented									
1	Complex breakpoint block implemented									
PCS	9	Indicates if the PC Sampling feature is implemented. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No PC Sampling implemented</td> </tr> <tr> <td>1</td> <td>PC Sampling implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No PC Sampling implemented	1	PC Sampling implemented	R	1
Encoding	Meaning									
0	No PC Sampling implemented									
1	PC Sampling implemented									
PCR	8:6	PC Sampling rate. Values 0 to 7 map to values $2^5$ to $2^{12}$ cycles, respectively. That is, a PC sample is written out every 32, 64, 128, 256, 512, 1024, 2048, or 4096 cycles respectively. The external probe or software is allowed to set this value to the desired sample rate.	R/W	7						
PCSe	5	If the PC sampling feature is implemented, then indicates whether PC sampling is initiated or not. That is, a value of 0 indicates that PC sampling is not enabled and when the bit value is 1, then PC sampling is enabled and the counters are operational.	R/W	0						
IntE	4	Hardware and software interrupt enable for Non-Debug Mode, in conjunction with other disable mechanisms: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt disabled</td> </tr> <tr> <td>1</td> <td>Interrupt enabled depending on other enabling mechanisms</td> </tr> </tbody> </table>	Encoding	Meaning	0	Interrupt disabled	1	Interrupt enabled depending on other enabling mechanisms	R/W	1
Encoding	Meaning									
0	Interrupt disabled									
1	Interrupt enabled depending on other enabling mechanisms									
NMIE	3	Non-Maskable Interrupt (NMI) enable for Non-Debug Mode: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>NMI disabled</td> </tr> <tr> <td>1</td> <td>NMI enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	NMI disabled	1	NMI enabled	R/W	1
Encoding	Meaning									
0	NMI disabled									
1	NMI enabled									
NMIpend	2	Indication for pending NMI: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No NMI pending</td> </tr> <tr> <td>1</td> <td>NMI pending</td> </tr> </tbody> </table>	Encoding	Meaning	0	No NMI pending	1	NMI pending	R	0
Encoding	Meaning									
0	No NMI pending									
1	NMI pending									

**Table 16.10 Debug Control Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
SRstE	1	Controls soft reset enable: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Soft reset masked for soft reset sources dependent on implementation</td> </tr> <tr> <td>1</td> <td>Soft reset is fully enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	Soft reset masked for soft reset sources dependent on implementation	1	Soft reset is fully enabled	R/W	1
Encoding	Meaning									
0	Soft reset masked for soft reset sources dependent on implementation									
1	Soft reset is fully enabled									
ProbEn	0	Indicates value of the ProbEn value in the ECR register: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No access should occur to the dmseg segment</td> </tr> <tr> <td>1</td> <td>Probe services accesses to the dmseg segment</td> </tr> </tbody> </table> Bit is read-only (R) and reads as zero if not implemented.	Encoding	Meaning	0	No access should occur to the dmseg segment	1	Probe services accesses to the dmseg segment	R	Same value as ProbEn in ECR
Encoding	Meaning									
0	No access should occur to the dmseg segment									
1	Probe services accesses to the dmseg segment									

#### 16.14.1.2 DebugVectorAddr Register

This register allows an alternate debug exception vector address to be specified, which can enable placing a debug monitor program into RAM for much faster execution than the default ROM address. This register is memory mapped at an offset of 0x00020 within the DRSEG memory segment.

Figure 16.8 shows the register format and Table 16.12 describes the fields in this register.


Table 16.11 shows the different debug exception vector locations that are possible.

**Table 16.11 Debug Exception Vectors**

ECR <sub>ProbTrap</sub>	DCR <sub>RdVec</sub>	Config <sub>5K</sub>	SI_UseExceptionBase	Cache Error?	Debug Exception Vector
1	x	x	x	x	0xFF20_0200
0	1	0	x	0	2'b10    DebugVectorAddr[29:0]
0	1	1	x	0	DebugVectorAddr[31:0]
0	1	0	x	1	3'b101    DebugVectorAddr[28:0]
0	1	1	x	1	DebugVectorAddr[31:0]
0	0	0	1	0	2'b10    SI_ExceptionBase[29:12]    0x480
0	0	1	1	0	SI_ExceptionBase[31:12]    0x480
0	0	0	1	1	3'b101    SI_ExceptionBase[28:12]    0x480
0	0	1	1	1	SI_ExceptionBase[31:12]    0x480
0	0	x	0	x	0xBF0_0480



**Figure 16.8 DebugVectorAddr Register Format**



**Table 16.12 DebugVectorAddr Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DebugVectorAddr	31	Programmable Debug Exception Vector Address. Note that bits 31:30 have default values of 1 and 0 respectively and can only be written when the WG bit is set. If the WG bit is cleared, these bits are read-only and retain their previous values. These two bits can be written whenever the WG bit is set, regardless of the state of <i>Config5<math>\kappa</math></i> .	R/W	1
	30		R/W	0
	29:7		R/W	0x7f8009 (corresponds to 0xbfc00480)
WG	6	Write gate.  When the WG bit is set, the DebugVectorAddr field is expanded to include bits 31:30 to facilitate programmable memory segmentation controlled by the <i>SegCtl0</i> through <i>SegCtl2</i> registers.  When the WG bit is cleared, bits 31:30 of this register are not writeable and remain unchanged from the last time that WG was cleared.	R/W	Externally Set
0	5:1	Ignored on write, returns zero on read.	R	0
ISA	0	ISA mode to be used for debug exception handler. Only used on cores implementing microMIPS.	R	0

### 16.14.2 Instruction Breakpoint Registers

The registers for instruction breakpoints are described below. These registers have implementation information and are used to set up the instruction breakpoints. All registers are in drseg with addresses as shown in [Table 16.13](#).

**Table 16.13 Addresses for Instruction Breakpoint Registers**

Offset in drseg	Register Mnemonic	Register Name and Description
0x1000	<i>IBS</i>	Instruction Breakpoint Status
0x1100 + n * 0x100	<i>IBAn</i>	Instruction Breakpoint Address n
0x1108 + n * 0x100	<i>IBMn</i>	Instruction Breakpoint Address Mask n
0x1110 + n * 0x100	<i>IBASIDn</i>	Instruction Breakpoint ASID n
0x1118 + n * 0x100	<i>IBCn</i>	Instruction Breakpoint Control n
n is breakpoint number in range 0 to 3		

An example of some of the registers; *IBA0* is at offset 0x1100 and *IBC2* is at offset 0x1318.


### 16.14.2.1 Instruction Breakpoint Status (IBS) Register

**Compliance Level:** Implemented only if instruction breakpoints are implemented.

The Instruction Breakpoint Status (*IBS*) register holds implementation and status information about the instruction breakpoints.

The ASID applies to all the instruction breakpoints.

**Figure 16.9 IBS Register Format**



**Table 16.14 IBS Register Field Descriptions**


Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
Res	31	Must be written as zero; returns zero on read.	R	0						
ASIDsup	30	Hardware and software interrupt enable for Non-Debug Mode, in conjunction with other disable mechanisms: <table border="1" style="margin: 5px auto; border-collapse: collapse;"> <thead> <tr> <th style="width: 10%;">Encoding</th> <th style="width: 80%;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>ASID compare not supported</td> </tr> <tr> <td style="text-align: center;">1</td> <td>ASID compare supported (IBASIDn register implemented)</td> </tr> </tbody> </table>	Encoding	Meaning	0	ASID compare not supported	1	ASID compare supported (IBASIDn register implemented)	R	1
Encoding	Meaning									
0	ASID compare not supported									
1	ASID compare supported (IBASIDn register implemented)									
Res	29:28	Must be written as zero; returns zero on read.	R	0						
BCN	27:24	Number of instruction breakpoints implemented.	R	2 or 4						
Res	23:4	Must be written as zero; returns zero on read.	R	0						
BS	3:0	Break status for breakpoint <i>n</i> is at BS[ <i>n</i> ], with <i>n</i> from 0 to 3. The bit is set to 1 when the corresponding breakpoint is enabled and the condition has matched. If only two instruction breakpoints are implemented, bits 2 and 3 must be written as zero and will return zero on read.	R/W	Undefined						

### 16.14.2.2 Instruction Breakpoint Address *n* (IBAn) Register

**Compliance Level:** Implemented only for implemented instruction breakpoints.

The Instruction Breakpoint Address *n* (*IBAn*) register has the address used in the condition for instruction breakpoint *n*.

**Figure 16.10 IBAn Register Format**



**Table 16.15 IBAn Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bit(s)			
IBA	31:1	Instruction breakpoint address for condition.	R/W	Undefined
ISA	0	Instruction breakpoint ISA mode for condition	R/W	Undefined

### 16.14.2.3 Instruction Breakpoint Address Mask n (IBMn) Register

**Compliance Level:** Implemented only for implemented instruction breakpoints.

The Instruction Breakpoint Address Mask *n* (*IBMn*) register has the mask for the address compare used in the condition for instruction breakpoint *n*.

**Figure 16.11 IBMn Register Format**



**Table 16.16 IBMn Register Field Descriptions**


Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
IBM	31:1	Instruction breakpoint address mask for condition: <table border="1" style="margin-left: 20px; margin-top: 10px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>Corresponding address bit not masked</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Corresponding address bit masked</td> </tr> </tbody> </table>	Encoding	Meaning	0	Corresponding address bit not masked	1	Corresponding address bit masked	R/W	Undefined
Encoding	Meaning									
0	Corresponding address bit not masked									
1	Corresponding address bit masked									
ISAM	0	Instruction breakpoint ISA mode mask for condition:                     condition: <table border="1" style="margin-left: 20px; margin-top: 10px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>Corresponding address bit not masked</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Corresponding address bit masked</td> </tr> </tbody> </table> <p style="margin-left: 20px; margin-top: 10px;">0: ISA mode considered for match condition 1: ISA mode masked</p>	Encoding	Meaning	0	Corresponding address bit not masked	1	Corresponding address bit masked	R/W	Undefined
Encoding	Meaning									
0	Corresponding address bit not masked									
1	Corresponding address bit masked									

### 16.14.2.4 Instruction Breakpoint ASID n (IBASIDn) Register

**Compliance Level:** Implemented only for implemented instruction breakpoints.

This register is used to define an ASID value to be used in the match expression.

**Figure 16.12 IBASIDn Register Format**



**Table 16.17 IBASIDn Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Res	31:8	Must be written as zero; returns zero on read.	R	0
ASID	7:0	Instruction breakpoint ASID value for a compare.	R/W	Undefined

**16.14.2.5 Instruction Breakpoint Control n (IBCN) Register**

**Compliance Level:** Implemented only for implemented instruction breakpoints.

The Instruction Breakpoint Control *n* (IBCN) register controls the setup of instruction breakpoint *n*.

**Figure 16.13 IBCn Register Format**



**Table 16.18 IBCn Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
TC	31:24	The value of TC (thread context) to match in the comparison to determine if the instruction break is to be taken. TC value is ignored if TCuse is set to 0	R/W	Undefined						
ASIDuse	23	Use ASID value in compare for instruction breakpoint n: <table border="1" style="margin: 5px auto; border-collapse: collapse;"> <thead> <tr> <th style="width: 100px;">Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>Don't use ASID value in compare</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Use ASID value in compare</td> </tr> </tbody> </table>	Encoding	Meaning	0	Don't use ASID value in compare	1	Use ASID value in compare	R/W	Undefined
Encoding	Meaning									
0	Don't use ASID value in compare									
1	Use ASID value in compare									
TCUse	22	Must be written as zero; returns zero on read.	R	0						
TE	2	Trigger-only Enable. This field is ignored when BE is set. When BE is cleared and TE is set, instruction breakpoint n is enabled, but will not signal a debug exception.	R/W	0						
R	1	Must be written as zero; returns zero on read.	R	0						
BE	0	Breakpoint Enable. When set, instruction breakpoint n is enabled and will signal a debug exception when its condition matches.	R/W	0						

### 16.14.3 Data Breakpoint Registers

The registers for data breakpoints are described below. These registers have implementation information and are used to setup the data breakpoints. All registers are in drseg, and the addresses are shown in [Table 16.19](#).

**Table 16.19 Addresses for Data Breakpoint Registers**

Offset in drseg	Register Mnemonic	Register Name and Description
0x2000	<i>DBS</i>	Data Breakpoint Status
0x2100 + 0x100 * n	<i>DBAn</i>	Data Breakpoint Address n
0x2108 + 0x100 * n	<i>DBMn</i>	Data Breakpoint Address Mask n
0x2110 + 0x100 * n	<i>DBASIDn</i>	Data Breakpoint ASID n
0x2118 + 0x100 * n	<i>DBCn</i>	Data Breakpoint Control n
0x2120 + 0x100 * n	<i>DBVn</i>	Data Breakpoint Value n
0x2124 + 0x100 * n	<i>DBVHn</i>	Data Breakpoint Value High n

n is breakpoint number as 0 or 1

An example of some of the registers; *DBM0* is at offset 0x2108 and *DBV1* is at offset 0x2220.


#### 16.14.3.1 Data Breakpoint Status (DBS) Register

**Compliance Level:** Implemented if data breakpoints are implemented.

The Data Breakpoint Status (*DBS*) register holds implementation and status information about the data breakpoints.

The ASIDsup field indicates whether ASID compares are supported.

**Figure 16.14 DBS Register Format**



**Table 16.20 DBS Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
Res	31	Must be written as zero; returns zero on read.	R	0						
ASID	30	Indicates that ASID compares are supported in data breakpoints. n: <table border="1" style="margin-left: 40px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Don't use ASID value in compare</td> </tr> <tr> <td>1</td> <td>Use ASID value in compare</td> </tr> </tbody> </table> 0: Not supported 1: Supported	Encoding	Meaning	0	Don't use ASID value in compare	1	Use ASID value in compare	R	TLB MMU - 1
Encoding	Meaning									
0	Don't use ASID value in compare									
1	Use ASID value in compare									

**Table 16.20 DBS Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Res	29:28	Must be written as zero; returns zero on read.	R	0
BCN	27:24	Number of data breakpoints implemented.	R	1 or 2
Res	23:2	Must be written as zero; returns zero on read.	R	0
BS	1:0	Break status for breakpoint n is at BS[n], with n from 0 to 1. The bit is set to 1 when the condition for the corresponding breakpoint has matched and the condition has matched. If only one data breakpoint is implemented, bit 1 must be written as 0 and will return 0 on reads.	R/W0	Undefined

### 16.14.3.2 Data Breakpoint Address n (DBAn) Register

**Compliance Level:** Implemented only for implemented data breakpoints.

The Data Breakpoint Address n (*DBAn*) register has the address used in the condition for data breakpoint n.

**Figure 16.15 DBAn Register Format**



**Table 16.21 DBAn Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DBA	31:0	Data breakpoint address for condition.	R/W	Undefined

### 16.14.3.3 Data Breakpoint Address Mask n (DBMn) Register

**Compliance Level:** Implemented only for implemented data breakpoints.

The Data Breakpoint Address Mask n (*DBMn*) register has the mask for the address compare used in the condition for data breakpoint n.

**Figure 16.16 DBMn Register Format**



**Table 16.22 DBMn Register Field Descriptions**


Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
DBM	31:0	Data breakpoint address mask for condition: <i>n</i> : <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Don't use ASID value in compare</td> </tr> <tr> <td>1</td> <td>Use ASID value in compare</td> </tr> </tbody> </table> 0: Corresponding address bit not masked 1: Corresponding address bit masked	Encoding	Meaning	0	Don't use ASID value in compare	1	Use ASID value in compare	R/W	Undefined
Encoding	Meaning									
0	Don't use ASID value in compare									
1	Use ASID value in compare									

**16.14.3.4 Data Breakpoint ASID n (DBASIDn) Register**

**Compliance Level:** Implemented only for implemented data breakpoints.

For processors with a TLB based MMU, this register is used to define an ASID value to be used in the match expression. For cores with the FM MMU, this register is reserved and reads as 0.

**Figure 16.17 DBASIDn Register Format**



**Table 16.23 DBASIDn Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Res	31:8	Must be written as zero; returns zero on read.	R	0
ASID	7:0	Data breakpoint ASID value for compares.	R/W	Undefined

**16.14.3.5 Data Breakpoint Control n (DBCn) Register**

**Compliance Level:** Implemented only for implemented data breakpoints.

The Data Breakpoint Control *n* (DBC*n*) register controls the setup of data breakpoint *n*.

**Figure 16.18 DBCn Register Format**



**Table 16.24 DBCn Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
TC	31:24	The value of TC to match in the comparison to determine if data break is to be taken. TC value is ignored if TCuse is set to 0	R/W	Undefined						
ASIDuse	23	Use ASID value in compare for data breakpoint n: <table border="1" data-bbox="522 445 1024 564"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Don't use ASID value in compare</td> </tr> <tr> <td>1</td> <td>Use ASID value in compare</td> </tr> </tbody> </table>	Encoding	Meaning	0	Don't use ASID value in compare	1	Use ASID value in compare	R/W	Undefined
Encoding	Meaning									
0	Don't use ASID value in compare									
1	Use ASID value in compare									
TCuse	22	Use TC value in comparison for data breakpoint n: 0: Do not use TC value in the compare. 1: Use TC value in the compare.	R/W	Undefined						
BAI	21:14	Byte access ignore controls ignore of access to a specific byte. <i>BAI</i> [0] ignores access to byte at bits [7:0] of the data bus, <i>BAI</i> [1] ignores access to byte at bits [15:8], etc.: <table border="1" data-bbox="522 793 1024 968"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Condition depends on access to corresponding byte</td> </tr> <tr> <td>1</td> <td>Access for corresponding byte is ignored</td> </tr> </tbody> </table>	Encoding	Meaning	0	Condition depends on access to corresponding byte	1	Access for corresponding byte is ignored	R/W	Undefined
Encoding	Meaning									
0	Condition depends on access to corresponding byte									
1	Access for corresponding byte is ignored									
NoSB	13	Controls if condition for data breakpoint is fulfilled on a store transaction: <table border="1" data-bbox="522 1077 1024 1251"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Condition may be fulfilled on store transaction</td> </tr> <tr> <td>1</td> <td>Condition is never fulfilled on store transaction</td> </tr> </tbody> </table>	Encoding	Meaning	0	Condition may be fulfilled on store transaction	1	Condition is never fulfilled on store transaction	R/W	Undefined
Encoding	Meaning									
0	Condition may be fulfilled on store transaction									
1	Condition is never fulfilled on store transaction									
NoLB	12	Controls if condition for data breakpoint is fulfilled on a load transaction: <table border="1" data-bbox="522 1360 1024 1535"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Condition may be fulfilled on load transaction</td> </tr> <tr> <td>1</td> <td>Condition is never fulfilled on load transaction</td> </tr> </tbody> </table>	Encoding	Meaning	0	Condition may be fulfilled on load transaction	1	Condition is never fulfilled on load transaction	R/W	Undefined
Encoding	Meaning									
0	Condition may be fulfilled on load transaction									
1	Condition is never fulfilled on load transaction									
BLM	11:4	Byte lane mask for value compare on data breakpoint. <i>BLM</i> [0] masks byte at bits [7:0] of the data bus, <i>BLM</i> [1] masks byte at bits [15:8], etc.: <table border="1" data-bbox="522 1675 1024 1795"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Compare corresponding byte lane</td> </tr> <tr> <td>1</td> <td>Mask corresponding byte lane</td> </tr> </tbody> </table>	Encoding	Meaning	0	Compare corresponding byte lane	1	Mask corresponding byte lane	R/W	Undefined
Encoding	Meaning									
0	Compare corresponding byte lane									
1	Mask corresponding byte lane									
R	3	Must be written as zero; returns zero on reads.	R	0						



**Table 16.24 DBCn Register Field Descriptions(continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
TE	2	Trigger-only Enable. This field is ignored when <i>BE</i> is set. When <i>BE</i> is cleared and TE is set, data breakpoint n is enabled, but will not signal a debug exception.	R/W	0
R	1	Must be written as zero; returns zero on reads.	R	0
BE	0	Breakpoint Enable. When set, data breakpoint n is enabled and will signal a debug exception when its condition matches.	R/W	0

### 16.14.3.6 Data Breakpoint Value n (DBVn) Register

**Compliance Level:** Implemented only for implemented data breakpoints.

The Data Breakpoint Value n (*DBVn*) register has the value used in the condition for data breakpoint n.

**Figure 16.19 DBVn Register Format**



**Table 16.25 DBVn Register Field Descriptions**


Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DBV	31:0	Data breakpoint value for condition.	R/W	Undefined

### 16.14.3.7 Data Breakpoint Value High n (DBVHn) Register

**Compliance Level:** Implemented only for implemented data breakpoints.

The Data Breakpoint Value High n (*DBVHn*) register has the value used in the condition for data breakpoint n.

**Figure 16.20 DBVHn Register Format**



**Table 16.26 DBVHn Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
DBVH	31:0	Data breakpoint value high for condition. This register provides the high order bits [63:32] for data value on double-word floating point loads and stores.	R/W	Undefined

## 16.14.4 EJTAG TAP Registers

The EJTAG TAP Module has one Instruction register and a number of data registers, all accessible through the TAP:

### 16.14.4.1 Instruction Register

The Instruction register is accessed when the TAP receives an Instruction register scan protocol. During an Instruction register scan operation the TAP controller selects the output of the Instruction register to drive the *TDO* pin. The shift register consists of a series of bits arranged to form a single scan path between *TDI* and *TDO*. During an Instruction register scan operations, the TAP controls the register to capture status information and shift data from *TDI* to *TDO*. Both the capture and shift operations occur on the rising edge of *TCK*. However, the data shifted out from the *TDO* occurs on the falling edge of *TCK*. In the Test-Logic-Reset and *Capture-IR* state, the instruction shift register is set to 00001<sub>2</sub>, as for the IDCODE instruction. This forces the device into the functional mode and selects the Device ID register. The Instruction register is 5 bits wide. The instruction shifted in takes effect for the following data register scan operation. A list of the implemented instructions are listed in [Table 16.8](#).

### 16.14.4.2 Data Registers Overview

The EJTAG uses several data registers that are arranged in parallel from the primary *TDI* input to the primary *TDO* output. The Instruction register supplies the address that allows one of the data registers to be accessed during a data register scan operation. During a data register scan operation, the addressed scan register receives TAP control signals to capture the register and shift data from *TDI* to *TDO*. During a data register scan operation, the TAP selects the output of the data register to drive the *TDO* pin. The register is updated in the *Update-DR* state with respect to the write bits.

This description applies in general to the following data registers:

- Bypass Register
- Device Identification Register
- Implementation Register
- EJTAG Control Register (ECR)
- Address Register
- Data Register
- FastData Register

### 16.14.4.3 Bypass Register


The *Bypass* register consists of a single scan register bit. When selected, the Bypass register provides a single bit scan path between *TDI* and *TDO*. The Bypass register allows abbreviating the scan path through devices that are not involved in the test. The Bypass register is selected when the Instruction register is loaded with a pattern of all ones to satisfy the IEEE 1149.1 Bypass instruction requirement.

### 16.14.4.4 Device Identification (ID) Register

The *Device Identification* register is defined by IEEE 1149.1, to identify the device's manufacturer, part number, revision, and other device-specific information. [Table 16.27](#) shows the bit assignments defined for the read-only Device Identification Register, and inputs to the core determine the value of these bits. These bits can be scanned out of the

ID register after being selected. The register is selected when the Instruction register is loaded with the IDCODE instruction. Note that this register contains only device manufacturer information and should not be used in an attempt to determine the EJTAG or PDTrace revisions of the device.

**Figure 16.21 Device Identification Register Format**




**Table 16.27 Device Identification Register**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Version	31:28	<b>Version</b> (4 bits) This field identifies the version number of the processor derivative.	R	<i>EJ_Version[3:0]</i>
PartNumber	27:12	<b>Part Number</b> (16 bits) This field identifies the part number of the processor derivative.	R	<i>EJ_PartNumber[15:0]</i>
ManufID	11:1	<b>Manufacturer Identity</b> (11 bits) Accordingly to IEEE 1149.1-1990, the manufacturer identity code shall be a compressed form of the JEDEC Publications 106-A.	R	<i>EJ_ManufID[10:0]</i>
R	0	reserved	R	1

**16.14.4.5 Implementation Register**

This 32-bit read-only register is used to identify the features of the EJTAG implementation. Some of the reset values are set by inputs to the core. The register is selected when the Instruction register is loaded with the IMPCODE instruction. The EJTAG probe uses this TAP register to determine the EJTAG version of the device. Software has no access to this register and must use the CP0 Debug register to determine the EJTAG version.

**Figure 16.22 Implementation Register Format**



**Table 16.28 Implementation Register Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
EJTAGver	31:29	Indicates EJTAG version 5.0.	R	5
reserved	28:25	Reserved. Must be written as zeros; returns zeros on reads.	R	0

**Table 16.28 Implementation Register Descriptions(continued)**

Fields		Description	Read / Write	Reset State										
Name	Bit(s)													
DINTsup	24	DINT Signal Supported from Probe This bit indicates if the DINT signal from the probe is supported:  <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>DINT signal from the probe is not supported.</td> </tr> <tr> <td>1</td> <td>Probe can use DINT signal to make debug interrupt.</td> </tr> </tbody> </table>	Encoding	Meaning	0	DINT signal from the probe is not supported.	1	Probe can use DINT signal to make debug interrupt.	R	<i>EJ_DINTsup</i>				
Encoding	Meaning													
0	DINT signal from the probe is not supported.													
1	Probe can use DINT signal to make debug interrupt.													
ASIDsize	23:21	Size of ASID field in implementation:  <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No ASID in implementation</td> </tr> <tr> <td>1</td> <td>Reserved</td> </tr> <tr> <td>2</td> <td>8-bit ASID</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	0	No ASID in implementation	1	Reserved	2	8-bit ASID	3	Reserved	R	2
Encoding	Meaning													
0	No ASID in implementation													
1	Reserved													
2	8-bit ASID													
3	Reserved													
R	20:17	Reserved	R	0										
MIPS16	16	Indicates whether MIPS16 is implemented:  <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No MIPS16 support</td> </tr> <tr> <td>1</td> <td>MIPS16 implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	No MIPS16 support	1	MIPS16 implemented	R	1				
Encoding	Meaning													
0	No MIPS16 support													
1	MIPS16 implemented													
R	15	Reserved. Must be written as zeros; returns zeros on reads.	R	0										
NoDMA			R	1										
TYPE	13:11	Indicates what type of entity is associated with this TAP and whether the TypeInfo field exists.  <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>TYPEINFO field not implemented. Legacy value - probably attached to a CPU.</td> </tr> <tr> <td>1</td> <td>This TAP is attached to a CPU and the TYP-EINFO field reflects <i>EBase<sub>CPUNUM</sub></i>.</td> </tr> <tr> <td>2</td> <td>This TAP is attached to a Trace-Master and the TypeInfo field is not used.</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	0	TYPEINFO field not implemented. Legacy value - probably attached to a CPU.	1	This TAP is attached to a CPU and the TYP-EINFO field reflects <i>EBase<sub>CPUNUM</sub></i> .	2	This TAP is attached to a Trace-Master and the TypeInfo field is not used.	3	Reserved	R	1
Encoding	Meaning													
0	TYPEINFO field not implemented. Legacy value - probably attached to a CPU.													
1	This TAP is attached to a CPU and the TYP-EINFO field reflects <i>EBase<sub>CPUNUM</sub></i> .													
2	This TAP is attached to a Trace-Master and the TypeInfo field is not used.													
3	Reserved													

**Table 16.28 Implementation Register Descriptions(continued)**

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
TYPEINFO	10:1	Identifier information specific to the type of entity associated with this TAP. The attached entity is specified by the TYPE field. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>CPU</td> <td>Reflects <math>EBase_{CPUNUM}</math> of the associated CPU.</td> </tr> <tr> <td>Others</td> <td>Reserved.</td> </tr> </tbody> </table>	Encoding	Meaning	CPU	Reflects $EBase_{CPUNUM}$ of the associated CPU.	Others	Reserved.	R	1
Encoding	Meaning									
CPU	Reflects $EBase_{CPUNUM}$ of the associated CPU.									
Others	Reserved.									
R	0	Reserved. Must be written as zeros; returns zeros on reads.	R	0						


#### 16.14.4.6 EJTAG Control Register

This 32-bit register controls the various operations of the TAP modules. This register is selected by shifting in the CONTROL instruction. Bits in the EJTAG Control register can be set/cleared by shifting in data; status is read by shifting out the contents of this register. This EJTAG Control register can only be accessed by the TAP interface.

The EJTAG Control register is not updated in the *Update-DR* state unless the Reset occurred (Rocc) bit 31, is either 0 or written to 0. This is in order to ensure proper handling of processor accesses.

The value used for reset indicated in the table below takes effect on CPU resets, but not on TAP controller resets (e.g. *TRST\_N*). *TCK* clock is not required when the CPU reset occurs, but the bits are still updated to the reset value when the *TCK* is supplied. The first 5 *TCK* clocks after CPU reset may result in reset of the bits, due to synchronization between clock domains.

**Figure 16.23 EJTAG Control Register Format**



**Table 16.29 EJTAG Control Register Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
Rocc	31	Reset Occurred The bit indicates if a CPU reset has occurred: <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No reset occurred since bit last cleared</td> </tr> <tr> <td>1</td> <td>Reset occurred since bit last cleared</td> </tr> </tbody> </table> <p>The Rocc bit will remain set to 1 as long as reset is applied. This bit must be cleared by the probe to acknowledge that the incident was detected.</p> <p>The EJTAG Control register is not updated in the <i>Update-DR</i> state unless Rocc is 0 or written to 0, in order to ensure proper handling of processor access following reset.</p>	Encoding	Meaning	0	No reset occurred since bit last cleared	1	Reset occurred since bit last cleared	R/W	1
Encoding	Meaning									
0	No reset occurred since bit last cleared									
1	Reset occurred since bit last cleared									

**Table 16.29 EJTAG Control Register Descriptions (continued)**

Fields		Description	Read / Write	Reset State																																	
Name	Bit(s)																																				
Psz[1:0]	30:29	<p>Processor Access Transfer Size</p> <p>These bits are used in combination with the lower two address bits of the Address register to determine the size of a processor access transaction. The bits are only valid when processor access is pending.</p> <table border="1"> <thead> <tr> <th>PAA[1:0]</th> <th>Psz[1:0]</th> <th>Transfer Size</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>00</td> <td>Byte (LE, byte 0; BE, byte 3)</td> </tr> <tr> <td>01</td> <td>00</td> <td>Byte (LE, byte 1; BE, byte 2)</td> </tr> <tr> <td>10</td> <td>00</td> <td>Byte (LE, byte 2; BE, byte 1)</td> </tr> <tr> <td>11</td> <td>00</td> <td>Byte (LE, byte 3; BE, byte 0)</td> </tr> <tr> <td>00</td> <td>01</td> <td>Halfword (LE, bytes 1:0; BE, bytes 3:2)</td> </tr> <tr> <td>10</td> <td>01</td> <td>Halfword (LE, bytes 3:2; BE, bytes 1:0)</td> </tr> <tr> <td>00</td> <td>10</td> <td>Word (LE, BE; bytes 3, 2, 1, 0)</td> </tr> <tr> <td>00</td> <td>11</td> <td>Triple (LE, bytes 2, 1, 0; BE, bytes 3, 2, 1)</td> </tr> <tr> <td>01</td> <td>11</td> <td>Triple (LE, bytes 3, 2, 1; BE, bytes 2, 1, 0)</td> </tr> <tr> <td colspan="2">All others</td> <td>Reserved</td> </tr> </tbody> </table> <p>Note: LE=little endian, BE=big endian, the byte# refers to the byte number in a 32-bit register, where byte 3 = bits 31:24; byte 2 = bits 23:16; byte 1 = bits 15:8; byte 0=bits 7:0, independently of the endianness.</p>	PAA[1:0]	Psz[1:0]	Transfer Size	00	00	Byte (LE, byte 0; BE, byte 3)	01	00	Byte (LE, byte 1; BE, byte 2)	10	00	Byte (LE, byte 2; BE, byte 1)	11	00	Byte (LE, byte 3; BE, byte 0)	00	01	Halfword (LE, bytes 1:0; BE, bytes 3:2)	10	01	Halfword (LE, bytes 3:2; BE, bytes 1:0)	00	10	Word (LE, BE; bytes 3, 2, 1, 0)	00	11	Triple (LE, bytes 2, 1, 0; BE, bytes 3, 2, 1)	01	11	Triple (LE, bytes 3, 2, 1; BE, bytes 2, 1, 0)	All others		Reserved	R	Undefined
PAA[1:0]	Psz[1:0]	Transfer Size																																			
00	00	Byte (LE, byte 0; BE, byte 3)																																			
01	00	Byte (LE, byte 1; BE, byte 2)																																			
10	00	Byte (LE, byte 2; BE, byte 1)																																			
11	00	Byte (LE, byte 3; BE, byte 0)																																			
00	01	Halfword (LE, bytes 1:0; BE, bytes 3:2)																																			
10	01	Halfword (LE, bytes 3:2; BE, bytes 1:0)																																			
00	10	Word (LE, BE; bytes 3, 2, 1, 0)																																			
00	11	Triple (LE, bytes 2, 1, 0; BE, bytes 3, 2, 1)																																			
01	11	Triple (LE, bytes 3, 2, 1; BE, bytes 2, 1, 0)																																			
All others		Reserved																																			
Res	28:24	Reserved.	R	0																																	
VPED	23	<p>VPE Disable state</p> <p>EJTAG state is not valid if this bit is set:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>VPE is currently enabled</td> </tr> <tr> <td>1</td> <td>VPE is currently disabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	VPE is currently enabled	1	VPE is currently disabled	R	1																											
Encoding	Meaning																																				
0	VPE is currently enabled																																				
1	VPE is currently disabled																																				
Doze	22	<p>Doze state</p> <p>The Doze bit indicates any type of low-power mode. The value is sampled in the Capture-DR state of the TAP controller:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>CPU not in low power mode</td> </tr> <tr> <td>1</td> <td>CPU is in low power mode</td> </tr> </tbody> </table> <p>Doze includes the Reduced Power (RP) and WAIT power-reduction modes.</p>	Encoding	Meaning	0	CPU not in low power mode	1	CPU is in low power mode	R	0																											
Encoding	Meaning																																				
0	CPU not in low power mode																																				
1	CPU is in low power mode																																				

**Table 16.29 EJTAG Control Register Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
Halt	21	<p>Halt state</p> <p>The Halt bit indicates if the internal system bus clock is running or stopped. The value is sampled in the Capture-DR state of the TAP controller:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Internal system clock is running</td> </tr> <tr> <td>1</td> <td>Internal system clock is stopped</td> </tr> </tbody> </table>	Encoding	Meaning	0	Internal system clock is running	1	Internal system clock is stopped	R	0
Encoding	Meaning									
0	Internal system clock is running									
1	Internal system clock is stopped									
PerRst	20	<p>Peripheral Reset</p> <p>When the bit is set to 1, it is only guaranteed that the peripheral reset has occurred in the system when the read value of this bit is also 1. This is to ensure that the setting from the <i>TCK</i> clock domain takes effect in the CPU clock domain and in peripherals.</p> <p>When the bit is written to 0, it must also be read as 0 before it is guaranteed that the indication is also cleared in the CPU clock domain.</p> <p>This bit controls the <i>EJ_PerRst</i> signal on the core.</p>	R/W	0						
PRnW	19	<p>Processor Access Read and Write</p> <p>This bit indicates if the pending processor access is for a read or write transaction, and the bit is only valid while <i>PrAcc</i> is set:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Read transaction</td> </tr> <tr> <td>1</td> <td>Write transaction</td> </tr> </tbody> </table>	Encoding	Meaning	0	Read transaction	1	Write transaction	R	Undefined
Encoding	Meaning									
0	Read transaction									
1	Write transaction									
PrAcc	18	<p>Processor Access (PA)</p> <p>Read value of this bit indicates if a Processor Access (PA) to the EJTAG memory is pending:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No pending processor access</td> </tr> <tr> <td>1</td> <td>Pending processor access</td> </tr> </tbody> </table> <p>The probe's software must clear this bit to 0 to indicate the end of the PA. A write of 1 is ignored.</p> <p>A pending Processor Access is cleared when <i>Rocc</i> is set, but another PA may occur just after the reset if a debug exception occurs.</p> <p>Finishing a Processor Access is not accepted while the <i>Rocc</i> bit is set. This is to avoid a Processor Access occurring after the reset is finished because of an indication of a Processor Access that occurred before the reset.</p> <p>The FASTDATA access can clear this bit.</p>	Encoding	Meaning	0	No pending processor access	1	Pending processor access	R/W0	0
Encoding	Meaning									
0	No pending processor access									
1	Pending processor access									
Res	17	Reserved	R	0						

**Table 16.29 EJTAG Control Register Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
PrRst	16	<p>Processor Reset (Implementation dependent behavior)</p> <p>When the bit is set to 1, then it is only guaranteed that this setting has taken effect in the system when the read value of this bit is also 1. This is to ensure that the setting from the <i>TCK</i> clock domain gets effect in the CPU clock domain, and in peripherals.</p> <p>When the bit is written to 0, then the bit must also be read as 0 before it is guaranteed that the indication is cleared in the CPU clock domain also.</p> <p>This bit controls the <i>EJ_PrRst</i> signal. If the signal is used in the system, then it must be ensured that both the processor and all devices required for a reset are properly reset. Otherwise the system may fail or hang. The bit resets itself, since the <i>EJTAG Control</i> register is reset by a reset.</p>	R/W	0						
ProbEn	15	<p>Probe Enable</p> <p>This bit indicates to the CPU if the EJTAG memory is handled by the probe so processor accesses are answered:</p> <table border="1" data-bbox="522 827 1024 1003"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Probe does not handle EJTAG memory transactions</td> </tr> <tr> <td>1</td> <td>Probe does handle EJTAG memory transactions</td> </tr> </tbody> </table> <p>It is an error by the software controlling the probe if it sets the ProbTrap bit to 1, but resets the <i>ProbEn</i> to 0. The operation of the processor is UNDEFINED in this case.</p> <p>The ProbEn bit is reflected as a read-only bit in the ProbEn bit, bit0, in the Debug Control Register (DCR).</p> <p>The read value indicates the effective value in the DCR, due to synchronization issues between <i>TCK</i> and CPU clock domains; however, it is ensured that change of the ProbEn prior to setting the <i>EjtagBrk</i> bit will have effect for the debug handler executed due to the debug exception.</p> <p>The reset value of the bit depends on whether the EJTAGBOOT indication is given or not:                      No EJTAGBOOT indication given: 0                      EJTAGBOOT indication given: 1</p>	Encoding	Meaning	0	Probe does not handle EJTAG memory transactions	1	Probe does handle EJTAG memory transactions	R/W	0 or 1 from EJTAGBOOT
Encoding	Meaning									
0	Probe does not handle EJTAG memory transactions									
1	Probe does handle EJTAG memory transactions									



**Table 16.29 EJTAG Control Register Descriptions (continued)**

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
ProbTrap	14	<p>Probe Trap This bit controls the location of the debug exception vector:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>In normal memory. Vector is located as described in <a href="#">Section 16.14.1.2 “DebugVectorAddr Register”</a></td> </tr> <tr> <td>1</td> <td>In EJTAG memory at 0xFF20.0200 in dmseg</td> </tr> </tbody> </table> <p>Valid setting of the ProbTrap bit depends on the setting of the ProbEn bit, see comment under ProbEn bit. The ProbTrap should not be set to 1 unless the ProbEn bit is also set to 1 to indicate that the EJTAG memory may be accessed. The read value indicates the effective value to the CPU, due to synchronization issues between <i>TCK</i> and CPU clock domains; however, it is ensured that change of the ProbTrap bit prior to setting the EjtagBrk bit will have effect for the EjtagBrk. The reset value of the bit depends on whether the EJTAGBOOT indication is given or not:</p>	Encoding	Meaning	0	In normal memory. Vector is located as described in <a href="#">Section 16.14.1.2 “DebugVectorAddr Register”</a>	1	In EJTAG memory at 0xFF20.0200 in dmseg	R/W	0 or 1 from EJTAGBOOT
Encoding	Meaning									
0	In normal memory. Vector is located as described in <a href="#">Section 16.14.1.2 “DebugVectorAddr Register”</a>									
1	In EJTAG memory at 0xFF20.0200 in dmseg									
Res	13	reserved	R	0						
EjtagBrk	12	<p>EJTAG Break Setting this bit to 1 causes a debug exception to the processor, unless the CPU was in debug mode or another debug exception occurred. When the debug exception occurs, the processor core clock is restarted if the CPU was in low power mode. This bit is cleared by hardware when the debug exception is taken. The reset value of the bit depends on whether the EJTAGBOOT indication is given or not:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No EJTAGBOOT indication given</td> </tr> <tr> <td>1</td> <td>EJTAGBOOT indication given</td> </tr> </tbody> </table>	Encoding	Meaning	0	No EJTAGBOOT indication given	1	EJTAGBOOT indication given	R/W	0 or 1 from EJTAGBOOT
Encoding	Meaning									
0	No EJTAGBOOT indication given									
1	EJTAGBOOT indication given									
Res	11:4	Reserved	R	0						
DM	3	<p>Debug Mode This bit indicates the debug or non-debug mode:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Processor is in non-debug mode</td> </tr> <tr> <td>1</td> <td>Processor is in debug mode</td> </tr> </tbody> </table> <p>The bit is sampled in the <i>Capture-DR</i> state of the TAP controller.</p>	Encoding	Meaning	0	Processor is in non-debug mode	1	Processor is in debug mode	R	0
Encoding	Meaning									
0	Processor is in non-debug mode									
1	Processor is in debug mode									
Res	2:0	Reserved	R	0						

## 16.14.5 Processor Access Registers

### 16.14.5.1 Processor Access Address Register

The Address register is used to provide the address of the processor access in the dmseg, and the register is only valid when a processor access is pending. The length of the Address register is 32 bits, and this register is selected by shifting in the ADDRESS instruction.

### 16.14.5.2 Processor Access Data Register


The Data register is used to provide data value to and from a processor access. The length of the Data register is 32 bits, and this register is selected by shifting in the DATA instruction.

The register has the written value for a processor access write due to a CPU store to the dmseg, and the output from this register is only valid when a processor access write is pending. The register is used to provide the data value for a processor access read due to a CPU load or fetch from the dmseg. The register will be updated with a new value when a processor access write is pending.

The Data register is 32 bits wide. Data alignment is not used for this register, so the value in the Data register matches data on the internal bus. The unused bytes for a processor access write are undefined, and for a Data register read, 0 (zero) must be shifted in for the unused bytes.

The organization of bytes in the Data register depends on the endianness of the core, as shown in [Figure 16.24](#). The endian mode for debug/kernel mode is determined by the state of the *SI\_Endian* input at power-up.

**Figure 16.24 Endian Formats for the Data Register**



The size of the transaction and thus the number of bytes available/required for the Data register is determined by the *Psz* field in the *ECR*.

## 16.14.6 Fastdata Registers

### 16.14.6.1 Fastdata Register (TAP Instruction FASTDATA)

The width of the Fastdata register is 1 bit. During a Fastdata access, the Fastdata register is written and read, i.e., a bit is shifted in and a bit is shifted out. During a Fastdata access, the Fastdata register value shifted in specifies whether

the Fastdata access should be completed or not. The value shifted out is a flag that indicates whether the Fastdata access was successful or not (if completion was requested).

**Figure 16.25 Fastdata Register Format**




**Table 16.30 Fastdata Register Field Description**

Fields		Description	Read / Write	Power-up State
Name	Bits			
SPrAcc	0	Shifting in a zero value requests completion of the Fastdata access. The PrAcc bit in the EJTAG Control register is overwritten with zero when the access succeeds. (The access succeeds if PrAcc is one and the operation address is in the legal dmseg Fastdata area.) When successful, a one is shifted out. Shifting out a zero indicates a Fastdata access failure.  Shifting in a one does not complete the Fastdata access and the PrAcc bit is unchanged. Shifting out a one indicates that the access would have been successful if allowed to complete and a zero indicates the access would not have successfully completed.	R/W	Undefined

### 16.14.7 FDC TAP Register

The FDC TAP instruction performs a 38 bit bidirectional transfer of the FDC TAP register. The register format is shown in [Figure 16.26](#) and the fields are described in [Figure 16.31](#)

**Figure 16.26 FDC TAP Register Format**



**Table 16.31 FDC TAP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
Probe Data Accept	37	Indicates to core that the probe is accepting the data that was scanned out.	W	Undefined
Data In Valid	36	Indicates to core that the probe is sending new data to the receive FIFO.	W	Undefined
Receive Buffer Full	37	Indicates to probe that the receive buffer is full and the core will not accept the data being scanned in. Analogous to ProbeDataAccept, but opposite polarity	R	0

**Table 16.31 FDC TAP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
Data Out Valid	36	Indicates to probe that the core is sending new data from the transmit FIFO	R	0
ChannelID	35:32	Channel number associated with the data being scanned in or out. This field can be used to indicate the type of data that is being sent and allow independent communication channels  Scanning in a value with ChannelID=0xd and Data In Valid = 0 will generate a receive interrupt. This can be used when the probe has completed sending data to the core.	R/W	Undefined
Data	31:0	Data value being scanned in or out	R/W	Undefined

### 16.14.8 Fast Debug Channel Registers

This section describes the Fast Debug Channel registers. CPU access to FDC is via loads and stores to the FDC device in the Common Device Memory Map (CDMM) region. These registers provide access control, configuration and status information, as well as access to the transmit and receive FIFOs. The registers and their respective offsets are shown in [Table 16.32](#)


**Table 16.32 FDC Register Mapping**

Offset in CDMM device block	Register Mnemonic	Register Name and Description
0x0	FDACSR	FDC Access Control and Status Register
0x8	FDCFG	FDC Configuration Register
0x10	FDSTAT	FDC Status Register
0x18	FDRX	FDC Receive Register
0x20 + 0x8* n	FDTXn	FDC Transmit Register n (0 ≤ n ≤ 15)

#### 16.14.8.1 FDC Access Control and Status (FDACSR) Register (Offset 0x0)

This is the general CDMM Access Control and Status register which defines the device type and size and controls user and supervisor access to the remaining FDC registers. The Access Control and Status register itself is only accessible in kernel mode. [Figure 16.27](#) has the format of an Access Control and Status register (shown as a 64-bit register), and [Table 16.33](#) describes the register fields.

**Figure 16.27 FDC Access Control and Status Register**




**Table 16.33 FDC Access Control and Status Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
DevType	31:24	This field specifies the type of device.	R	0xfd
DevSize	21:16	This field specifies the number of extra 64-byte blocks allocated to this device. The value 0x2 indicates that this device uses 2 extra, or 3 total blocks.	R	0x2
DevRev	15:12	This field specifies the revision number of the device. The value 0x0 indicates that this is the initial version of FDC	R	0x0
Uw	3	This bit indicates if user-mode write access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to write to the device while in user mode with access disabled is ignored.	R/W	0
Ur	2	This bit indicates if user-mode read access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to read from the device while in user mode with access disabled will return 0 and not change any state.	R/W	0
Sw	1	This bit indicates if supervisor-mode write access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to write to the device while in supervisor mode with access disabled is ignored.	R/W	0
Sr	0	This bit indicates if supervisor-mode read access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to read from the device while in supervisor mode with access disabled will return 0 and not change any state..	R/W	0
0	11:4	Reserved for future use. Ignored on write; returns zero on read.	R	0

**16.14.8.2 FDC Configuration (FDCFG) Register (Offset 0x8)**

The FDC configuration register holds information about the current configuration of the Fast Debug Channel mechanism. [Figure 16.28](#) has the format of the FDC Configuration register, and [Table 16.34](#) describes the register fields.

**Figure 16.28 FDC Configuration Register**




**Table 16.34 FDC Configuration Register Field Descriptions**

Fields		Description	Read / Write	Reset State										
Name	Bits													
0	31:20	Reserved for future use. Read as zeros, must be written as zeros.	R	0										
TxIntThresh	19:18	Controls whether transmit interrupts are enabled and the state of the TxFIFO needed to generate an interrupt. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Transmit Interrupt Disabled</td> </tr> <tr> <td>1</td> <td>Empty</td> </tr> <tr> <td>2</td> <td>Not Full</td> </tr> <tr> <td>3</td> <td>Almost Empty - zero or one entry in use*(see 16.15.2 for specifics)</td> </tr> </tbody> </table>	Encoding	Meaning	0	Transmit Interrupt Disabled	1	Empty	2	Not Full	3	Almost Empty - zero or one entry in use*(see 16.15.2 for specifics)	R/W	0
Encoding	Meaning													
0	Transmit Interrupt Disabled													
1	Empty													
2	Not Full													
3	Almost Empty - zero or one entry in use*(see 16.15.2 for specifics)													
RxIntThresh	17:16	Controls whether receive interrupts are enabled and the state of the RxFIFO needed to generate an interrupt. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Receive Interrupt Disabled</td> </tr> <tr> <td>1</td> <td>Full</td> </tr> <tr> <td>2</td> <td>Not empty</td> </tr> <tr> <td>3</td> <td>Almost Full - zero or one entry free</td> </tr> </tbody> </table>	Encoding	Meaning	0	Receive Interrupt Disabled	1	Full	2	Not empty	3	Almost Full - zero or one entry free	R/W	0
Encoding	Meaning													
0	Receive Interrupt Disabled													
1	Full													
2	Not empty													
3	Almost Full - zero or one entry free													
TxFIFOSize	15:8	This field holds the total number of entries in the transmit FIFO.	R	Preset										
RxFIFOSize	7:0	This field holds the total number of entries in the receive FIFO.	R	Preset										

**16.14.8.3 FDC Status (FDSTAT) Register (Offset 0x10)**

The FDC Status register holds up to date state information for the FDC mechanism. [Figure 16.29](#) has the format of the FDC Status register, and [Table 16.35](#) describes the register fields.

**Figure 16.29 FDC Status Register**



**Table 16.35 FDC Status Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
Tx_Count	31:24	This optional field is not implemented and will read as 0	R	0
Rx_Count	23:16	This optional field is not implemented and will read as 0	R	0
0	15:8	Reserved for future use. Must be written as zeros and read as zeros.	R	0


**Table 16.35 FDC Status Register Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
RxChan	7:4	This field indicates the channel number used by the top item in the receive FIFO. This field is only valid if RxE=0.	R	Undefined
RxE	3	If RxE is set, the receive FIFO is empty. If RxE is not set, the FIFO is not empty.	R	1
RxF	2	If RxF is set, the receive FIFO is full. If RxF is not set, the FIFO is not full.	R	0
TxE	1	If TxE is set, the transmit FIFO is empty. If TxE is not set, the FIFO is not empty.	R	1
TxF	0	If TxF is set, the transmit FIFO is full. If TxF is not set, the FIFO is not full.	R	0

**16.14.8.4 FDC Receive (FDRX) Register (Offset 0x18)**

This register exposes the top entry in the receive FIFO. A read from this register returns the top item in the FIFO and removes it from the FIFO itself. The result of a write to this register is **UNDEFINED**. The result of a read when the FIFO is empty is also **UNDEFINED** so software must check the  $FDSTAT_{RxE}$  flag prior to reading. Figure 16.30 has the format of the FDC Receive register, and Table 16.36 describes the register fields.

**Figure 16.30 FDC Receive Register**




**Table 16.36 FDC Receive Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
RxData	31:0	This register holds the top entry in the receive FIFO	R	Undefined

**16.14.8.5 FDC Transmit n (FDTXn) Registers (Offset 0x20 + 0x8\*n)**

These sixteen registers all access the bottom entry in the transmit FIFO. The different addresses are used to generate a 4b channel identifier that is attached to the data value. This allows software to track different event types without needing to reserve a portion of the 32b data as a tag. A write to one of these registers results in a write to the transmit FIFO of the data value and channel ID corresponding to the register being written. Reads from these registers are **UNDEFINED**. Attempting to write to the transmit FIFO if it is full has **UNDEFINED** results. Hence, the software running on the core must check the  $FDSTAT_{TxF}$  flag to ensure that there is space for the write. Figure 16.31 has the format of the FDC Transmit register, and Table 16.37 describes the register fields.

**Figure 16.31 FDC Transmit Register**



**Table 16.37 FDC Transmit Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
TxData	31:0	This register holds the bottom entry in the transmit FIFO	W, Undefined value on read	Undefined

**Table 16.38 FDTXn Address Decode**

Addr	Chan	Addr	Chan	Addr	Chan	Addr	Chan
0x20	0x0	0x40	0x4	0x60	0x8	0x80	0xc
0x28	0x1	0x48	0x5	0x68	0x9	0x88	0xd
0x30	0x2	0x50	0x6	0x70	0xa	0x90	0xe
0x38	0x3	0x58	0x7	0x78	0xb	0x98	0xf

### 16.14.9 PDtrace™ Registers (Software Control)

The CP0 registers associated with PDtrace are listed in [Table 16.39](#) and described in [Chapter 2, “CP0 Registers”](#) on [page 53](#).

**Table 16.39 A List of Coprocessor 0 Trace Registers**

Register Number	Sel	Register Name
23	1	<i>TraceControl</i>
23	2	<i>TraceControl2</i>
24	2	<i>TraceControl3</i>
23	3	<i>UserTraceData1</i>
24	3	<i>UserTraceData2</i>

### 16.14.10 Trace Control Block (TCB) Registers (Hardware Control)

The TCB registers used to control its operation are listed in [Table 16.40](#) and [Table 16.42](#). These registers are accessed via the EJTAG TAP interface, or by software through mapping to drseg memory space.

**Table 16.40 TCB EJTAG Registers**

EJTAG Register	Name	Description	Implemented
0x10	TCBCTRL0LA	Control register in the TCB mainly used for controlling the trace input signals to the core on the PDtrace interface. See <a href="#">Section 16.14.10.1 “TCBCTRL0LA Register”</a> .	Yes



**Table 16.40 TCB EJTAG Registers (continued)**

EJTAG Register	Name	Description	Implemented
0x11	TCBCONTROLB	Control register in the TCB that is mainly used to specify what to do with the trace information. The <i>REG</i> [25:21] field in this register specifies the number of the TCB internal register accessed by the <i>TCBDATA</i> register. A list of all the registers that can be accessed by the <i>TCBDATA</i> register is shown in . See <a href="#">Section 16.14.10.2 “TCBCONTROLB Register”</a> .	Yes
0x12	TCBDATA	This is used to access registers specified by the <i>REG</i> field in the <i>TCBCONTROLB</i> register. See <a href="#">Section 16.14.10.3 “TCBDATA Register”</a> .	Yes
0x13	TCBCONTROLC	Control Register in the TCB used to control and hold tracing information. See <a href="#">Section 16.14.10.4 “TCBCONTROLC Register”</a> .	Yes
0x16	TCBCONTROLE	Control Register in the TCB used to control tracing for the performance counter tracing feature. See <a href="#">Section 16.14.10.6 “TCBCONTROLE Register”</a> .	Yes

**Table 16.41 Registers Selected by TCBCONTROLB<sub>REG</sub>**

<i>TCBCONTROLB</i> <sub>REG</sub> field	Name	Reference	Implemented
0	TCBCONFIG	<a href="#">Section 16.14.10.7 “TCBCONFIG Register (Reg 0)”</a>	Yes
4 - 7		Values are undefined.	No
16-23	TCBTRIGx	<a href="#">Section 16.14.10.8 “TCBTRIGx Register (Reg 16-23)”</a>	Only the number indicated by <i>TCBCONFIG</i> <sub>TRIG</sub> are implemented.


### 16.14.10.1 TCBCONTROLA Register

The TCB is responsible for asserting or de-asserting the trace input control signals on the PDtrace interface to the core’s tracing logic. Most of the control is done using the *TCBCONTROLA* register.

The *TCBCONTROLA* register is written by an EJTAG TAP controller instruction, TCBCONTROLA (0x10). This register is also mapped to offset 0x3000 in drseg. .

The format of the *TCBCONTROLA* register is shown below, and the fields are described in [Table 16.42](#).

**Figure 16.32 TCBCONTROLA Register Format**



**Table 16.42 TCBCONTROLA Register Field Descriptions**

Fields		Description	Read / Write	Reset State																		
Name	Bits																					
SyPExt	31:30	<p>These two bits used to be Implementation specific until PDtrace spec revision 06.00 when it reverts to architecturally defined bits to extend the SyP (sync period) field for implementations that need higher numbers of cycles between synchronization events.</p> <p>The value of SyP is extended by assuming that these two bits are juxtaposed to the left of the three bits of SyP (SyPExt.SyP). When only SyP was used to specify the synchronization period, the value was <math>2^x</math>, where x was computed from SyP by adding 5 to the actual value represented by the bits. A similar formula is applied to the 5 bits just obtained by the juxtaposition of SyPExt and SyP. Sync period values greater than <math>2^{31}</math> are UNPREDICTABLE. Since the value of 11010 represents the value of 31 (with +5), all values greater than 11010 are UNPREDICTABLE.</p> <p>Note that with these new bits, a sync period range of <math>2^5</math> to <math>2^{31}</math> cycles can now be obtained.</p>	R/W	0																		
0	29:26	Reserved. Must be written as zero; returns zero on read.	R	0																		
VModes	25:24	<p>This field specifies the type of tracing that is supported by the processor, as follows:</p> <table border="1" data-bbox="462 924 1084 1117"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>PC tracing only</td> </tr> <tr> <td>01</td> <td>PC and Load and store address tracing only</td> </tr> <tr> <td>10</td> <td>PC, load and store address, and load and store data.</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	00	PC tracing only	01	PC and Load and store address tracing only	10	PC, load and store address, and load and store data.	11	Reserved	R	10								
Encoding	Meaning																					
00	PC tracing only																					
01	PC and Load and store address tracing only																					
10	PC, load and store address, and load and store data.																					
11	Reserved																					
ADW	23	<p><i>PDO_AD</i> bus width.</p> <p>0: The width is 16 bits.</p> <p>1: The width is 32 bits.</p>	R	1																		
SyP	22:20	<p>Used to indicate the synchronization period.</p> <p>The period (in cycles) between which the periodic synchronization information is to be sent is defined as shown in the table below.</p> <table border="1" data-bbox="618 1339 930 1732"> <thead> <tr> <th>SyP</th> <th>Sync Period</th> </tr> </thead> <tbody> <tr> <td>000</td> <td><math>2^5</math></td> </tr> <tr> <td>001</td> <td><math>2^6</math></td> </tr> <tr> <td>010</td> <td><math>2^7</math></td> </tr> <tr> <td>011</td> <td><math>2^8</math></td> </tr> <tr> <td>100</td> <td><math>2^9</math></td> </tr> <tr> <td>101</td> <td><math>2^{10}</math></td> </tr> <tr> <td>110</td> <td><math>2^{11}</math></td> </tr> <tr> <td>111</td> <td><math>2^{12}</math></td> </tr> </tbody> </table> <p>This field defines the value on the <i>PDI_SyncPeriod</i> signal.</p>	SyP	Sync Period	000	$2^5$	001	$2^6$	010	$2^7$	011	$2^8$	100	$2^9$	101	$2^{10}$	110	$2^{11}$	111	$2^{12}$	R/W	000
SyP	Sync Period																					
000	$2^5$																					
001	$2^6$																					
010	$2^7$																					
011	$2^8$																					
100	$2^9$																					
101	$2^{10}$																					
110	$2^{11}$																					
111	$2^{12}$																					

**Table 16.42 TCBCONTROLA Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
TB	19	Trace All Branches. When set to one, this field indicates that the core must trace either full or incremental PC values for all branches. When set to zero, only the unpredictable branches are traced.	R/W	Undefined
IO	18	Inhibit Overflow. This bit is used to indicate to the core trace logic that slow but complete tracing is desired. Hence, the core tracing logic must not allow a FIFO overflow and discard trace data. This is achieved by stalling the pipeline when the FIFO is nearly full so that no trace records are ever lost.	R/W	Undefined
D	17	When set to one, this enables tracing in Debug mode, i.e., when the <i>DM</i> bit is one in the <i>Debug</i> register. For trace to be enabled in Debug mode, the <i>On</i> bit must be one, and either the <i>G</i> bit must be one, or the current process must match the <i>ASID</i> field in this register. When set to zero, trace is disabled in Debug mode, irrespective of other bits.	R/W	Undefined
E	16	This controls when tracing is enabled. When set, tracing is enabled when either of the <i>EXL</i> or <i>ERL</i> bits in the <i>Status</i> register is one, provided that the <i>On</i> bit (bit 0) is also set, and either the <i>G</i> bit is set, or the current process ASID matches the <i>ASID</i> field in this register.	R/W	Undefined
S	15	When set, this enables tracing when the core is in Supervisor mode as defined in the MIPS32 or MIPS64 architecture specification. This is provided the <i>On</i> bit (bit 0) is also set, and either the <i>G</i> bit is set, or the current process ASID matches the <i>ASID</i> field in this register. If TraceControl3.GV is set, the GuestID of instruction execution must match the TraceControl3.GuestID register field for tracing to be enabled.	R/W	Undefined
K	14	When set, this enables tracing when the <i>On</i> bit is set and the core is in Kernel mode. Unlike the usual definition of Kernel Mode, this bit enables tracing only when the <i>ERL</i> and <i>EXL</i> bits in the <i>Status</i> register are zero. This is provided the <i>On</i> bit (bit 0) is also set, and either the <i>G</i> bit is set, or the current process ASID matches the <i>ASID</i> field in this register. If TraceControl3.GV is set, the GuestID of instruction execution must match the TraceControl3.GuestID register field for tracing to be enabled.	R/W	Undefined
U	13	When set, this enables tracing when the core is in User mode as defined in the MIPS32 or MIPS64 architecture specification. This is provided the <i>On</i> bit (bit 0) is also set, and either the <i>G</i> bit is set, or the current process ASID matches the <i>ASID</i> field in this register. If TraceControl3.GV is set, the GuestID of instruction execution must match the TraceControl3.GuestID register field for tracing to be enabled.	R/W	Undefined
ASID	12:5	The ASID field to match when the <i>G</i> bit is zero. When the <i>G</i> bit is one, this field is ignored.	R/W	Undefined
G	4	When set, this implies that tracing is to be enabled for all processes, provided that other enabling functions (like U, S, etc.,) are also true.	R/W	Undefined

**Table 16.42 TCBCONTROLA Register Field Descriptions (continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
TFCR	3	When set, this indicates to the PDtrace interface that complete information about instruction if it can be a function call or return should be traced. It also indicates to the TCB that the optional Fcr bit must be traced in the appropriate trace formats	R/W	Undefined
TLSM	2	When set, this indicates to the PDtrace interface that complete information about Load and Store data cache miss should be traced. It also indicates to the TCB that the optional LSm bit must be traced in the appropriate trace formats.	R/W	Undefined
TIM	1	When set, this indicates to the PDtrace interface that complete information about instruction cache miss should be traced. It also indicates to the TCB that the optional Im bit must be traced in the appropriate trace formats.	R/W	Undefined
On	0	This is the global trace enable switch to the core. When zero, tracing from the core is always disabled, unless enabled by core internal software override. When set to one, tracing is enabled whenever the other enabling functions are also true.	R/W	0

**16.14.10.2 TCBCONTROLB Register**

The TCB includes a second control register, *TCBCONTROLB* (0x11). This register generally controls what to do with the trace information received. This register is also mapped to offset 0x3008 in drseg.

The format of the *TCBCONTROLB* register is shown below, and the fields are described in [Table 16.43](#).

**Figure 16.33 TCBCONTROLB Register Format**



**Table 16.43 TCBCONTROLB Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
WE	31	Write Enable. Only when set to 1 will the other bits be written in <i>TCBCONTROLB</i> . This bit will always read 0.	R	0
0	30:28	Reserved. Must be written as zero; returns zero on read.	R	0
TWSrc-Width	27:26	Used to indicate the number of bits used in the source field of the Trace Word, this is a configuration option of the core that cannot be modified by software. 00 - zero source field width 01 - two bit source field width 10 - four bit source field width 11 - reserved for future use This field can only be 10 for the interAptiv core.	R	10

**Table 16.43 TCBCONTROLB Register Field Descriptions (continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
REG	25:21	Register select: This field select the registers accessible through the <i>TCBDATA</i> register. Legal values are shown in <a href="#">Table 16.42</a> .	R/W	0
WR	20	Write Registers: When set, the register selected by REG field is read and written when <i>TCBDATA</i> is accessed. Otherwise the selected register is only read.	R/W	0
0	19	Reserved. Must be written as zero; returns zero on read.	R	0
TRPAD	18	Trace RAM access disable bit, disables program software access to the on-chip trace RAM using load/store instructions. If probe access is not provided in the implementation, then this register bit must be tied to zero value to allow software to control access.	R/W	0
FDT	17	Filtered Data Trace Mode enable bit. When the bit is 0, this mode is disabled, reset value is disable. When set to 1, this mode is enabled.	R/W	0
0	16:12	Reserved. Must be written as zero; returns zero on read.	R	0
TLSIF	11	When set, this indicates to the TCB that information about Load and Store data cache miss, instruction cache miss, and function call are to be taken from the PDtrace interface and trace them out in the appropriate trace formats as the three optional bits LSm, Im, and Fcr.	R/W	0
0	10:7	Reserved. Must be written as zero; returns zero on read.	R	0
TWSrcVal	6:3	These bits are used to indicate the value of the TW source field that will be traced if TWSrcWidth indicates a source bit field width of 2 or 4 bits. Note that if the field is 2 bits, then only bits 4:3 of this field will be used in the TW.	R	Preset
CA	2	Cycle accurate trace. When set to 1, the trace will include stall information. When set to 0, the trace will exclude stall information, and remove bit zero from all transmitted TF's. The stall information included/excluded is: <ul style="list-style-type: none"> <li>• TF6 formats with TCBcode 0001 and 0101.</li> <li>• All TF1 formats.</li> </ul>	R/W	0
OfC	1	This bit is always set to 1, indicating that the trace is sent to off-chip memory using <i>TR_DATA</i> pins.	R	1
EN	0	Enable trace. This is the master enable for trace to be generated from the TCB. This bit can be set or cleared, either by writing this register or from a start/stop/about trigger. When set to 1, Trace Words are generated and sent to the trace funnel. When set to 0, trace information is ignored. A potential TF6-stop (from a stop trigger) is generated as the last information, the TCB pipe-line is flushed, and trace output is stopped.	R/W	0

### 16.14.10.3 TCBDATA Register

The *TCBDATA* register (0x12) is used to access the registers defined by the *TCBCONTROLB<sub>REG</sub>* field; see [Table 16.40](#). Regardless of which register or data entry is accessed through *TCBDATA*, the register is only written if the *TCBCONTROLB<sub>WR</sub>* bit is set. For read-only registers, *TCBCONTROLB<sub>WR</sub>* is a don't care.

The format of the *TCBDATA* register is shown below, and the field is described in [Table 16.44](#). The width of *TCBDATA* is 64 bits when on-chip trace words (TWs) are accessed (*TCBTW* access).

**Figure 16.34 TCBDATA Register Format**



**Table 16.44 TCBDATA Register Field Descriptions**

Fields		Description	Read/Write	Reset State
Names	Bits			
Data	31:0 63:0	Register fields or data as defined by the <i>TCBCONTROLB<sub>REG</sub></i> field	Only writable if <i>TCBCONTROLB<sub>WR</sub></i> is set	0


#### 16.14.10.4 TCBCONTROLC Register

The trace output from the processor on the PDtrace interface can be controlled by the trace input signals to the processor from the TCB. The TCB uses a control register, *TCBCONTROLC*, whose values are used to change the signal values on the PDtrace input interface. External software (i.e., debugger) can therefore manipulate the trace output by writing to this register.

The *TCBCONTROLC* register is written by the EJTAG TAP controller instruction, *TCBCONTROLC* (0x13). This register is also mapped to offset 0x3010 in *drseg*.

The format of the *TCBCONTROLC* register is shown below, and the fields are described in [Table 16.45](#).

**Figure 16.35 TCBCONTROLC Register Format**



**Table 16.45 TCBCONTROLC Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
Res	31:30	Reserved for future use. Must be written as zero; returns zero on read.	0	0
NumDO	29:28	Specifies the number of bits needed by this implementation to specify the DataOrder: 10 - Six bits	R	10

**Table 16.45 TCBCONTROL Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State												
Name	Bits															
Mode	27:23	<p>When tracing is turned on, this signal specifies what information is to be traced by the core. It uses 5 bits, where each bit turns on a tracing of a specific tracing mode.</p> <table border="1"> <thead> <tr> <th>Bit # Set</th> <th>Trace The Following</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>PC</td> </tr> <tr> <td>1</td> <td>Load address</td> </tr> <tr> <td>2</td> <td>Store address</td> </tr> <tr> <td>3</td> <td>Load data</td> </tr> <tr> <td>4</td> <td>Store data</td> </tr> </tbody> </table> <p>The table shows what trace value is turned on when that bit value is a 1. If the corresponding bit is 0, then the Trace Value shown in column two is not traced by the processor. On the interAptiv core, PC tracing is always enabled, regardless of the value on bit 23. This field defines the value on the <i>PDI_TraceMode</i> signal.</p>	Bit # Set	Trace The Following	0	PC	1	Load address	2	Store address	3	Load data	4	Store data	R/W	0
Bit # Set	Trace The Following															
0	PC															
1	Load address															
2	Store address															
3	Load data															
4	Store data															
CPUvalid	22	<p>This bit enables VPE based tracing. This bit is ignored if TCvalid field is set 0: Instructions for all VPEs are traced 1: Instructions from only one VPE specified in CPUId field are traced This field defines the value on the <i>PDI_CPUIdValid</i> signal</p>	R/W	0												
Res	21:15	Reserved for future use.	R/W	0												
CPUId	14	<p>This bit indicates the value of the VPEid to be traced if CPUValid field is set 0: Instructions from VPE0 are traced 1: Instructions from VPE1 are traced This field defines the value on the <i>PDI_CPUId</i> signal</p>	R/W	Undefined												
TCvalid	13	<p>This bit enables TC based tracing 0: Instructions are traced based on CPUValid/CPUId settings 1: Instructions from only one TC specified in TCnum field are traced This field defines the value on the <i>PDI_TCNumValid</i></p>	R/W	0												
Res	12:9	Reserved for future use.	R/W	Undefined												
TCnum	8:5	<p>This field indicates the value of the TC to be traced if TCvalid is set This field defines the value on the <i>PDI_TCNum</i> signal</p>	R/W	Undefined												
TCbits	4:2	<p>This value is used by the TCB to determine the number of bits needed to represent TC value in a Trace Format(TF). Returns 3 on reads indicating 4 bits are needed to represent 9 TC value.</p>	R	Preset												
MTtraceType	1	<p>This bit indicates the type of implemented multi-threading 0: Fine grained, i.e., switch threads every cycle. If MTtraceTC field is set then each Trace format is augmented by TC information 1: Coarse-grained, also known as block multi-threading. If MTtraceTC field is set then TF7 is used and each TF is not augmented. Returns 0 on read indicating that processor may switch threads every cycle if needed.</p>	0	Preset												

**Table 16.45 TCBCONTROLC Register Field Descriptions (continued)**


Fields		Description	Read / Write	Reset State
Name	Bits			
MTtraceTC	0	This bit controls TC value tracing. 0: TC value is not traced 1: TC value is tracing by augmenting TCId on each Trace format	R/W	Undefined

#### 16.14.10.5 TCBCONTROLD Register

The TCB includes a control register, TCBCONTROLD, whose values are used to enable tracing of the Coherence Manager. External software (i.e., debugger) can therefore manipulate the trace output by writing to this register. Each of the cores in the system has this register, and the *Core\_CM\_En* field is considered from each of the cores.

The *TCBCONTROLD* register is written by an EJTAG TAP controller instruction, *TCBCONTROLD* (0x14). This register is also mapped to offset 0x3018 in drseg. The format of the *TCBCONTROLD* register is shown below, and the fields are described in [Table 16.46](#).

**Figure 16.36 TCBCONTROLD Register Format**



**Table 16.46 TCBCONTROLD Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:2	Reserved. Must be written as zero; returns zero on read.	R	0
Core_CM_En	1	Core_CM_Enable: The CM looks at this bit coming from each of the cores. Allows cores other than the master to enable tracing if other conditions are met.	R/W	0
0	0	Reserved. Must be written as zero; returns zero on read.	R	0

#### 16.14.10.6 TCBCONTROLE Register

The trace output from the processor on the PDtrace interface can be controlled by the trace input signals to the processor from the TCB. The TCB uses a control register, *TCBCONTROLE*, whose values are used to change the signal values on the PDtrace input interface. External software (i.e., debugger), can therefore manipulate the trace output by writing the *TCBCONTROLE* register.

The *TCBCONTROLE* register is written by an EJTAG TAP controller instruction, *TCBCONTROLE* (0x16). This register is also mapped to offset 0x3020 in drseg.

The format of the *TCBCONTROLE* register is shown below, and the fields are described in [Table 16.47](#).



**Figure 16.37 TCBCONTROLE Register Format**

31	9	8	7	6	5	4	3	2	1	0
0	TdIDLE	0	PecOvf	PeCFCR	PeCBP	PeCSync	PeCE	PeC		

**Table 16.47 TCBCONTROLE Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:9	Reserved for future use. Must be written as zero; returns zero on read.	0	0
<i>TrIDLE</i>	8	Trace Unit Idle. This bit indicates if the trace hardware is currently idle (not processing any data). This can be useful when switching control of trace from hardware to software and vice versa. The bit is read-only and updated by the trace hardware.	R	1
0	7:6	Reserved for future use; Must be written as zero; returns zero on read. (Hint to architect, Reserved for future expansion of performance counter trace events).	0	0
<i>PeCOvf</i>	5	Trace performance counters when one of the performance counters overflows its count value. Enabled when set to 1.	R/W	0
<i>PeCFCR</i>	4	Trace performance counters on function call/return or on an exception handler entry. Enabled when set to 1.	R/W	0
<i>PeCBP</i>	3	Trace performance counters on hardware breakpoint match trigger. Enabled when set to 1.	R/W	0
<i>PeCSync</i>	2	Trace performance counters on synchronization counter expiration. Enabled when set to 1.	R/W	0
<i>PeCE</i>	1	Performance counter tracing enable. If performance counter hardware is present, this field is read/write. If not present, this field is read-only. When set to 0, the tracing out of performance counter values as specified is disabled. To enable, this bit must be set to 1. This bit is used under software control. When trace is controlled by an external probe, this enabling is done via the <i>TCB Control</i> register.	Config Option	0
<i>PeC</i>	0	Specifies whether or not Performance Control Tracing is implemented. This is an optional feature that may be omitted by implementation choice.	R	Preset

The following registers are accessed by the TCBCONTROLB<sub>REG</sub> field.

#### 16.14.10.7 TCBCONFIG Register (Reg 0)

The *TCBCONFIG* register holds information about the hardware configuration of the TCB. The format of the *TCBCONFIG* register is shown below, and the fields are described in [Table 16.49](#).

**Figure 16.38 TCBCONFIG Register Format**

31	30	25	24	21	20	17	16	14	13	11	10	9	8	6	5	4	3	0	
CF1	0	TRIG	SZ	CRMax	CRMin	PW	PiN	OnT	OfT	REV									

**Table 16.48 TCBCONFIG Register Field Descriptions**

Fields		Description	Read / Write	Reset State										
Name	Bits													
CF1	31	This bit is set if a <i>TCBCONFIG1</i> register exists. In this revision, <i>TCBCONFIG1</i> does not exist and this bit always reads zero.	R	0										
0	30:25	Reserved. Must be written as zero; returns zero on read.	R	0										
TRIG	24:21	Number of triggers implemented. This also indicates the number of <i>TCBTRIGx</i> registers that exist.	R	Preset Legal values are 0 - 8										
SZ	20:17	On-chip trace memory size. This field holds the encoded size of the on-chip trace memory. The size in bytes is given by $2^{(SZ+8)}$ , implying that the minimum size is 256 bytes and the largest is 8 Mb. This bit is reserved if on-chip memory is not implemented.	R	Preset										
CRMax	16:14	Off-chip Maximum Clock Ratio. This field indicates the maximum ratio of the CPU clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 16.49</a> . This bit is reserved if off-chip trace option is not implemented.	R	Preset										
CRMin	13:11	Off-chip Minimum Clock Ratio. This field indicates the minimum ratio of the CPU clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 16.49</a> . This bit is reserved if off-chip trace option is not implemented.	R	Preset										
PW	10:9	ProbeWidth: Number of bits available on the off-chip trace interface <i>TR_DATA</i> pins. The number of <i>TR_DATA</i> pins is encoded, as shown in the table.  <table border="1" data-bbox="500 1108 1000 1297"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>4-bits</td> </tr> <tr> <td>01</td> <td>8-bits</td> </tr> <tr> <td>10</td> <td>16-bits</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	00	4-bits	01	8-bits	10	16-bits	11	Reserved	R	Preset
Encoding	Meaning													
00	4-bits													
01	8-bits													
10	16-bits													
11	Reserved													
PiN	8:6	Pipe number. Indicates the number of execution pipelines.	R	0										
OnT	5	When set, this bit indicates that on-chip trace memory is present. This bit is preset based on the selected option when the TCB is implemented.	R	Preset										
OffT	4	When set, this bit indicates that off-chip trace interface is present. This bit is preset based on the selected option when the TCB is implemented, and on the existence of a PIB module ( <i>TC_PibPresent</i> asserted).	R	Preset										
REV	3:0	Revision of TCB.	R	0x3										

**Table 16.49 Clock Ratio Encoding of the CR Field**

CR/CRMin/CRMax	Clock Ratio
000	8:1 (Trace clock is eight times that of CPU clock)
001	4:1 (Trace clock is four times that of CPU clock)
010	2:1 (Trace clock is two times that of CPU clock)
011	1:1 (Trace clock is same as CPU clock)
100	1:2 (Trace clock is one half of CPU clock)
101	1:4 (Trace clock is one fourth of CPU clock)
110	1:6 (Trace clock is one sixth of CPU clock)
111	1:8 (Trace clock is one eighth of CPU clock)


### 16.14.10.8 TCBTRIGx Register (Reg 16-23)

Up to eight Trigger Control registers are possible. Each register is named *TCBTRIGx*, where *x* is a single digit number from 0 to 7 (*TCBTRIG0* is Reg 16). The actual number of trigger registers implemented is defined in the *TCBCONFIG<sub>TRIG</sub>* field. An unimplemented register will read all zeros and writes are ignored.

Each Trigger Control register controls when an associated trigger is fired, and the action to be taken when the trigger occurs. Please also read [Section 16.16 “TCB Trigger Logic”](#), for detailed description of trigger logic issues.

The format of the *TCBTRIGx* register is shown below, and the fields are described in [Table 16.50](#).

**Figure 16.39 TCBTRIGx Register Format**



**Table 16.50 TCBTRIGx Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
TCBInfo	31:24	This field is to be used in a possible TF6 trace format when this trigger fires.	R/W	0
Trace	23	When set, generate TF6 trace information when this trigger fires. Use <i>TCBInfo</i> field for the TCBInfo of TF6 and use <i>Type</i> field for the two MSB of the TCBtype of TF6. The two LSB of <i>TCBtype</i> are 00. The write value of this bit always controls the behavior of this trigger. When this trigger fires, the read value will change to indicate if the TF6 format was ever suppressed by a simultaneous trigger. If so, the read value will be 0. If the write value was 0, the read value is always 0. This special read value is valid until the <i>TCBTRIGx</i> register is written.	R/W	0
0	22:16	Reserved. Must be written as zero; returns zero on read.	R	0

**Table 16.50 TCBTRIGx Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State
Names	Bits			
CHTro	15	When set, generate a single cycle strobe on <i>TC_ChipTrigOut</i> when this trigger fires.	R/W	0
PDTro	14	When set, generate a single cycle strobe on <i>TC_ProbeTrigOut</i> when this trigger fires.	R/W	0
0	13:7	Reserved. Must be written as zero; returns zero on read.	R	0
DM	6	When set, this Trigger will fire when a rising edge on the Debug mode indication from the core is detected. The write value of this bit always controls the behavior of this trigger. When this trigger fires, the read value will change to indicate if this source was ever the cause of a trigger action (even if the action was suppressed). If so the read value will be 1. If the write value was 0 the read value is always 0. This special read value is valid until the <i>TCBTRIGx</i> register is written.	R/W	0
CHTri	5	When set, this Trigger will fire when a rising edge on <i>TC_ChipTrigIn</i> is detected. The write value of this bit always controls the behavior of this trigger. When this trigger fires, the read value will change to indicate if this source was ever the cause of a trigger action (even if the action was suppressed). If so the read value will be 1. If the write value was 0 the read value is always 0. This special read value is valid until the <i>TCBTRIGx</i> register is written.	R/W	0
PDTri	4	When set, this Trigger will fire when a rising edge on <i>TC_ProbeTrigIn</i> is detected. The write value of this bit always controls the behavior of this trigger. When this trigger fires, the read value will change to indicate if this source was ever the cause of a trigger action (even if the action was suppressed). If so the read value will be 1. If the write value was 0 the read value is always 0. This special read value is valid until the <i>TCBTRIGx</i> register is written.	R/W	0

**Table 16.50 TCBTRIGx Register Field Descriptions (continued)**

Fields		Description	Read / Write	Reset State										
Names	Bits													
Type	3:2	<p>Trigger Type: The Type indicates the action to take when this trigger fires. The table below show the Type values and the Trigger action.</p> <table border="1"> <thead> <tr> <th>Type</th> <th>Trigger action</th> </tr> </thead> <tbody> <tr> <td>00</td> <td><b>Trigger Start:</b> Trigger start-point of trace.</td> </tr> <tr> <td>01</td> <td><b>Trigger End:</b> Trigger end-point of trace.</td> </tr> <tr> <td>10</td> <td><b>Trigger About:</b> Trigger center-point of trace.</td> </tr> <tr> <td>11</td> <td><b>Trigger Info:</b> No action trigger, only for trace info.</td> </tr> </tbody> </table> <p>The actual action is to set or clear the <math>TCBCONTROLB_{EN}</math> bit. A Start trigger will set <math>TCBCONTROLB_{EN}</math>, a End trigger will clear <math>TCBCONTROLB_{EN}</math>. The About trigger will clear <math>TCBCONTROLB_{EN}</math> half way through the trace memory, from the trigger. The size determined by the <math>TCBCONFIG_{SZ}</math> field for on-chip memory. Or from the <math>TCBCONTROLA_{SyP}</math> field for off-chip trace.</p> <p>If Trace is set, then a TF6 format is added to the trace words. For Start and Info triggers this is done before any other TF's in that same cycle. For End and About triggers, the TF6 format is added after any other TF's in that same cycle.</p> <p>If the <math>TCBCONTROLB_{TM}</math> field is implemented it must be set to Trace-To mode (00), for the <i>Type</i> field to control on-chip trace fill. The write value of this bit always controls the behavior of this trigger.</p> <p>When this trigger fires, the read value will change to indicate if the trigger action was ever suppressed. If so the read value will be 11. If the write value was 11 the read value is always 11. This special read value is valid until the <math>TCBTRIGx</math> register is written.</p>	Type	Trigger action	00	<b>Trigger Start:</b> Trigger start-point of trace.	01	<b>Trigger End:</b> Trigger end-point of trace.	10	<b>Trigger About:</b> Trigger center-point of trace.	11	<b>Trigger Info:</b> No action trigger, only for trace info.	R/W	0
Type	Trigger action													
00	<b>Trigger Start:</b> Trigger start-point of trace.													
01	<b>Trigger End:</b> Trigger end-point of trace.													
10	<b>Trigger About:</b> Trigger center-point of trace.													
11	<b>Trigger Info:</b> No action trigger, only for trace info.													
FO	1	<p>Fire Once. When set, this trigger will not re-fire until the <math>TR</math> bit is de-asserted. When de-asserted this trigger will fire each time one of the trigger sources indicates trigger.</p>	R/W	0										
TR	0	<p>Trigger happened. When set, this trigger fired since the <math>TR</math> bit was last written 0.</p> <p>This bit is used to inspect whether the trigger fired since this bit was last written zero.</p> <p>When set, all the trigger source bits (bit 4 to 13) will change their read value to indicate if the particular bit was the source to fire this trigger. Only enabled trigger sources can set the read value, but more than one is possible.</p> <p>Also when set the <i>Type</i> field and the <i>Trace</i> field will have read values which indicate if the trigger action was ever suppressed by a higher priority trigger.</p>	R/W0	0										

### 16.14.11 Register Reset State

Reset state for all register fields is entered when either of the following occur:

1. TAP controller enters/is in Test-Logic-Reset state.
2. *EJ\_TRST\_N* input is asserted low.

## 16.15 Fast Debug Channel

The Fast Debug Channel (FDC) mechanism provides an efficient means to transfer data between the core and an external device using the EJTAG TAP pins. The FDC was created to allow for faster communication between the core and the probe. In previous generation MIPS processors, whenever the core wanted to communicate with the probe, the core would be halted and data sent to the probe because the probe had no way to read the core. The FDC provides a mechanism using FIFO's, whereby the probe can read the core without requiring that the core be halted. These FIFO's provide a cross boundary between the core and the EJTAG regions of the interAptiv core.

In the FDC, when the probe wishes to read an FDC register, the core gets an interrupt from the probe requesting this information. The core then places the requested information into the FIFO and continues operation. The core places information in the top of the FIFO, and the probe reads information from the bottom of the FIFO. The data contains information such as transmit versus receive, status of the operation, etc.

The external device would typically be an EJTAG probe and that is the term used here, but it could be something else. FDC utilizes two First In First Out (FIFO) structures to buffer data between the core and probe. The probe uses the FDC TAP instruction to access these FIFOs, while the core itself accesses them using memory accesses. To transfer data out of the core, the core writes one or more pieces of data to the transmit FIFO. At this time, the core can resume doing other work. An external probe would examine the status of the transmit FIFO periodically. If there is data to be read, the probe starts to receive data from the FIFO, one entry at a time. When all data from the FIFO has been drained, the probe goes back to waiting for more data. The core can either choose to be informed of the empty transmit FIFO via an interrupt, or it can choose to periodically check the status. Receiving data works in a similar manner - the probe writes to the receive FIFO. At that time, the core is either interrupted, or finds out via polling a status bit. The core can then do load accesses to the receive FIFO and receive data being sent to it by the probe. The TAP transfer is bidirectional - a single shift can be pulling transmit data and putting receive data at the same time.

The primary advantage of FDC over normal processor accesses or fastdata accesses is that it does not require the core to be blocked when the probe is reading or writing to the data transfer FIFOs. This significantly reduces the core overhead and makes the data transfer far less intrusive to the code executing on the core.

The FDC memory mapped registers are located in the common device memory map (CDMM) region. FDC has a device ID of 0xFD.

### 16.15.1 Common Device Memory Map

Software on the core accesses FDC through memory mapped registers. These memory mapped registers are located within the Common Device Memory Map (CDMM). The CDMM is a region of physical address space that is reserved for mapping IO device configuration registers within a MIPS processor. The base address and enabling of this region is controlled by the CDMMBase CP0 register.

### 16.15.2 Fast Debug Channel Interrupt

The FDC block can generate an interrupt to inform software of incoming data being available or space being available in the outgoing FIFO. This interrupt is handled similarly to the timer or performance counter interrupts. The *Cause<sub>FDCI</sub>* bit indicates that the interrupt is pending. The interrupt is also sent to the core outputs *SI\_FDCI[1]* where it is combined with one of the *SI\_Int* pins. For non-EIC mode, the *SI\_IPFDCI* input indicates which interrupt pin is has

been combined with and this information is reflected in the *IntCtl<sub>IPFDCI</sub>* field. Note that this interrupt is a regular interrupt and not a debug interrupt.

The FDC Configuration Register (see [Section 16.14.8.2 “FDC Configuration \(FDCFG\) Register \(Offset 0x8\)”](#)) includes fields for enabling and setting the threshold for generating each interrupt. Receive and transmit interrupt thresholds are specified independently, but transmit/receive interrupts are ORed together to form a single interrupt per VPE.

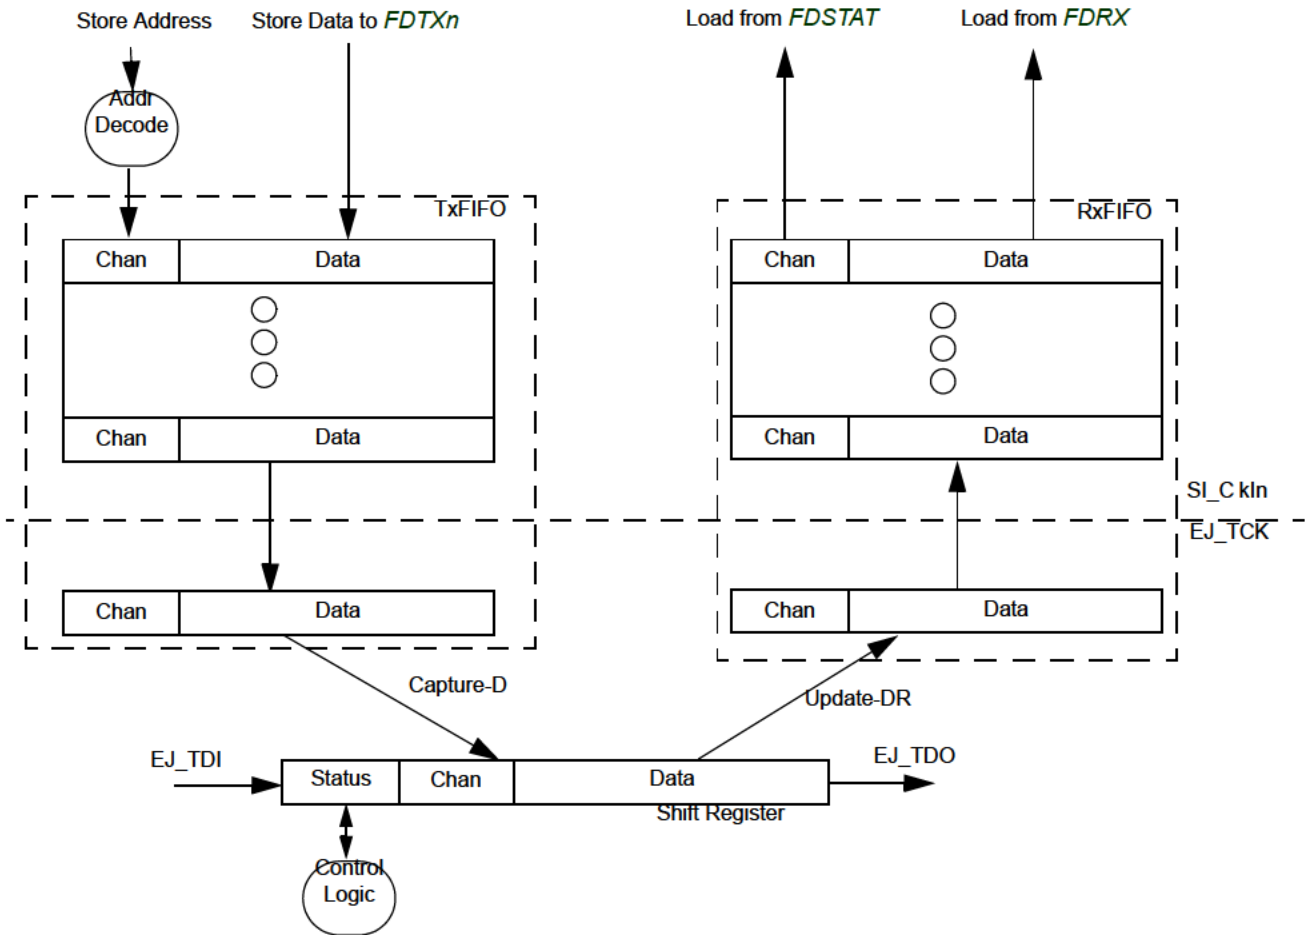
The following interrupt thresholds are supported:

- **Interrupts Disabled:** No interrupt will be generated and software must poll the status registers to determine if incoming data is available or if there is space for outgoing data.
- **Minimum core Overhead:** This setting minimizes the core overhead by not generating an interrupt until the receive FIFO (RxFIFO) is completely full or the transmit FIFO (TxFIFO) is completely empty.
- **Minimum latency:** To have the core take data as soon as it is available, the receive interrupt can be fired whenever the RxFIFO is not empty. There is a complimentary TxFIFO not full setting although that may not be quite as useful.
- **Maximum bandwidth:** When configured for minimum core overhead, bandwidth between the probe and core can be wasted if the core does not service the interrupt before the next transfer occurs. To reduce the chances of this happening, the interrupt threshold can be set to almost full or almost empty to generate an interrupt earlier. This setting causes receive interrupts to be generated when there are 0 or 1 unused RxFIFO entries. Transmit interrupts are generated when there are 0 or 1 used TxFIFO entries (see note in following section about this condition)

### 16.15.3 Core FDC Buffers

[Figure 16.40](#) shows the general organization of the transmit and receive buffers on the interAptiv core.

Figure 16.40 Fast Debug Channel Buffer Organization



One particular thing to note is the asynchronous crossings between the  $EJ\_TCK$  and  $SI\_ClkIn$  clock domains. This crossing is handled with a handshaked interface that safely transfers data between the domains. Two data registers are included in this interface, one in the source domain and one in the destination domain. The control logic actively manages these registers so that they can be used as FIFO entries. The fact that one FIFO entry is in the  $EJ\_TCK$  clock domain is normally transparent, but it can create some unexpected behavior:

- **TxFIFO availability:** Data is first written into the  $SI\_ClkIn$  FIFO entries, then it will move into the  $EJ\_TCK$  FIFO entry. But, it takes several  $EJ\_TCK$  cycles to complete the handshake and move the data.  $EJ\_TCK$  is generally much slower than  $SI\_ClkIn$  and may even be stopped (although that would be uncommon when this feature is in use). This can result in there not being space for new data, even though there are only  $N-1$  data values queued up. To prevent the loss of data, the  $FDSTAT_{Tx}$  bit is set when all of the  $SI\_ClkIn$  FIFO entries are full. Software writing to the FIFO should always check the  $FDSTAT_{Tx}$  bit prior to attempting a write and should not make any assumptions about being able to arbitrarily use all entries. ie. software seeing the  $FDSTAT_{Fx}$  bit set should not assume that it can write  $FDCFG_{TxCnt}$  data words without checking for full.
- **TxFIFO Almost Empty Interrupt:** As transmit data moves from  $SI\_ClkIn$  to  $EJ\_TCK$ , both of the flops will temporarily look full. This makes it difficult to determine when just 1 FIFO entry is in use. To enable a simpler condition, the almost empty TxInterrupt condition is set when all of the  $SI\_ClkIn$  FIFO entries are empty. When this



condition is met, there will be 0 or 1 valid entries. However, the interrupt will not be asserted when there is only one valid entry if it is an *SL\_ClkIn* entry

- The RxFIFO has similar characteristics but these are even less visible to software since *SL\_ClkIn* must be running to access the FDC registers.

#### 16.15.4 Sleep mode

FDC data transfers do not prevent the core from entering sleep mode and will proceed normally in sleep mode. The FDC block monitors the TAP interface signals with a free-running clock. When new receive data is available or transmit data can be sent, the gated clock will be enabled for a few cycles to transfer the data and then allowed to stop again. If FDC interrupts are enabled, transferring data may cause an interrupt to be generated which can wake the core up.

### 16.16 TCB Trigger Logic

The TCB is optionally implemented with trigger unit. If this is the case, then the *TCBCONFIG<sub>TRIG</sub>* field is non-zero. This section will explain some of the issues around triggers in the TCB.

#### 16.16.1 TCB Trace Enabling

The TCB must be enabled in order to produce a trace to the trace funnel, when trace information is sent on the PDtrace interface. The main switch for this is the *TCBCONTROLB<sub>EN</sub>* bit. When set, the TCB will send trace information to the trace funnel.

The TCB can optionally include trigger logic, which can control the *TCBCONTROLB<sub>EN</sub>* bit. Please see [Section 16.16 “TCB Trigger Logic”](#) for details.

#### 16.16.2 Tracing a Reset Exception

Tracing a reset exception is possible. However, the *TraceControl<sub>TS</sub>* bit is reset to 0 at core reset, so all the trace control must be from the TCB (using *TCBCONTROLA* and *TCBCONTROLB*). The PDtrace fifo and the entire TCB are reset based on an EJTAG reset. It is thus possible to set up the trace modes, etc., using the TAP controller, and then reset the core.

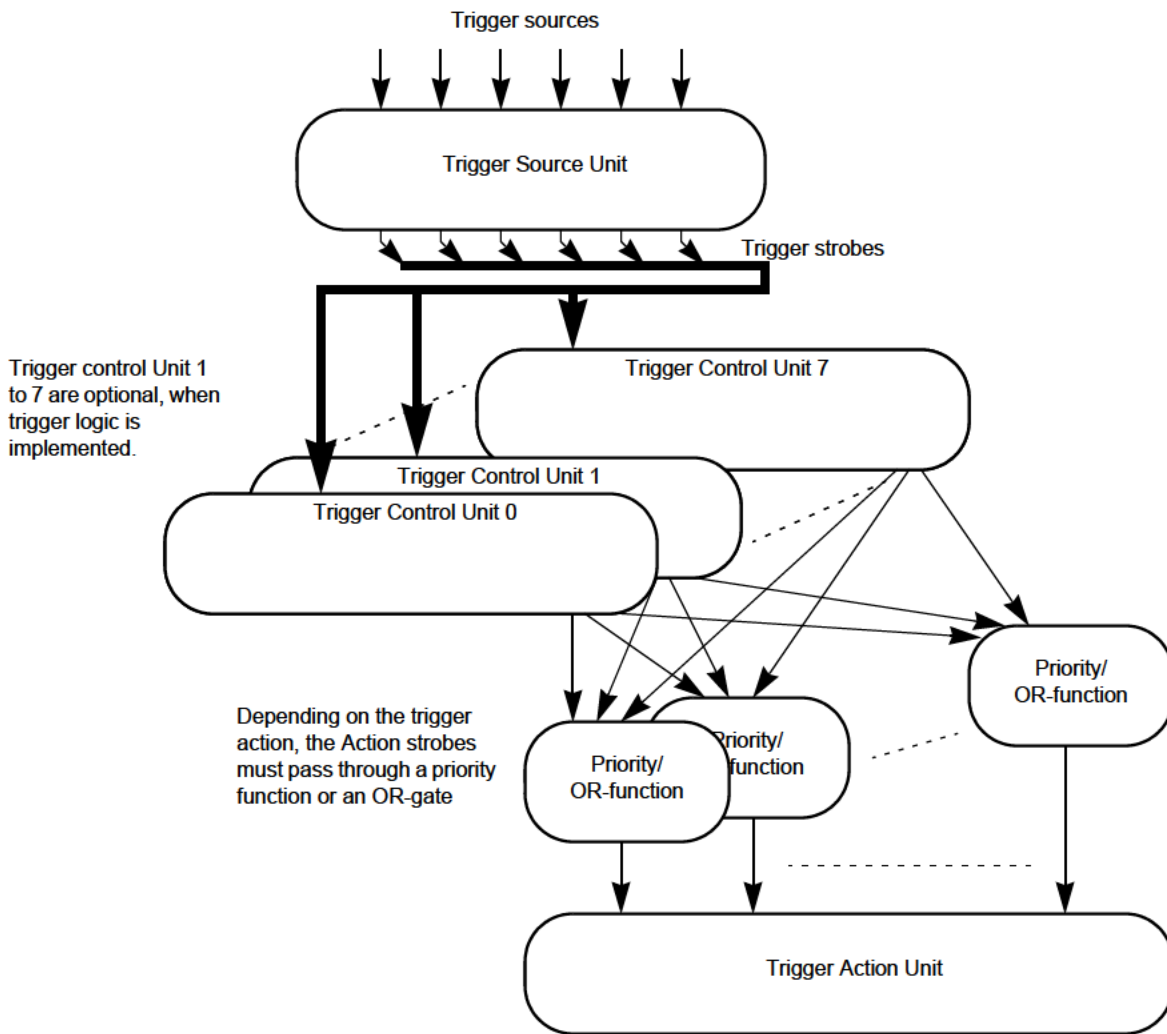
#### 16.16.3 Trigger Units Overview

TCB trigger logic features three main parts:

1. A common Trigger Source detection unit.
2. 1 to 8 separate Trigger Control units.
3. A common Trigger Action unit.

[Figure 16.41](#) show the functional overview of the trigger flow in the TCB.

**Figure 16.41 TCB Trigger Processing Overview**



### 16.16.4 Trigger Source Unit

The TCB has three trigger sources:

1. Chip-level trigger input (*TC\_ChipTrigIn*).
2. Probe trigger input (*TR\_TRIGIN*).
3. Debug Mode (DM) entry indication from the core.

The input triggers are all rising-edge triggers, and the Trigger Source Units convert the edge into a single cycle strobe to the Trigger Control Units.

## 16.16.5 Trigger Control Units

Up to eight Trigger Control Units are possible. Each of them has its own Trigger Control Register ( $TCBTRIGx$ ,  $x=\{0..7\}$ ). Each of these registers controls the trigger fire mechanism for the unit. Each unit has all of the Trigger Sources as possible trigger event and they can fire one or more of the Trigger Actions. This is all defined in the Trigger Control register  $TCBTRIGx$  (see [Section 16.14.10.8 “TCBTRIGx Register \(Reg 16-23\)”](#)).

## 16.16.6 Trigger Action Unit

The TCB has four possible trigger actions:

1. Chip-level trigger output ( $TC\_ChipTrigOut$ ).
2. Probe trigger output ( $TR\_TRIGOUT$ ).
3. Trace information. Put a programmable byte into the trace stream from the TCB.
4. Start, End or About (delayed end) control of the  $TCBCONTROLB_{EN}$  bit.

The basic function of the trigger actions is explained in [Section 16.14.10.8 “TCBTRIGx Register \(Reg 16-23\)”](#). Please also read the next [Section 16.16.7 “Simultaneous Triggers”](#).

## 16.16.7 Simultaneous Triggers

Two or more triggers can fire simultaneously. The resulting behavior depends on trigger action set for each of them, and whether they should produce a TF6 trace information output or not. There are two groups of trigger actions: Prioritized and OR'ed.

### 16.16.7.1 Prioritized Trigger Actions

For prioritized simultaneous trigger actions, the trigger control unit which has the lowest number takes precedence over the higher numbered units. The  $x$  in  $TCBTRIGx$  registers defines the number. The oldest trigger takes precedence over everything.

The following trigger actions are prioritized when two or more units fire simultaneously:

- Trigger Start, End and About type triggers ( $TCBTRIGx_{Type}$  field set to 00, 01 or 10), which will assert/de-assert the  $TCBCONTROLB_{EN}$  bit. The About trigger is delayed and will always change  $TCBCONTROLB_{EN}$  because it is the oldest trigger when it de-asserts  $TCBCONTROLB_{EN}$ . An About trigger will not start the countdown if an even older About trigger is using the Trace Word counter.
- Triggers which produce TF6 trace information in the trace flow (Trace bit is set).

Regardless of priority, the  $TCBTRIGx_{TR}$  bit is set when the trigger fires. This is so even if a trigger action is suppressed by a higher priority trigger action. If the trigger is set to only fire once (the  $TCBTRIGx_{FO}$  bit is set), then the suppressed trigger action will not happen until after  $TCBTRIGx_{TR}$  is written 0.

If a Trigger action is suppressed by a higher priority trigger, then the read value, when the  $TCBTRIGx_{TR}$  bit is set, for the  $TCBTRIGx_{Trace}$  field will be 0 for suppressed TF6 trace information actions. The read value in the  $TCBTRIGx_{Type}$  field for suppressed Start/End/About triggers will be 11. This indication of a suppressed action is sticky. If any of the two actions (Trace and Type) are ever suppressed for a multi-fire trigger (the  $TCBTRIGx_{FO}$  bit is zero), then the read values in Trace and/or Type are set to indicate any suppressed action.

### **About Trigger**

The About triggers delayed de-assertion of the *TCBCONTROLB<sub>EN</sub>* bit is always executed, regardless of priority from another Start trigger at the time of the *TCBCONTROLB<sub>EN</sub>* change. This means that if a simultaneous About trigger action on the *TCBCONTROLB<sub>EN</sub>* bit (*n/2* Trace Words after the trigger) and a Start trigger hit the same cycle, then the About trigger wins, regardless of which trigger number it is. The oldest trigger takes precedence.

However, if an About trigger has started the count down from *n/2*, but not yet reached zero, then a new About trigger, will NOT be executed. Only one About trigger can have the cycle counter. This second About trigger will store 11 in the *TCBTRIGx<sub>Type</sub>* field. But, if the *TCBTRIGx<sub>Trace</sub>* bit is set, a TF6 trace information will still go in the trace.

#### **16.16.7.2 OR'ed Trigger Actions**

The simple trigger actions CHTro and PDTro from each trigger unit, are effectively OR'ed together to produce the final trigger. One or more expected trigger strobes on i.e. *TC\_ChipTrigOut* can thus disappear. External logic should not rely on counting of strobes, to predict a specific event, unless simultaneous triggers are known not to occur.

## Multi-CPU Debug

This section describes the debug features of the interAptiv Multiprocessing System. The following sections are included in this chapter:

- [Section 17.1 “CM Performance Counters”](#)
- [Section 17.2 “Debug Mode Triggering”](#)
- [Section 17.3 “PDTrace Software Architecture”](#)

### 17.1 CM Performance Counters

#### 17.1.1 CM Performance Counter Functionality

Performance characteristics of the CM can be measured via the CM performance counters. Two sets of identical programmable 32-bit performance counters in addition to a 32-bit cycle counter are implemented. The counters are controlled and accessed via GCR registers described in [Chapter 8, “Coherency Manager” on page 379](#). This section describes the operation of those registers.

The counters are started by writing a 1 to the *P0\_CountOn*, *P1\_CountOn* and *Cycl\_Cnt\_CountOn* bits in the *CM Performance Counter Control Register* (see [Table 8.49](#) for a description of this register). Each counter can be reset to 0, and the corresponding overflow bit (*P0\_Overflow*, *P1\_Overflow*, *Cycl\_Cnt\_Overflow*) is reset to 0 prior to the start of counting by writing a 1 to the *P0\_Reset*, *P1\_Reset* and *Cycl\_Cnt\_Reset* bits in the same access that sets the corresponding start bits. This functionality allows all three counters to be reset and started with a single GCR write.

The *CM Performance Counter Control Register* also controls how a counter overflow is handled. If the *Perf\_Ovf\_Stop* bit is set to 1, then all CM Performance counters will stop when one of the counters (including the Cycle Counter) reaches its maximum value of 0xFFFFFFFF. If instead the *Perf\_Ovf\_Stop* bit is set to 0, when a counter overflows, it rolls over and continues counting from 0.

If the *Perf\_Int\_En* bit is set to 1, an interrupt is generated when one of the counters (including the cycle counter) reaches its maximum value of 0xFFFFFFFF. The CM asserts the *CM\_PCInt* signal which generates an interrupt only if the System Integrator has connected *CM\_PCInt* to one bit of *SI\_CMIInt*.

When a performance counter overflows, the corresponding bit is automatically set in the *CM Performance Counter Overflow Status Register*. A status bit is cleared by writing a 1 to it.

The event to be counted by each performance counter is designated by the event number set in the *Event\_Sel\_0* and *Event\_Sel\_1* fields of the *CM Performance Counter Event Selection Register*. The events corresponding to the event numbers are listed and described in [Table 17.1](#). Each event is further specified by the *CM Performance Counter Qualifier Register*. The meaning of the *CM Performance Counter Qualifier Register* is different for each event. The column labeled “Qualifier” in [Table 17.1](#) shows the qualifiers that can be specified for each event. For example, the qualifiers for the *Request\_Count* event (Event 0) are the request port, CCA, Burst Length, Command, and Target. The details of the qualifiers for the *Request\_Count* event are defined in [Table 17.2](#).

The qualifiers for some events are composed of several groups. A performance counter will increment if the specified event occurs and the qualifier criteria is matched in all groups. For example, assume the *Event\_Sel\_0* field in the *CM Performance Counter Event Selection Register* is set to 0 (Request\_Count). This event occurs when the CM serializes a request. However, the performance counter for this event will only count if the request meets the criteria programmed in all 5 groups in the Request Qualifier (see [Table 17.2](#)):

```

    The port that issued the request has the corresponding Request Port qualifier bit
    set to 1
AND
    The Cacheability attribute (CCA) for the request has the corresponding CCA
    qualifier bit set to 1
AND
    The Burst Length of the request (in dwords) has the corresponding qualifier bit set
    to 1
AND
    The OCP MCmd Type for the request has the corresponding Request Command qualifier
    bit set to 1
AND
    The target of the request has the corresponding Target qualifier bit set to 1

```

Multiple bits within a qualification group may be set. In this case, the OR of all bits set within the group. For example, by setting the request port qualifier for Port 0 and Port 1, then a request will be counted if it originated from Port 0 or Port 1.

A qualifier group can be set to “don’t care” by setting all bits within the group to 1. For example, to have performance counter 0 count all requests from port 1, program the *CM Performance Counter Event Selection Register* and *CM Performance Counter Qualifier 0 Register* as follows:

```

Set Event_Sel_0 to 0 (Request_Count)
Set Request Port Qualifier bit to 1 for Port 1
Set Request Port Qualifier bits to 0 for all other Ports
Set all other qualifier bits to 1 (causing the CCA, Burst Length, Command and Target
to be ignored)

```

The two counters can be programmed to count a different event or the same event with different qualifiers. For example, to measure the ratio of requests from Port 1 vs. all Ports, set program Counter 0 to count requests from Port 1 (see previous example) and program Counter 1 to count all request from all Ports by setting *Event\_Sel\_1* to 0 (Request\_Count) and set *all* bits in the *CM Performance Counter Qualifier 1 Register* to 1.

The cycle counter can be used to calculate the average rates of specified events. Continuing the above example, assuming the cycle counter is reset, started, and stopped simultaneously with the two performance counters, then the rate of requests from port 1 and all ports can be easily computed (value of each performance counter / value in cycle counter).

### 17.1.2 Performance Counter Usage Models

There are several model for using the CM performance counters. This sections discusses 3 possible models:

- Periodic Sampling - take many measurement samples of specific duration
- Stop and Interrupt when counter overflows - counters run until one overflows, then interrupt CPU
- Large count capability - enables unrestricted sample periods

One model for making performance measurements is for the software to set up and gather samples for a set period of time. The code sequence could follow the following steps:

```
start:
Write CM Event and Qualifier Registers for particular event of interest
Write CM Performance Counter Control Register to reset and start counters
    Perf_Int_En = 0 (no interrupt on overflow)
    Perf_Ovf_Stop = 0(no stop on overflow).
    P1_Reset = 1, P1_CountOn = 1
    P0_Reset = 1, P0_CountOn = 1
    Cycl_Cnt_Reset = 1, Cycl_Cnt_CountOn = 1
Wait for some relatively small period of time (i.e., 2 seconds)
Write CM Performance Counter Control Register to stop counters
    P1_Counton = 0, P0_CountOn=0, Cycl_Cnt_CountOn = 0
Read CM Performance Counter 0, Counter 1, and Cycle Counter Registers
If more events, go to start (or if measuring same counter go to step 2 instead)
```

A second CM performance counter usage model involves setting up the counters to stop and interrupt on overflow. This runs the counters until one of the counters (usually the cycle counter) reaches the 32-bit limit. An example of such a code sequence is:

```
start:
Write CM Event and Qualifier Registers for particular event of interest
Write CM Performance Counter Control Register to reset and start counters
    Perf_Int_En = 1 (interrupt on overflow)
    Perf_Ovf_Stop = 1(stop on overflow).
    P1_Reset = 1, P1_CountOn = 1
    P0_Reset = 1, P0_CountOn = 1
    Cycl_Cnt_Reset = 1, Cycl_Cnt_CountOn = 1
When interrupt occurs:
Read CM Performance Counter Status Register
Read CM Performance Counter 0, Counter 1, and Cycle Counter Registers
Write CM Performance Counter Control Register to reset counters
    (clears status register and interrupt)
    P0_Reset = 1, P1_Reset = 1, Cycl_Cnt_Reset = 1
If more events, go to start (or if measuring same counter go to step 2 instead)
```

If larger counts than can fit into the 32-bit counters are required, the counters can be set up to interrupt, but not stop, on overflow. Memory variables can then count the number of overflows, as shown below:

```
start:
Write CM Event and Qualifier Registers for particular event of interest
Write CM Performance Counter Control Register to reset and start counters
    Perf_Int_En = 1 (interrupt on overflow)
    Perf_Ovf_Stop = 0 (do not stop on overflow).
    P1_Reset = 1, P1_CountOn = 1
    P0_Reset = 1, P0_CountOn = 1
    Cycl_Cnt_Reset = 1, Cycl_Cnt_CountOn = 1
When interrupt occurs:
<status>=Read CM Performance Counter Status Register
Increment <overflow_count>[counter] for each counter with <status> = 1
Write <status> to CM Performance Counter Status Register to clear interrupt
```

```
When run limit is reached then :
Write CM Performance Counter Control Register to stop counters
    P1_Counton = 0, P0_CountOn=0, Cycl_Cnt_CountOn = 0
Read CM Performance Counter 0, Counter 1, and Cycle Counter Registers
Write CM Performance Counter Control Register to reset counters
    (clears status register and interrupt)
    P0_Reset = 1, P1_Reset = 1, Cycl_Cnt_Reset = 1
If more events, go to start (or if measuring same counter go to step 2 instead)
```

In the above model, the final counts are calculated for each counter by multiplying <overflow\_count>[counter] by 4G and adding the final values in the performance counter register.



### 17.1.3 CM Performance Counter Event Types and Qualifiers

This section describes the Performance Counter Event Types and associated qualifiers.

**Table 17.1 CM Performance Counter Event Types**

Event #	Related Events	Use	Qualifiers	Description/Comments
0	Request_Count	Measuring Load	Request Port Request CCA Request Cmd Request Length Request Target See <a href="#">Table 17.2</a>	Can be used in conjunction with a cycle count to determine number of requests received in a given period of time.
1	Coh_Req_Resp	Track coherent requests or responses, and measure sharing	Intervention State Speculation Intervention Cmd Store Conditional See <a href="#">Table 17.3</a>	Gives a count of the specified coherent request and response types.
2	L2_WR_Data_Util	L2 Write Data Bus Usage	Accept State See <a href="#">Table 17.4</a>	Counts number of cycles the L2/Memory write data bus is occupied. The qualifier determines if stall cycles are counted or not.
3	L2_Cmd_Util	L2 Command Bus Usage	Accept State See <a href="#">Table 17.4</a>	Counts number of cycles the L2/Memory command data bus is occupied. The qualifier determines if stall cycles are counted or not.
4	L2_RD_Data_Util	L2 Read Data Bus Usage	None	Counts number of cycles the L2/Memory read data bus is occupied.
5	Sharing_Miss	Sharing Frequency	Request Source Port Data Source Port See <a href="#">Table 17.5</a>	Counts source of data for coherent read requests only (i.e., CohReadShare, CohReadDiscard, CohReadOwn, and CohReadAlways).  Useful to determine how many cache misses were satisfied by other processors.
6	RSU_Util	RSU Usage	Port to measure Response Type See <a href="#">Table 17.6</a>	Counts number of d-words on the processor/iocu read data bus. A counter can only measure one port at a time. The port number is specified as the qualifier.
8	L2_Util	L2 Pipeline Usage	L2 Pipeline starts See <a href="#">Table 17.7</a>	Counts starts into the TA stage of the L2 pipeline.
9	L2_Hit	L2 Hit/Miss Usage	Hit/Miss Type Source Port See <a href="#">Table 17.8</a>	Counts different types of L2 Cache Hits and Misses, crossed with Source Port ID.
16	IOCU_Request	IOCU Request	Transaction ID I/O Parking CM Transaction Cnt BurstLength L2 allocation Posted Cacheability Request Type See <a href="#">Table 17.9</a>	Counts requests receive by the IOCU. The CM receives a sideband signal, SI_CMP_IOC_PerfInfo from the IOCU as described in <a href="#">Table 17.9</a> .

**Table 17.1 CM Performance Counter Event Types (continued)**

Event #	Related Events	Use	Qualifiers	Description/Comments
17	IOCU1_Request	2nd IOCU Request	Transaction ID I/O Parking CM Transaction Cnt BurstLength L2 allocation Posted Cacheability Request Type See <a href="#">Table 17.9</a>	Counts requests receive by the 2nd IOCU. The CM receives a sideband signal, SI_CMP_IOC1_PerfInfo from the 2nd IOCU as described in <a href="#">Table 17.9</a> .

**Table 17.2 CM Performance Counter Request Count Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31	Request Port	Port 7	Request originated from port 7
30		Port 6	Request originated from port 6
29		Port 5	Request originated from port 5
28		Port 4	Request originated from port 4
27		Port 3	Request originated from port 3
26		Port 2	Request originated from port 2
25		Port 1	Request originated from port 1
24		Port 0	Request originated from port 0
23	Request CCA <sup>1</sup>	WT	Request had Write Through Cacheability Attribute
22		UC/UCA	Request had Uncached Cacheability Attribute
21		WB	Request had Cached (non-coherent) Attribute
20		CWBE	Request had Coherent (Exclusive) Attribute
19		CWB	Request had Coherent (Shared) Attribute
18	Burst Length <sup>2</sup> (# of dwords)	1 dword	Request was for 1 dword of data Note: This counts the burst length as seen by the Coherent Manager. Requests from the I/O Subsystem may be longer, but the IOCU may break these into multiple smaller requests.
17		2 dwords	Request was for 2 dwords of data See Note for 1 dword.
16		4 dwords	Request was for 4 dwords of data See Note for 1 dword

**Table 17.2 CM Performance Counter Request Count Qualifier (continued)**

Bit	Qualifier Group	Qualifier Value	Description/Comments
15	Request Command	Legacy WR	Request is a legacy Write command. This is used for all non-coherent writes. Note: When a processor is in coherent mode, L1 cache writebacks are always considered coherent, so they result in a cohWriteBack command, not a WR command.
14		Legacy RD	Request is a legacy Read command. This is used for all non-coherent reads, including code fetches.
13		CohReadShare CohReadShareAlways	Request is a coherent read share generated by the processor on a load that misses its L1 cache. Currently CohReadShareAlways is unused.
12		CohReadOwn	Request is a coherent read own generated by the processor on a store that misses its L1 cache.
11		CohReadDiscard	Request is a coherent read discard generated by the IOCU for coherent requests.
10		CohUpgrade	Request is a coherent upgrade request generated by the the processor on a store that hits a shared line in its L1 cache.
9		CohWriteBack	Request is coherent writeback generated by the processor when evicting a line from the L1 cache. The line may have been installed in the cache from a coherent or non-coherent transaction.
8		CohWriteInval (Partial Line)	Request is a coherent write invalidate (not a full line of data) generated by the IOCU.
7		CohWriteInval (Full Line)	Request is a coherent write invalidate (full line of data) generated by the IOCU.
6		CohInvalidate	Request is an invalidate request from a processor executing a PREF Prepare for Store or a CACHE Hit Invalidate.
5		CohCopyBack	Request from a processor executing a CACHE hit writeback
4		CohCopyBackInv	Request from a processor executing a CACHE hit CACHE Write-BackInvalidate
3		CohCompletionSync	Request is from a processor executing a SYNC instruction
2		Target	Memory
1	GCR/GIC/CPC		Request targets the Interrupt controller or Global Control Registers
0	MMIO		Request targets Memory Mapped I/O space

1. CCA qualifier group is ignored on non-coherent cache-ops
2. Burst Length only used when Request Command is Legacy Read, Legacy Write, CohReadDiscard or CohWriteInval.

**Table 17.3 CM Performance Counter Coherent Request/Response Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31:25	Reserved		

**Table 17.3 CM Performance Counter Coherent Request/Response Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
24	Intervention State	Exclusive with data	A processor has an exclusive copy in its L1 cache and returned data (all commands except CohInvalidate)
23		Exclusive with no data	A processor has an exclusive copy in its L1 cache but no data was returned (occurs on a CohInvalidate)
22		Modified with data	A processor has a modified copy in its L1 cache and returned data (all commands except CohInvalidate)
21		Modified with no data	A processor has a modified copy in its L1 cache but no data was returned (occurs on a CohInvalidate)
20		Shared	One or more processors have a shared copy in its L1 cache
19		Invalid	No processor has a copy of the data in its L1 cache
18	Speculation	Speculate	Request was a CohReadShare, CohReadOwn, CohReadDiscard or CohReadAlways and the CM issued a speculative read request to L2/Memory. This qualifier group is ignored when the request is not one of the commands listed above.
17		No Speculate	Request was a CohReadShare, CohReadOwn, CohReadDiscard or CohReadAlways and the CM did not issue a speculative read request to L2/Memory. This qualifier group is ignored when the request is not one of the commands listed above.
16	Intervention Cmd	Reserved	Currently a don't care.
15		Reserved	Currently a don't care.
14		CohReadShare	Request is a coherent read share generated by the processor on a load that misses its L1 cache.
13		CohReadShareAlways	Currently CohReadShareAlways is unused.
12		CohReadOwn	Request is a coherent read own generated by the processor on a store that misses its L1 cache.

**Table 17.3 CM Performance Counter Coherent Request/Response Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
11	Intervention Cmd (cont.)	CohReadDiscard	Request is a coherent read discard generated by the IOCU for coherent requests.
10		CohUpgrade (OK Response)	Request is a coherent upgrade request generated by the processor on a store that hits a shared line in its L1 cache. There is no intervening request to the same line so an OK response is given.
9		CohUpgrade (Data Response)	Request is a coherent upgrade request generated by the processor on a store that hits a shared line in its L1 cache. There is an intervening request to the same line so a data response is given.
8		CohWriteBack	Request is coherent writeback generated by the processor when evicting a line from the L1 cache. The line may have been installed in the cache from a coherent or non-coherent transaction.
7		CohWriteInval (Partial Line)	Request is a coherent write invalidate (not a full line of data) generated by the IOCU.
6		CohWriteInval (Full Line)	Request is a coherent write invalidate (full line of data) generated by the IOCU.
5		CohInvalidate	Request is an invalidate request from a processor executing a PREF Prepare for Store or a CACHE Hit Invalidate.
4		CohCopyBack	Request from a processor executing a CACHE hit writeback
3		CohCopyBackInv	Request from a processor executing a CACHE hit CACHE WriteBackInvalidate
2		Store Conditional (only used when cmd is CohUpgrade or CohReadOwn)	Not due to a Store Conditional
1	Store Conditional that was not Cancelled		CohUpgrade or CohReadOwn is due a store conditional instruction and the intervention was not cancelled. This qualifier group is ignored when the command is not a CohUpgrade or CohReadOwn.
0	Store Conditional that was Cancelled		CohUpgrade or CohReadOwn is due a store conditional instruction and the intervention was cancelled due to livelock avoidance scheme. This qualifier group is ignored when the command is not a CohUpgrade or CohReadOwn.

**Table 17.4 CM Performance Counter Accept State Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31:1	Reserved		
0	Accept State	Count_Stalls	Setting this value to 0 for the L2_WR_Data_Util or L2_Cmd_Util events cause a count of cycles when a data word or command is accepted by the L2/Memory.  Setting this value to 1 for L2_WR_Data_Util or L2_Cmd_Util cause a count of cycles when a data word or command is valid on the bus, i.e., the count includes cycles where the command or data bus is stalled.

**Table 17.5 CM Performance Counter CM Data Source Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31:15	Reserved		
14	Request Port	7	Request originated from port 7
13		6	Request originated from port 6
12		5	Request originated from port 5
11		4	Request originated from port 4
10		3	Request originated from port 3
9		2	Request originated from port 2
8		1	Request originated from port 1
7		0	Request originated from port 0
6	Response Port	5	Data returned by processor connected to port 5
5		4	Data returned by processor connected to port 4
4		3	Data returned by processor connected to port 3
3		2	Data returned by processor connected to port 2
2		1	Data returned by processor connected to port 1
1		0	Data returned by processor connected to port 0
0		L2/Mem	Data returned by L2/Memory

**Table 17.6 CM Performance Counter CM Port Response Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31:6	Reserved		
5	Response Type	Read Data Response	Response was a dword of data.
4		Write Acknowledge Response	Response was a write acknowledge (DVA response for a write).
3		OK Response	Response was an OK response (due to a CohUpgrade).
2:0	Port Number	Port to measure	Encoded value of port number to measure. For example, a value of 2 will only count responses on response port 2.

**Table 17.7 L2 Utilization Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31:6	Reserved		

**Table 17.7 L2 Utilization Qualifier (continued)**

Bit	Qualifier Group	Qualifier Value	Description/Comments
5	Pipeline Start Type	L2 Pipeline start was stalled	Any type of pipeline request start (new, replay,refill) was refused due to a stall (ram or global stall)
4		L2 Pipeline start is taken	Use to calculate L2 utilization Any type of pipeline request start (new, replay,refill)
3		New request waiting for Sync to clear	A new request is waiting to be dispatched to the L2 until a preceeding Sync has guaranteed ordering
2		New L2 request stalled	New request to the L2 was not accepted due to a stall (ram or global stall)
1		New L2 request denied	New request to the L2 was not accepted due to replay, refill, or a stall.
0		New L2 request started	Use to calculate L2 bandwidth

**Table 17.8 L2 Hit Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31:22	Reserved		
21	Partial Write Merge (write misses only)	Partial write merged	Partial write was merged into the SRT.
20		Partial write not merged	Partial write was not merged into the SRT.
19	Allocation (for Write or Read misses only)	Line allocated	A miss caused an allocation by the L2. This occurs either for a full line write miss or a read miss, depending on the L2 allocation policy.
18		Line not allocated	A miss did not cause an allocation by the L2.
17	Hit/Miss Type (these are mutually exclusive)	Other	Index L2 cacheop or Fetch & Lock.
16		Non-index cache-op hit	Non-index L2 cacheop hit the L2 cache.
15		Non-index cache-op miss	Non-index L2 cacheop missed the L2 cache.
14		Write hit without RMW	Write hit and update the L2 without a read modify write operation. When the data byte enables of all DWords are either all 0 or all 1, an L2 write operation is performed without doing a read modify write operation.
13		Write hit with RMW	Write hit that requires a read modify write operation. When the data byte enable of any DWord is not all 1, a read modify write operation is required.
12		Write miss, no memory read	Write miss that does not require a memory read request.
11		Write miss requiring memory read	Write miss that does require a memory read request.
10		Read into CRQ	Read matched a pending L2 miss. Data is returned when the pending line is refilled. It is not a Read hit or a Read miss.
9		Read hit	Read hit the L2 cache.
8		Read miss	Read missed the L2 cache. Either allocates or reads through to memory, depending on the L2 allocation policy.

**Table 17.8 L2 Hit Qualifier(continued)**

Bit	Qualifier Group	Qualifier Value	Description/Comments
7	Source Port	7	Request originated from port 7
6		6	Request originated from port 6
5		5	Request originated from port 5
4		4	Request originated from port 4
3		3	Request originated from port 3
2		2	Request originated from port 2
1		1	Request originated from port 1
0		0	Request originated from port 0

**Table 17.9 IOCU Performance Counter Request Count Qualifier**

Bit	Qualifier Group	Qualifier Value	Description/Comments
31	Reserved		
30:27	Transaction ID	TID	Value of IC_MTagID to match when the All_TID qualifier bit is set to 0. This field is unused when All_TID is 1.
26		All_TID	If 1 then the all values of IC_MTagID will match. If 0 then only transactions with IC_MTagID equal to the TID specified above will match.
25	I/O Parking	Start and Stop Parking	Request will start and stop I/O Parking.
24		Stop Parking	Request will stop I/O parking (but not start it).
23		Start Parking	Request will start I/O Parking (but not stop it).
22		No parking	Request will not start or stop I/O parking.
21	CM Transaction Count	5 CM Transactions	Request resulted in 5 CM transactions.
20		4 CM Transactions	Request resulted in 4 CM transactions.
19		3 CM Transactions	Request resulted in 3 CM transactions.
18		2 CM Transactions	Request resulted in 2 CM transactions.
17		1 CM Transaction	Request resulted in 1 CM transaction.
16	BurstLength	13-16	IC_MBurstLength is 13, 14, 15, or 16 dwords.
15		9-12	IC_MBurstLength is 9, 10, 11, or 12 dwords.
14		5-8	IC_MBurstLength is 5, 6, 7, or 8 dwords.
13		4	IC_MBurstLength is 4 dwords.
12		3	IC_MBurstLength is 3 dwords.
11		2	IC_MBurstLength is 2 dwords.
10		1	IC_MBurstLength is 1 dword.



**Table 17.9 IOCU Performance Counter Request Count Qualifier(continued)**

Bit	Qualifier Group	Qualifier Value	Description/Comments
9	L2 Allocation	L2 Allocation with Prepare for Store	Request will cause an L2 allocation and the request is a write with L2 Prepare For Store. This bit will never cause a match for read requests.
8		L2 Allocation without Prepare for Store	Request will cause an L2 allocation and the request is either a read or a write with L2 Prepare For Store not asserted.
7		No L2 Allocation	Request will not cause an L2 allocation.
6	Posted	Non-posted Write	Write is non-posted. Not used on reads.
5		Posted Write	Write is posted. Not used on reads.
4	Cacheability	Uncached	Request is uncached.
3		Cached	Request is Cached, non-coherent.
2		Coherent	Request is Coherent.
1	Request Type	Read	Request is a read.
0		Write	Request is a write.

## 17.2 Debug Mode Triggering

This section describes the how to control the cores when entering debug mode.

### 17.2.1 Selecting CPUs to Enter Debug Mode

The interAptiv Multiprocessing System contains a set of registers and logic that controls when the interAptiv cores enter Debug mode. The logic allows software to:

- Specify which interAptiv core enters debug mode on assertion of the *EJ\_DINT\_IN* signal (generally asserted by a debug probe).
- Force one or more interAptiv cores to enter debug mode by writing to the *DINT Send to Group Register*.

### 17.2.2 Debug Mode Groups and Cross Triggering

The interAptiv Multiprocessing System (MPS) allows software to define debug mode groups so that when one interAptiv core enters debug mode, all other cores within the group also enter debug mode.

Software creates debug mode groups by writing to each VPE's *VPE-Local DebugBreak Group Register*. Each bit in the *Join\_DebugM* field of the *VPE-Local DebugBreak Group Register* represents a VPE in the system. If the bit is set, the corresponding VPE will enter debug mode. If the bit is clear, the corresponding VPE is not affected by Debug Mode.

Only the positive edge of a VPE's *EJ<cpu>\_DebugM* signal can cause the other CPUs to also enter the Debug Mode as a group. When there is no positive edge on the *DebugM* signals, the *Join\_DebugM* fields in the *DebugBrk\_Group* registers can be written without causing spurious glitches on the *EJ<cpu>\_DINT* signals.

The bit which represents the local VPE cannot be used to disable Debug Mode for the local VPE. For example, if the local VPE is represented by bit *i*, clearing bit *i* will *not* disable Debug Mode for the local VPE.

Because each VPE has its own *DebugBreak Group Register*, there could be multiple debug groups of cores. One group of VPEs might be communicating with one debugger, while another group of CPUs can be communicating with a different debugger. This is useful when each group of VPEs is running different applications or operating systems.

### 17.2.3 VPE to CPU mapping in Debug Cross Trigger Facility

By convention, VPEs are numbered in ascending order, based on CPU numbers:

For example, a 4-CPU system, having 1 VPE per CPU, will be mapped as follows:

$$\{\text{VPE0, VPE1, VPE2, VPE3}\} \equiv \{\text{CPU0, CPU1, CPU2, CPU3}\}.$$

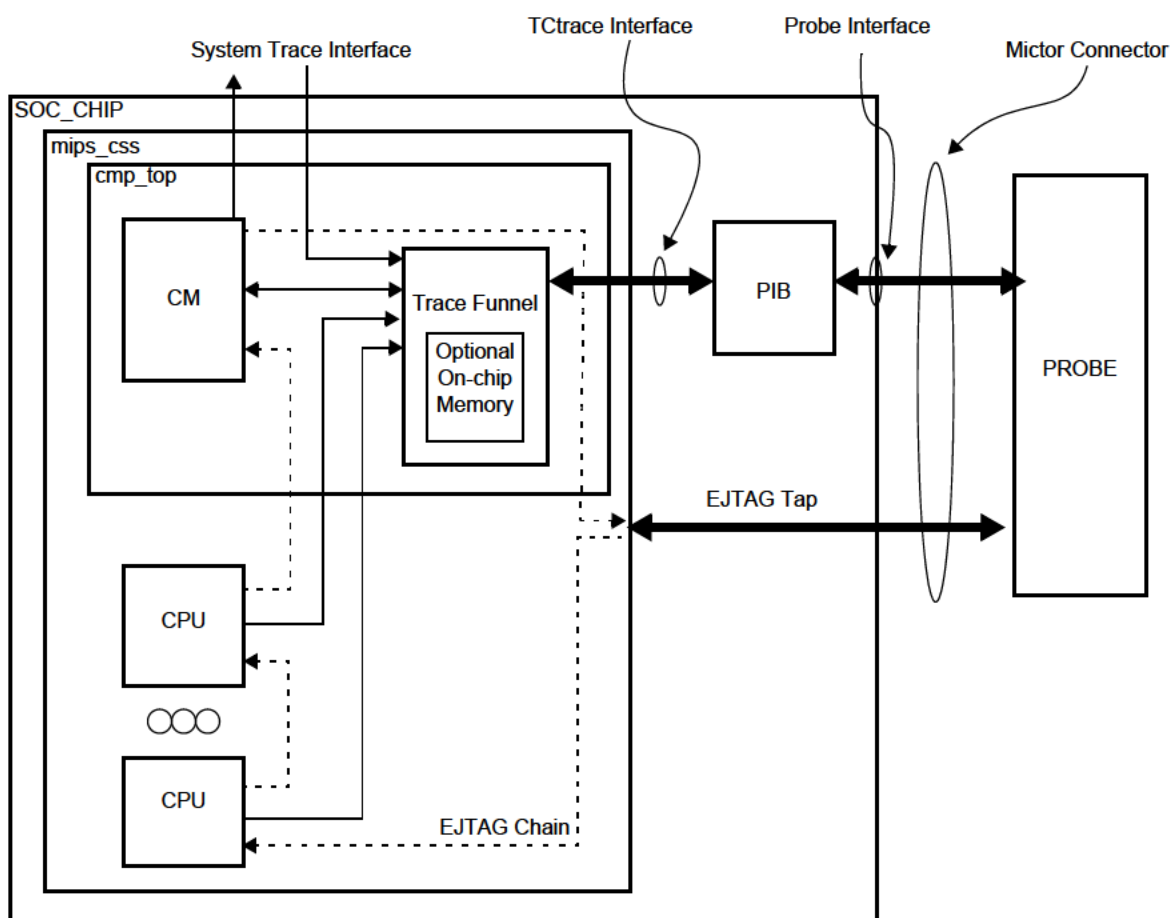
### 17.2.4 Debug Cross Trigger Facility and Power Management

Due to power management of interAptiv Multiprocessing System components, CPUs might not be powered or clocked when receiving a DINT via the debug cross trigger facility. However, the power controller observes all DINT events and will start up domains as requested. Depending on the programming of the power controller and time constants of the physical design, a delay between DINT event and a target CPU participating in the debug session might occur. To inquire about the current power status of a CPU, the debug handler can poll the power controller status registers. Generally, an EJTAG debug probe attached and recognized by the system will shorten the wake-up delay, while debug events without debug probe attachment might show more wake-up latency.

## 17.3 PDTrace Software Architecture

The interAptiv MPS enables debug trace information from the interAptiv cores, the Coherence Manager, and a System Trace Interface to be streamed off chip or stored in on-chip RAM. As shown in [Figure 17.1](#), each interAptiv core produces a 64-bit debug trace stream describing its program and data flow. The CM produces a 64-bit stream describing the flow of transactions within the CM. If a System Trace Interface is part of the build, it captures a 128-bit stream describing activity supplied externally by the System. The Trace Funnel muxes the CPU, CM, and System Trace streams into a single debug trace stream which is either stored in an on-chip buffer or passed onto a Probe Interface Block (PIB). A PIB is the on-chip link between the Trace Funnel and debug probe interface, and may include functionality such as time multiplexing the 64-bit TTrace data onto a narrower, slower probe interface.

**Figure 17.1 PD Trace Architecture**



The TTrace stream consists of 64-bit trace words (TW). Each trace word trace is packed with one or more Trace Formats (TF). There are many trace format types produced by CPUs and the CM. The CPU TFs allow tracing of information such as the program counter, load/store addresses, and load/address data values. The CM TFs produce information such as the serialization order of requests and the results of L1 cache interventions. See *The PDtrace® Interface and Trace Control Block Specification* for a detailed description of the TW and TF formats.

The trace output of each CPU can be controlled by a set of EJTAP accessible registers located in the Trace Control Block (TCB) associated with that CPU. Refer to *The MIPS32® interAptiv™ CPU Family Software User's Manual* for a detailed description of these registers and the CPU PDTrace functionality.

### 17.3.1 CM Trace Functionality

This section describes the configuration and functionality of the CM debug trace.

#### 17.3.1.1 CM Trace Configuration and Control

The CM Trace is controlled by the *CM TCBCONTROLD Register* as defined in [Section "TCBCONTROLD Register"](#). The enabling of the CM's Trace is determined by two fields in this register along with a field in each of the core's TCBCONTROLD register. [Figure 17.10](#) shows that there are two ways to enable CM Trace. First, CM Trace can be enabled independent of the Cores' state by setting both *CM\_EN* and *Global\_CM\_En* in the CM's *TCBCONTROLD*

Register. Alternatively, by setting *CM\_EN* and clearing *Global\_CM\_En*, the CM will only trace if at least one other core is tracing, i.e., *Core\_CM\_En* in at least one core's TCBCONTROLD register is set to 1. A core's *Core\_CM\_En* bit may be asserted/deasserted based on debug triggers as defined in *The MIPS32® interAptiv™ Processor Core Family Software User's Manual*. The value of each core's *Core\_CM\_En* bit is communicated to the CM on the *TC<core>\_Trace\_CM\_En* signal.

**Table 17.10 CM Trace Enable**

CM TCBCONTROLD Reg		Cores' TCBCONTROLD Reg	CM PDTrace Enabled/Disabled
CM_EN	Global_CM_En	Core_CM_En	
0	x	x	Disabled
1	1	x	Enabled
1	0	All 0	Disabled
1	0	not All 0	Enabled

### 17.3.1.2 System Trace Interface Configuration and Control

The System Trace Interface stream is generated and controlled by external logic. The CM has control output pins to support design of this logic. There are 2 specific control outputs and one 32-bit user-defined output. These outputs and the trace data/control pins associated with the trace stream are shown in [Table 17.11](#). All the signals are timed relative to the *SI\_CMClk*.

**Table 17.11 System Trace Interface Stream and Control Pins**

Signal	Direction/Type	Usage
SI_TC_Sys_Data[127:0]	CM stream input	System Trace stream data for 128-bit stream SI_TC_Sys_Data[71:68] must contain a Source Port ID and SI_TC_Sys_Data[7:4] must contain a Source Port ID. Legal values of either Source Port ID are: 4'hc or 4'hd. All other bits are completely user defined
SI_TC_Sys_Valid[1:0]	CM stream input	System Trace stream valid bits for upper and lower streams Bit 1 qualifies SI_TC_Sys_Data[127:64] Bit 0 qualifies SI_TC_Sys_Data[63:0] A value of 2'b10 is illegal
SI_TC_Sys_Stall	CM stream output	System Trace stream flow control.
SI_TC_Sys_Enable	CM control output	System Trace control advisory, driven from the <i>CM TCBCONTROLD<sub>ST_En</sub></i> . Its purpose is to advise the external logic of the state of this control bit. If desired, external logic can stop generation of the stream if this output is a zero, and allow generation of the stream if it is a 1. However, external logic may choose to continue sending stream data after de-assertion until it has flushed all its collected stream data.
<u>SI_TC_Sys_AnyCore_Enabled</u>	<u>CM control output</u>	<u>System Trace control advisory that at least one core is enabled to trace, derived from Cores' TCBCONTROLD</u>
<u>SI_TC_Sys_CM_Enabled</u>	<u>CM control output</u>	<u>System Trace control advisory that the CM2 is enabled to trace, derived from CM2's TCBCONTROLD</u>

**Table 17.11 System Trace Interface Stream and Control Pins**

Signal	Direction/Type	Usage
SI_TC_Sys_UserCtl[31:0]	CM control output	User defined control advisory bits, from TCBSYS. Bit 31 is a 1 when the Trace Funnel was configured with the System Trace present and is a 0 when the System Trace is not present. Bits [30:0] are completely user defined output values.

In addition to the System Trace Interface pins, there are internal control register bits that impact operation of the System Trace stream. Assertion of *CM TCBCONTROLB<sub>STCE</sub>* allows the System Trace funnel port to capture stream data; de-assertion of this bit causes the Trace Funnel to stop capturing the System Trace stream from within the Trace Funnel in case the external logic is problematic. In addition, de-assertion of *CM TCBCONTROLB<sub>EN</sub>* stops capture of all the streams (Cores, CM, System).

Thus the System Trace stream is enabled to capture the System Trace stream when these controls are asserted: *CM TCBCONTROLB<sub>STCE</sub>* and *CM TCBCONTROLB<sub>EN</sub>*. The control outputs *SI\_TC\_Sys\_Enable* and *SI\_TC\_Sys\_UserCtl[31:0]* are available to the external logic to further control generation of the System Trace stream by allowing or disallowing assertion of the *SI\_TC\_Sys\_Valid[1:0]* inputs. If any trace stream is being generated without enabling that stream to capture, then that stream is not captured and the data is dropped.

### 17.3.1.3 Trace Funnel Enable

When trace on the System, CM and/or Cores is enabled then trace information is continuously sent to the Trace Funnel. However, the trace funnel will only send the trace information to the trace probe or to the on-chip trace memory if it is enabled by setting the *CM TCBCONTROLB<sub>EN</sub>* bit. The Trace Funnel can be subsequently disabled by clearing the *CM TCBCONTROLB<sub>EN</sub>* bit. See “TCBCONTROLB Register Field Descriptions” on page 772 for more information.

### 17.3.1.4 CM Trace Formats

Trace information is captured at two points within the CM:

- Information about requests is captured by the Request Unit (RQU) after serialization, thus providing a view of the global order of requests.
- Information about L1 interventions is captured by the Intervention Unit (IVU) after all intervention responses have been received. This provides information about the state of the cache line in all L1 caches for coherent requests.

The type and amount of content in each Trace Format created by the CM depends on the source of the packet (RQU or IVU) and the configuration (TL<sub>ev</sub>, AE, P<port>\_Ctl control bits). Refer to *The PDtrace™ Interface and Trace Control Block Specification* for the detailed description of the CM Trace Formats.

### 17.3.1.5 CM / CPU Core Trace Correlation

In the interAptiv core, trace information is provided from each of the cores as well as the Coherence Manager. In order to correlate transactions from the CM to the instruction stream, an identifier is used in both the core and CM traces.

The CM trace includes the core ID and CosID for each request. The CosID changes relatively slowly - it is generally incremented after PCSync in the core or if an overflow is detected in the CM. Typically several requests in a row will use the same CosID value, and the intermediate correlation is enabled by the requests appearing in the same order in

the CM and core traces. Because of this, and the fact that the CosID is traced as a part of the instruction completion record, correlating instructions to CM transactions is possible only when PC tracing is enabled for all TCs executing on the core.

*The PDtrace™ Interface and Trace Control Block Specification* includes a more detailed description of the correlation process.

## 17.3.2 Controlling Trace in a Multi-CPU Multiprocessing System

The interAptiv MPS enables debug trace information from the interAptiv cores and the Coherence Manager to be streamed off chip or stored in on-chip RAM. As shown in [Figure 17.1](#), each interAptiv core produces a 64-bit debug trace stream describing its program and data flow. The CM produces a stream describing the flow of transactions within the CM2. The Trace Funnel muxes the CPU and CM trace streams into a single debug trace stream which is either stored in an on-chip buffer or passed onto a Probe Interface Block (PIB). A PIB is the on-chip link between the Trace Funnel and debug probe interface, and may include functionality such as time multiplexing the 64-bit TCtrace data onto a narrower, slower probe interface.

Since the interAptiv core streams PDTrace data directly to the trace funnel, the core TCB system is configured as if only off-chip trace is present. Core TCB register bits which refer to control of on-chip trace resources will behave as if on-chip trace is not implemented.

The CM has its own set of TCBControl registers. It is designated as the ‘master’ which controls trace functionality for the CM, the on-chip trace buffer, and the PIB interface. In addition to the CM2 as trace master, the GCR block itself can function as the trace master in the interAptiv core. This is done through memory mapped CM\_GCR global control registers.

## 17.3.3 EJTAG Debug Support in the interAptiv Coherence Manager

The EJTAG debug logic in the Coherence Manager is compliant with EJTAG Specification 5.0 and includes:

1. Standard Test Access Port (TAP) for a dedicated connection to a debug host
2. Optional PDtrace capability for program counter/data address/data value trace to On-chip memory or to Trace probe

The following sub-sections describe the TAP and EJTAG operation and registers.

### 17.3.3.1 Test Access Port (TAP)

The following main features are supported by the TAP module:

- 5-pin industry standard JTAG Test Access Port (*TCK*, *TMS*, *TDI*, *TDO*, *TRST\_N*) interface which is compatible with IEEE Std. 1149.1.
- Target chip and EJTAG feature identification available through the Test Access Port (TAP) controller.

## EJTAG Internal and External Interfaces

The external interface of the EJTAG module consists of the 5 signals defined by the IEEE standard.

**Table 17.12 EJTAG Interface Pins**

Pin	Type	Description
<i>TCK</i>	I	Test Clock Input Input clock used to shift data into or out of the Instruction or data registers. The <i>TCK</i> clock is independent of the CM clock, so the EJTAG probe can drive <i>TCK</i> independently of the CM clock frequency. The CM signal for this is called <i>EJ_TCK</i>
<i>TMS</i>	I	Test Mode Select Input The <i>TMS</i> input signal is decoded by the TAP controller to control test operation. <i>TMS</i> is sampled on the rising edge of <i>TCK</i> . The CM signal for this is called <i>EJ_TMS</i>
<i>TDI</i>	I	Test Data Input Serial input data ( <i>TDI</i> ) is shifted into the Instruction register or data registers on the rising edge of the <i>TCK</i> clock, depending on the TAP controller state. The CM signal for this is called <i>EJ_TDI</i>
<i>TDO</i>	O	Test Data Output Serial output data is shifted from the Instruction or data register to the <i>TDO</i> pin on the falling edge of the <i>TCK</i> clock. When no data is shifted out, the <i>TDO</i> is 3-stated. The CM signal for this is called <i>EJ_TDO</i> with output enable controlled by <i>EJ_TDOzstate</i> .
<i>TRST_N</i>	I	Test Reset Input (Optional pin) The <i>TRST_N</i> pin is an active-low signal for asynchronous reset of the TAP controller and instruction in the TAP module, independent of the main CM logic. The CM's transaction processing logic is not reset by the assertion of <i>TRST_N</i> . The CM signal for this is called <i>EJ_TRST_N</i> This signal is optional, but power-on reset must apply a low pulse on this signal at power-on and then leave it high, in case the signal is not available as a pin on the chip. If available on the chip, then it must be low on the board when the EJTAG debug features are unused by the probe.

### Test Access Port Operation

The TAP controller is controlled by the Test Clock (*TCK*) and Test Mode Select (*TMS*) inputs. These two inputs determine whether an the Instruction register scan or data register scan is performed. The TAP consists of a small controller, driven by the *TCK* input, which responds to the *TMS* input as shown in the state diagram in [Figure 17.2](#). The TAP uses both clock edges of *TCK*. *TMS* and *TDI* are sampled on the rising edge of *TCK*, while *TDO* changes on the falling edge of *TCK*.

At power-up the TAP is forced into the *Test-Logic-Reset* by low value on *TRST\_N*. The TAP instruction register is thereby reset to IDCODE. No other parts of the EJTAG hardware are reset through the *Test-Logic-Reset* state.

When test access is required, a protocol is applied via the *TMS* and *TCK* inputs, causing the TAP to exit the *Test-Logic-Reset* state and move through the appropriate states. From the *Run-Test/Idle* state, an Instruction register scan or a data register scan can be issued to transition the TAP through the appropriate states shown in [Figure 17.2](#).

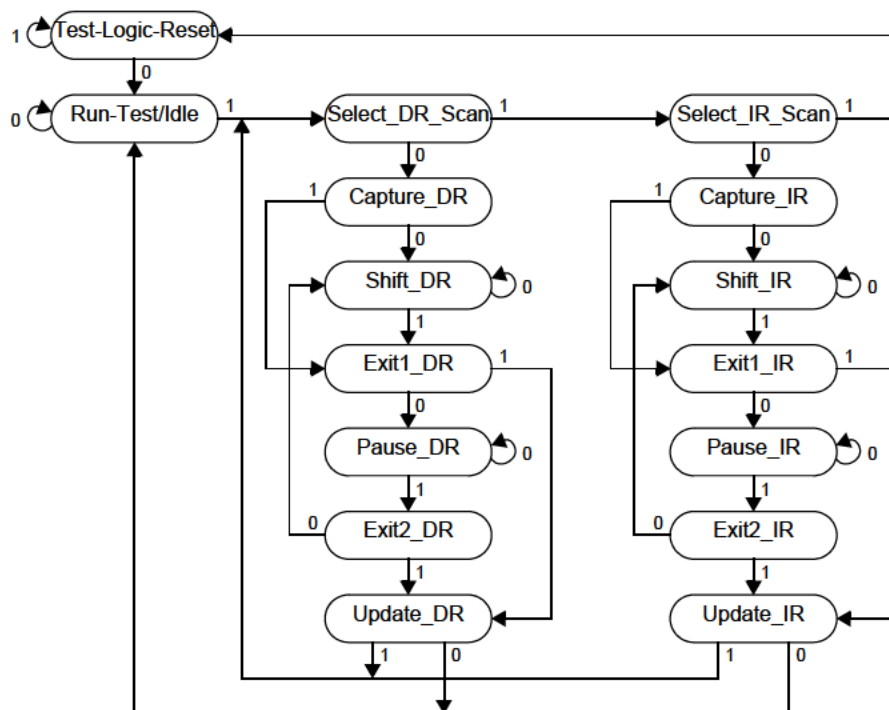
The states of the data and instruction register scan blocks are mirror images of each other adding symmetry to the protocol sequences. The first action that occurs when either block is entered is a capture operation. For the data registers, the *Capture-DR* state is used to capture (or parallel load) the data into the selected serial data path. In the Instruction register, the *Capture-IR* state is used to capture status information into the Instruction register.



From the *Capture* states, the TAP transitions to either the *Shift* or *Exit1* states. Normally the *Shift* state follows the *Capture* state so that test data or status information can be shifted out for inspection and new data shifted in. Following the *Shift* state, the TAP either returns to the *Run-Test/Idle* state via the *Exit1* and *Update* states or enters the *Pause* state via *Exit1*. The reason for entering the *Pause* state is to temporarily suspend the shifting of data through either the Data or Instruction Register while a required operation, such as refilling a host memory buffer, is performed. From the *Pause* state shifting can resume by re-entering the *Shift* state via the *Exit2* state or terminate by entering the *Run-Test/Idle* state via the *Exit2* and *Update* states.

Upon entering the data or Instruction register scan blocks, shadow latches in the selected scan path are forced to hold their present state during the *Capture* and *Shift* operations. The data being shifted into the selected scan path is not output through the shadow latch until the TAP enters the *Update-DR* or *Update-IR* state. The *Update* state causes the shadow latches to update (or parallel load) with the new data that has been shifted into the selected scan path.

**Figure 17.2 TAP Controller State Diagram**



### **Test-Logic-Reset State**

In the *Test-Logic-Reset* state the boundary scan test logic is disabled. The test logic enters the *Test-Logic-Reset* state when the *TMS* input is held HIGH for at least five rising edges of *TCK*. The *BYPASS* instruction is forced into the instruction register output latches during this state. The controller remains in the *Test-Logic-Reset* state as long as *TMS* is HIGH.

### **Run-Test/Idle State**

The controller enters the *Run-Test/Idle* state between scan operations. The controller remains in this state as long as *TMS* is held LOW. The instruction register and all test data registers retain their previous state. The instruction cannot change when the TAP controller is in this state.



When *TMS* is sampled HIGH on the rising edge of *TCK*, the controller transitions to the *Select\_DR* state.

### ***Select\_DR\_Scan State***

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, then the controller transitions to the *Capture\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Select\_IR* state. The instruction cannot change while the TAP controller is in this state.

### ***Select\_IR\_Scan State***

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW on the rising edge of *TCK*, the controller transitions to the *Capture\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Test-Reset-Logic* state. The instruction cannot change while the TAP controller is in this state.

### ***Capture\_DR State***

In this state the boundary scan register captures the value of the register addressed by the Instruction register, and the value is then shifted out in the *Shift\_DR*. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Shift\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_DR* state. The instruction cannot change while the TAP controller is in this state.

### ***Shift\_DR State***

In this state the test data register connected between *TDI* and *TDO* as a result of the current instruction shifts data one stage toward its serial output on the rising edge of *TCK*. If *TMS* is sampled LOW on the rising edge of *TCK*, the controller remains in the *Shift\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_DR* state. The instruction cannot change while the TAP controller is in this state.

### ***Exit1\_DR State***

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Pause\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Update\_DR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state.

### ***Pause\_DR State***

The *Pause\_DR* state allows the controller to temporarily halt the shifting of data through the test data register in the serial path between *TDI* and *TDO*. All test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW on the rising edge of *TCK*, the controller remains in the *Pause\_DR* state. A HIGH on *TMS* causes the controller to transition to the *Exit2\_DR* state. The instruction cannot change while the TAP controller is in this state.

### ***Exit2\_DR State***

This is a temporary controller state in which all test data registers selected by the current instruction retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Shift\_DR* state to allow another serial shift of data. A HIGH on *TMS* causes the controller to transition to the *Update\_DR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state.

### ***Update\_DR State***

When the TAP controller is in this state the value shifted in during the *Shift\_DR* state takes effect on the rising edge of the *TCK* for the register indicated by the Instruction register.

If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Run-Test/Idle* state. A HIGH on *TMS* causes the controller to transition to the *Select\_DR\_Scan* state. The instruction cannot change while the TAP controller is in this state and all shift register stages in the test data registers selected by the current instruction retain their previous state.

### ***Capture\_IR State***

In this state the shift register contained in the Instruction register loads a fixed pattern (00001<sub>2</sub>) on the rising edge of *TCK*. The data registers selected by the current instruction retain their previous state.

If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Shift\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_IR* state. The instruction cannot change while the TAP controller is in this state.

### ***Shift\_IR State***

In this state the instruction register is connected between *TDI* and *TDO* and shifts data one stage toward its serial output on the rising edge of *TCK*. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller remains in the *Shift\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Exit1\_IR* state.

### ***Exit1\_IR State***

This is a temporary controller state in which all registers retain their previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Pause\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Update\_IR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state and the instruction register retains its previous state.

### ***Pause\_IR State***

The *Pause\_IR* state allows the controller to temporarily halt the shifting of data through the instruction register in the serial path between *TDI* and *TDO*. If *TMS* is sampled LOW at the rising edge of *TCK*, the controller remains in the *Pause\_IR* state. A HIGH on *TMS* causes the controller to transition to the *Exit2\_IR* state. The instruction cannot change while the TAP controller is in this state.

### ***Exit2\_IR State***

This is a temporary controller state in which the instruction register retains its previous state. If *TMS* is sampled LOW at the rising edge of *TCK*, then the controller transitions to the *Shift\_IR* state to allow another serial shift of data. A HIGH on *TMS* causes the controller to transition to the *Update\_IR* state which terminates the scanning process. The instruction cannot change while the TAP controller is in this state.

### ***Update\_IR State***

The instruction shifted into the instruction register takes effect on the rising edge of *TCK*.

If *TMS* is sampled LOW at the rising edge of *TCK*, the controller transitions to the *Run-Test/Idle* state. A HIGH on *TMS* causes the controller to transition to the *Select\_DR\_Scan* state.

### 17.3.3.2 Test Access Port (TAP) Instructions

The TAP Instruction register allows instructions to be serially input into the device when TAP controller is in the *Shift-IR* state. Instructions are decoded and define the serial test data register path that is used to shift data between *TDI* and *TDO* during data register scanning.

The Instruction register is a 5-bit register. In the current EJTAG implementation only some instructions have been decoded; the unused instructions default to the BYPASS instruction.

**Table 17.13 Implemented EJTAG Instructions**

Value	Instruction	Function
0x01	IDCODE	Select Chip Identification data register.
0x03	IMPCODE	Select Implementation register.
0x08	Reserved	Instructions using this code select bypass register.
0x09	Reserved	Instructions using this code select bypass register.
0x0A	CONTROL	Select EJTAG Control register.
0x0B	Reserved	Instructions using this code select bypass register.
0x0C	Reserved	Instructions using this code select bypass register.
0x0D	Reserved	Instructions using this code select bypass register.
0x0E	Reserved	Instructions using this code select bypass register.
0x10	Reserved	Instructions using this code select bypass register.
0x11	TCBCONTROLB	Selects the <i>TCBCONTROLB</i> register in the Trace Control Block.
0x12	TCBDATA	Selects the <i>TCBDATA</i> register in the Trace Control Block.
0x13	Reserved	Instructions using this code select bypass register.
0x14	Reserved	Instructions using this code select bypass register.
0x15	TCBCONTROLD	Selects the <i>TCBCONTROLD</i> register in the Trace Control Block.
0x16	TCBCONTROLE	Selects the <i>TCBCONTROLE</i> register in the Trace Control Block.
0x17	Reserved	Instructions using this code select bypass register.
0x1F	BYPASS	Bypass register.

#### ***BYPASS Instruction***

The required BYPASS instruction selects the Bypass register to be connected between *TDI* and *TDO*. The BYPASS instruction allows serial data to be transferred through the CM from *TDI* to *TDO* without affecting its operation. The bit code of this instruction is defined to be all ones by the IEEE 1149.1 standard. Any unused instruction is defaulted to the BYPASS instruction.

#### ***IDCODE Instruction***

The IDCODE instruction selects the Device Identification (ID) register to be connected between *TDI* and *TDO*. The Device ID register is a 32-bit shift register containing information regarding the IC manufacturer, device type, and version code. Accessing the Identification Register does not interfere with the operation of the CM. Also, access to the Identification Register is immediately available, via a TAP data scan operation, after power-up when the TAP has been reset with on-chip power-on or through the optional *TRST\_N* pin.

### ***IMPCODE Instruction***

This instruction selects the Implementation register for output, which is always 32 bits.

### ***CONTROL Instruction***

This instruction is used to select the EJTAG Control register to be connected between *TDI* and *TDO*. The EJTAG Probe shifts 32 bits of *TDI* data into the EJTAG Control register and shifts out the EJTAG Control register bits via *TDO*.

### ***TCBCONTROLB Instruction***

This instruction is used to select the TCBCONTROLB register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

### ***TCBDATA Instruction***

This instruction is used to select the TCBDATA register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register. It should be noted that the TCBDATA register is only an access register to other TCB registers. The width of the TCBDATA register is dependent on the specific TCB register.

### ***TCBCONTROLD Instruction***

This instruction is used to select the TCBCONTROLD register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

### ***TCBCONTROLE Instruction***

This instruction is used to select the TCBCONTROLE register to be connected between *TDI* and *TDO*. This register is only implemented if the Trace Control Block is present. If no TCB is present, then this instruction will select the Bypass register.

## **17.3.3.3 EJTAG TAP Registers**

The EJTAG TAP Module has one Instruction register and a number of data registers, all accessible through the TAP:

### ***Instruction Register***

The Instruction register is accessed when the TAP receives an Instruction register scan protocol. During an Instruction register scan operation the TAP controller selects the output of the Instruction register to drive the *TDO* pin. The shift register consists of a series of bits arranged to form a single scan path between *TDI* and *TDO*. During an Instruction register scan operations, the TAP controls the register to capture status information and shift data from *TDI* to *TDO*. Both the capture and shift operations occur on the rising edge of *TCK*. However, the data shifted out from the *TDO* occurs on the falling edge of *TCK*. In the Test-Logic-Reset and *Capture-IR* state, the instruction shift register is set to 00001<sub>2</sub>, as for the IDCODE instruction. This forces the device into the functional mode and selects the Device ID register. The Instruction register is 5 bits wide. The instruction shifted in takes effect for the following data register scan operation. A list of the implemented instructions are listed in [Table 17.13](#).

### 17.3.3.4 Data Registers Overview

The EJTAG uses several data registers, which are arranged in parallel from the primary *TDI* input to the primary *TDO* output. The Instruction register supplies the address that allows one of the data registers to be accessed during a data register scan operation. During a data register scan operation, the addressed scan register receives TAP control signals to capture the register and shift data from *TDI* to *TDO*. During a data register scan operation, the TAP selects the output of the data register to drive the *TDO* pin. The register is updated in the *Update-DR* state with respect to the write bits.

This description applies in general to the following data registers:

- Bypass Register
- Device Identification Register
- Implementation Register
- EJTAG Control Register (ECR)

#### ***Bypass Register***

The *Bypass* register consists of a single scan register bit. When selected, the Bypass register provides a single bit scan path between *TDI* and *TDO*. The Bypass register allows abbreviating the scan path through devices that are not involved in the test. The Bypass register is selected when the Instruction register is loaded with a pattern of all ones to satisfy the IEEE 1149.1 Bypass instruction requirement.

#### ***Device Identification (ID) Register***

The *Device Identification* register is defined by IEEE 1149.1, to identify the device's manufacturer, part number, revision, and other device-specific information. Table 17.14 shows the bit assignments defined for the read-only Device Identification Register, and inputs to the CM determine the value of these bits. These bits can be scanned out of the *ID* register after being selected. The register is selected when the Instruction register is loaded with the IDCODE instruction.

**Figure 17.3 Device Identification Register Format**



**Table 17.14 Device Identification Register**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Version	31:28	<b>Version</b> (4 bits) This field identifies the version number of the CM.	R	<i>EJ_Version</i> [3:0]
PartNumber	27:12	<b>Part Number</b> (16 bits) This field identifies the part number of the CM.	R	<i>EJ_PartNumber</i> [15:0]
ManufID	11:1	<b>Manufacturer Identity</b> (11 bits) Accordingly to IEEE 1149.1-1990, the manufacturer identity code shall be a compressed form of the JEDEC Publications 106-A.	R	<i>EJ_ManufID</i> [10:0]
R	0	reserved	R	1

### Implementation Register

This 32-bit read-only register is used to identify the features of the EJTAG implementation. Some of the reset values are set by inputs to the CM2. The register is selected when the Instruction register is loaded with the IMPCODE instruction.

**Figure 17.4 Implementation Register Format**



**Table 17.15 Implementation Register Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
EJTAGver	31:29	Indicates EJTAG Version 5.0.	R	5
reserved	28:14	reserved	R	0
Type	13:10	Type of Entity associated with this TAP. 2: TAP is attached to a Trace-Master. TypeInfo field is not used.	R	2
TypeInfo	10:1	Identifier Information. Unused because this TAP is connected to a Trace-Master as indicated by the Type field.	R	0
reserved	0	reserved	R	0


### EJTAG Control Register

This 32-bit register controls the various operations of the TAP modules. This register is selected by shifting in the CONTROL instruction. Bits in the EJTAG Control register can be set/cleared by shifting in data; status is read by shifting out the contents of this register. This EJTAG Control register can only be accessed by the TAP interface.

The EJTAG Control register is not updated in the *Update-DR* state unless the Reset occurred (Rocc) bit 31, is either 0 or written to 0.

The value used for reset indicated in the table below takes effect on CM2 resets, but not on TAP controller resets by e.g. *TRST\_N*. *TCK* clock is not required when the CM2 reset occurs, but the bits are still updated to the reset value when the *TCK* is applied. The first 5 *TCK* clocks after CM2 resets may result in reset of the bits, due to synchronization between clock domains.

**Figure 17.5 EJTAG Control Register Format**



**Table 17.16 EJTAG Control Register Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Rocc	31	Reset Occurred The bit indicates if a CM reset has occurred: 0: No reset occurred since bit last cleared. 1: Reset occurred since bit last cleared. The Rocc bit will keep the 1 value as long as reset is applied. This bit must be cleared by the probe, to acknowledge that the incident was detected. The EJTAG Control register is not updated in the <i>Update-DR</i> state unless Rocc is 0, or written to 0. This is in order to ensure proper handling of processor access.	R/W	1
Res	30:23	reserved	R	0
Doze	22	Tied to 0.	R	0
Halt	21	Halt state The Halt bit indicates if the internal system bus clock is running or stopped. The value is sampled in the Capture-DR state of the TAP controller: 0: Internal CM clock is running 1: Internal CM clock is stopped	R	0
Res	20:0	reserved	R	0

### 17.3.3.5 CM2 Trace Control Block (TCB) Registers

The TCB registers used to control its operation are listed in [Table 17.17](#) and [Table 17.18](#). These registers, except for *TCBDATA*, are accessed via the EJTAG TAP interface as well as by the interAptiv core via memory-mapped accesses to the Global Debug Control Block in the CM GCRs. *TCBDATA* can only be accessed via the EJTAG TAP interface. Note that none of the TCB registers are implemented if PDTrace is not configured at build time.

**Table 17.17 TCB EJTAG Registers**

EJTAG Register	Memory-Mapped Address*	Name	Description
0x11	0x0008	<i>TCBCONTROLB</i>	Control register in the TCB that is mainly used to specify what to do with the trace information. The <i>REG</i> [25:21] field in this register specifies the number of the TCB internal register accessed by the <i>TCBDATA</i> register. A list of all the registers that can be accessed by the <i>TCBDATA</i> register is shown in <a href="#">Table 17.18</a> . See <a href="#">Section “TCBCONTROLB Register”</a> .
0x15	0x0010	<i>TCBCONTROLD</i>	Control register in the TCB used to control tracing from the Coherence Manager <a href="#">Section “TCBCONTROLD Register”</a>
0x16	0x0020	<i>TCBCONTROLE</i>	Control Register in the TCB used to control tracing for the performance counter tracing feature. See <a href="#">Section “TCBCONTROLE Register”</a> .

**Table 17.18 Registers Selected by TCBCONTROLB<sub>REG</sub>**

<i>TCBCONTROLB<sub>REG</sub></i> Field	Memory Mapped Address*	Name	Reference	Notes
0	0x0028	TCBCONFIG	Section “TCBCONFIG Register (Reg 0)”	
4	0x0200/0x0208**	TCBTW	Section “TCBTW Register (Reg 4)”	These registers have no function if on-chip memory does not exist.
5	0x0108	TCBRDP	Section “TCBRDP Register (Reg 5)”	
6	0x0110	TCBWRP	Section “TCBWRP Register (Reg 6)”	
7	0x0118	TCBSTP	Section “TCBSTP Register (Reg 7)”	
17-29		reserved		
30	0x0040	TCBSYS	Section “TCBSYS Register (Reg 30)”	
31		TCBBYPASS		

\* Memory-Mapped Address relative to the Global Debug Block in the CM GCRs.


\*\* Memory-Mapped Access for TCBTW is split into two 32-bit registers: TCBTW\_LO (address 0x0200) accesses TCBTW[31:0]. TCBTW\_HI (address 0x0208) accesses TCBTW[63:32]

***TCBCONTROLB Register***

The TCB includes a second control register, *TCBCONTROLB* (EJTAG Register 0x11). This register generally controls what to do with the trace information received. This register is also mapped to offset 0x0008 in the Global Debug Block of the CM GCRs.

The format of the *TCBCONTROLB* register is shown below, and the fields are described in [Table 17.19](#).

**Figure 17.6 TCBCONTROLB Register Format**



**Table 17.19 TCBCONTROLB Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
WE	31	Write Enable. Only when set to 1 will the other bits be written in <i>TCBCONTROLB</i> . This bit will always read 0.	R	0
Reserved	30:28	Reserved. Must be written as zero; returns zero on read.	R	0
TWSrc-Width	27:26	Used to indicate the number of bits used in the source field of the Trace Word. The value for the CM is always 0b10 indicating a four bit source field width.	R	10
REG	25:21	Register select: This field select the registers accessible through the <i>TCBDATA</i> register. Legal values are shown in <a href="#">Table 17.18</a> . Note: Although this field can be written via memory-mapped GCR or EJTAG accesses, the <i>TCBDATA</i> register is only accessible via EJTAG access.	R/W	0
WR	20	Write Registers: When set, the register selected by REG field is read and written when <i>TCBDATA</i> is accessed. Otherwise the selected register is only read. Note: Although this field can be written via memory-mapped GCR or EJTAG accesses, the <i>TCBDATA</i> register is only accessible via EJTAG access.	R/W	0



**Table 17.19 TCBCONTROLB Register Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State
Name	Bits			
STCE	19	System Trace capture enable. When asserted, the System Trace port of the Funnel is enabled to capture System Trace stream data. When not asserted, System Trace stream data is not captured regardless of <i>SI_TC_Sys_Valid[1:0]</i> input pin state.	R/W	0
TRPAD	18	Trace RAM access disable bit. When set to 1 core reads and writes to the on-chip trace RAM using GCR accesses are inhibited. If TRPAD is set, memory-mapped writes to the GCR_DB_TCBTW_LO and GCR_DB_TCBTW_HI registers have no effect, and memory-mapped reads from GCR_DB_TCBTW_LO and GCR_DB_TCBTW_HI do not access the Trace RAM and 0 is returned. Also, when TRPAD is set, then memory-mapped writes to all CM TCB registers listed in Table 17.18 are inhibited.	R/W	0
Reserved	17	Reserved. Must be written as zero; returns zero on read.	R	0
RM	16	Read on-chip trace memory. When written to 1, the read address-pointer of the on-chip memory in register <i>TCBRDP</i> is set to the value held in <i>TCBSTP</i> . Subsequent access to the <i>TCBTW</i> register (through the <i>TCBDATA</i> register), will automatically increment the read pointer in register <i>TCBRDP</i> after each read. When the write pointer is reached, this bit is automatically reset to 0, and the <i>TCBTW</i> register will read all zeros. Once set to 1, writing 1 again will have no effect. The bit is reset by setting the TR bit or by reading the last Trace word in <i>TCBTW</i> . This bit has no function if on-chip memory is not implemented.	R/W1	0
TR	15	Trace memory reset. Trace memory reset. When written to one, the address pointers for the on-chip trace memory <i>TCBSTP</i> , <i>TCBRDP</i> and <i>TCBWRP</i> are reset to zero. Also the RM and BF bits are reset to 0. This bit is automatically reset back to 0, when the reset specified above is completed.	R/W1	0
BF	14	Buffer Full indicator that the TCB uses to communicate to external software in the situation that the on-chip trace memory is being deployed in the <b>trace-from</b> and <b>trace-to</b> mode. This bit is cleared when writing 1 to the <i>TR</i> bit. This bit has no function if on-chip memory is not implemented.	R	0

**Table 17.19 TCBCONTROLB Register Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State										
Name	Bits													
TM	13:12	<p>Trace Mode. This field determines how the trace memory is filled when using the simple-break control in the PDtrace interface to start or stop trace.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>TM</th> <th>Trace Mode</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Trace-To</td> </tr> <tr> <td>01</td> <td>Trace-From</td> </tr> <tr> <td>10</td> <td>Reserved</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table> <p>In Trace-To mode, the on-chip trace memory is filled, continuously wrapping around and overwriting older Trace Words, as long as there is trace data coming from the core.                      In Trace-From mode, the on-chip trace memory is filled from the point that the core starts tracing until the on-chip trace memory is full.                      In both cases, de-asserting the EN bit in this register will also stop fill to the trace memory.                      If a <i>TCBTRIGx</i> trigger control register is used to start/stop tracing, then this field should be set to Trace-To mode.                      These bits have no function if on-chip memory is not implemented.</p>	TM	Trace Mode	00	Trace-To	01	Trace-From	10	Reserved	11	Reserved	R/W	0
TM	Trace Mode													
00	Trace-To													
01	Trace-From													
10	Reserved													
11	Reserved													
Reserved	11	Reserved. Must be written as zero; returns zero on read.	R	0										
CR	10:8	<p>Off-chip Clock Ratio. Writing this field, sets the ratio of the core clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 17.20</a>.</p> <p><b>Note:</b> As the Probe interface works in double data rate (DDR) mode, a 1:2 ratio indicates one data packet sent per core clock rising edge.                      These bits have no function if off-chip memory is not implemented.</p>	R/W	100 <sub>2</sub>										

**Table 17.19 TCBCONTROLB Register Field Descriptions(continued)**

Fields		Description	Read / Write	Reset State																																																												
Name	Bits																																																															
Cal	7	<p>Calibrate off-chip trace interface.</p> <p>If set to one, the off-chip trace pins will produce the following pattern in consecutive trace clock cycles. If more than 4 data pins exist, the pattern is replicated for each set of 4 pins. The pattern repeats from top to bottom until the Cal bit is de-asserted.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th colspan="4">Calibrations pattern</th> </tr> <tr> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr><td>0</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>1</td><td>1</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>1</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>1</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>1</td></tr> <tr><td>1</td><td>1</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>1</td><td>0</td><td>1</td><td>1</td></tr> <tr><td>0</td><td>1</td><td>1</td><td>1</td></tr> </tbody> </table> <p style="text-align: center; margin-left: 100px;">This pattern is replicated for every 4 bits of TR_DATA pins.</p> <p><b>Note:</b> The clock source of the TCB and PIB must be running. These bits have no function if off-chip memory is not implemented.</p>	Calibrations pattern				3	2	1	0	0	0	0	0	1	1	1	1	0	0	0	0	0	1	0	1	1	0	1	0	1	0	0	0	0	1	0	0	0	0	1	0	0	0	0	1	1	1	1	0	1	1	0	1	1	0	1	1	0	1	1	1	R/W	0
Calibrations pattern																																																																
3	2	1	0																																																													
0	0	0	0																																																													
1	1	1	1																																																													
0	0	0	0																																																													
0	1	0	1																																																													
1	0	1	0																																																													
1	0	0	0																																																													
0	1	0	0																																																													
0	0	1	0																																																													
0	0	0	1																																																													
1	1	1	0																																																													
1	1	0	1																																																													
1	0	1	1																																																													
0	1	1	1																																																													
Reserved	6:2	Reserved. Must be written as zero; returns zero on read.	R	0																																																												
OfC	1	<p>If set to 1, trace is sent to off-chip memory using <i>TR_DATA</i> pins.</p> <p>If set to 0, trace info is sent to on-chip memory.</p> <p>This bit is read only if a single memory option exists (either off-chip or on-chip only).</p>	R/W	Preset																																																												
EN	0	<p>Funnel Trace Enable. When this bit is set, the trace funnels accepts trace information from the CM and/or cores and writes the information to off-chip or on-chip memory.</p> <p>When this bit is cleared, the trace funnel drops all new trace information from the CM and/or cores . The trace information already accepted by the trace funnel is sent to the off-chip or on-chip memory, but new trace information is dropped and not written out.</p>	R/W	0																																																												

The Probe Interface Block (PIB) has been an available component with many previous MIPS cores, including the interAptiv core. The interAptiv core architecture brings two significant changes to the PIB. First, the PIB is now instantiated in `mips_css`. Second, this new version of the PIB, referred to as PIB2, provides additional clock ratios.

The PIB2 provides available TR\_CLK to processor clock ratios of 1:2, 1:4, 1:6, 1:8, 1:10, 1:12, 1:16, and 1:20. The PIB1 supplied by MIPS has only the ratios 1:2, 1:4, 1:6, and 1:8. The PIB1 architecture also has provision for clock multiples, 1:1, 2:1, 4:1, and 8:1, but these are not supported in PIB2.

The PIB2 reports the minimum CR (TC\_CRMin) as 3'b111 and maximum (TC\_CRMax) as 3'b000 as shown in the table below. This is how software identifies a PIB2 as opposed to PIB.

**Table 17.20 Clock Ratio Encoding of the CR field**


TC_ClockRatio	TR_CLK : gclk
3'b000	1:20
3'b001	1:16
3'b010	1:12
3'b011	1:10
3'b100	1:2
3'b101	1:4
3'b110	1:6
3'b111	1:8

### TCBDATA Register

The *TCBDATA* register (0x12) is used to access the registers defined by the *TCBCONTROLB<sub>REG</sub>* field; see [Table 17.18](#). Regardless of which register or data entry is accessed through *TCBDATA*, the register is only written if the *TCBCONTROLB<sub>WR</sub>* bit is set. For read-only registers, *TCBCONTROLB<sub>WR</sub>* is a don't care.

The format of the *TCBDATA* register is shown below, and the field is described in [Table 17.21](#). The width of *TCBDATA* is 64 bits when on-chip trace words (TWs) are accessed (*TCBTW* access).

**Figure 17.7 TCBDATA Register Format**



**Table 17.21 TCBDATA Register Field Descriptions**


Fields		Description	Read/Write	Reset State
Names	Bits			
Data	31:0 63:0	Register fields or data as defined by the <i>TCBCONTROLB<sub>REG</sub></i> field	Only writable if <i>TCBCONTROLB<sub>WR</sub></i> is set	0

### TCBCONTROLD Register

The TCB includes a second control register, *TCBCONTROLD* (EJTAG Register 0x14), whose values are used to control the tracing functions of the Coherence Manager. External software (i.e., debugger) can therefore manipulate the trace output by writing to this register. This register is also mapped to offset 0x0010 in the Global Debug Block of the CM GCRs.

The format of the *TCBCONTROLD* register is shown below, and the fields are described in [Table 17.22](#)

**Figure 17.8 TCBCONTROLD Register Format.**



**Table 17.22 TCBCONTROLD Register Definition**

Fields		Description	Read / Write	Reset State										
Name	Bits													
Reserved	31:30	Reserved for future use. Must be written as 0.	R	0										
P6_Ctl	29:28	Implementation specific finer grained control over tracing Port 6 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
P5_Ctl	27:26	Implementation specific finer grained control over tracing Port 5 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
P4_Ctl	25:24	Implementation specific finer grained control over tracing Port 4 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
P3_Ctl	23:22	Implementation specific finer grained control over tracing Port 3 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
P2_Ctl	21:20	Implementation specific finer grained control over tracing Port 2 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
P1_Ctl	19:18	Implementation specific finer grained control over tracing Port 1 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
P0_Ctl	17:16	Implementation specific finer grained control over tracing Port 0 traffic at the CM. See <a href="#">Table 17.23</a> .	R/W	0										
Reserved	15:12	Reserved for future use. Must be written as 0 and read as 0.	R	0										
TWSrcVal	11:8	The source ID of the CM.	R/W	0										
WB	7	When this bit is set, Coherent Writeback requests are traced. If this bit is not set, all Coherent Writeback requests are suppressed from the CM trace stream.	R/W	0										
ST_En	6	System Trace Enable. Driven to the CM output pin <i>SI_TC_Sys_Enable</i> . External logic can use this output to control generation of the System Trace stream.	R/W	0										
IO	5	Inhibit Overflow on CM FIFO full condition. When set to 1 the CM never drops trace words, but instead will stall the request and/or intervention processing until forward progress can be made. When set to 0 the CM will drop trace words when the trace word FIFO overflows.	R/W	0										
TLev	4:3	This defines the current trace level being used by CM tracing <table border="1" style="margin-left: 20px; border-collapse: collapse; width: 150px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>No Timing Information</td> </tr> <tr> <td>01</td> <td>Include Stall Times, Causes</td> </tr> <tr> <td>10</td> <td>Reserved</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	00	No Timing Information	01	Include Stall Times, Causes	10	Reserved	11	Reserved	R/W	0
Encoding	Meaning													
00	No Timing Information													
01	Include Stall Times, Causes													
10	Reserved													
11	Reserved													

Fields		Description	Read / Write	Reset State
Name	Bits			
AE	2	When set to 1, address tracing is always enabled for the CM. This affects trace output from the serialization unit of the CM. When set to 0, address tracing may be enabled through the implementation specific P[x]_Ctl bits.	R/W	0
Global_CM_En	1	Each CPU core can enable or disable CM tracing using this bit. This bit is not routed through the master core, but is individually controlled by each core. Setting this bit can enable tracing from the CM even if tracing is being controlled through software, if all other enabling functions are true.	R/W	0
CM_EN	0	This is the master trace enable switch to the CM. When zero, tracing from the CM is always disabled. When set to one, tracing is enabled if other enabling functions are true.	R/W	0

**Table 17.23 P<port>\_Ctl Trace Control Bits**

Value	Meaning
0	Tracing Enabled, No Address Tracing, assuming AE = 0
1	Tracing Enabled, Address Tracing Enabled, independent of AE
2	Reserved
3	Tracing Disabled

The *TCBCONTROL.D.AE* bit enables addresses to be supplied when any request is serialized. This is not typically required because addresses issued from processor CPUs can be inferred from the CPU PDTrace stream.

The *TCBCONTROL.B.TLev* bit controls the amount of information to be included the CM trace. Setting *TLev* to 1 may be useful when debugging performance problems.

The *TCBCONTROL.IO* bit determines the action taken by the CM with its internal trace buffers overflow. If the *IO* bit is 0 then trace information is lost when the trace buffer overflows. In this case, the CM temporarily stops producing trace messages, waits until the trace buffer becomes empty, performs a trace resynchronization with the CPUs and then starts producing new trace words.

However, if *TCBCONTROL.IO* bit is 1 then trace information is never lost, but the system performance may be impacted when the trace buffer becomes full and the additional trace words are required. In this case, the CM stalls the processing of requests and/or L1 intervention responses until a trace buffer becomes available.

The *TCBCONTROL.WB* determines if L1 writebacks are traced or not. L1 writebacks are not software visible and do not appear in the CPU PDTrace, so typically writebacks are not traced in the CM (*WB* set to 0).


The value in the *TCBCONTROL.D.TWSrcVal* field appears in all trace words produced by the CM, thus tagging the trace word as coming from the CM. A unique value must be programmed in this field and *TCBCONTROL.B.TWSrcVal* for all cores.

The five *P<port>\_Ctl* fields in *TCBCONTROLD* give the ability to control the amount of trace information provided for requests received on the specified port. As shown in [Table 17.23](#), requests from a given CM request port can be traced normally, always traced with addresses, or not traced. Typically, the CM request ports connected to CPUs will be traced normally (P0\_Ctl, P1\_Ctl, P3\_Ctl, P4\_Ctl set to 0) because the address is traced by the CPU itself. However, requests from the IOCU are only traced by the CM and therefore should have their addresses traced by the CM (P4\_Ctl should be set to 2).

### TCBCONTROLE Register

The *TCBCONTROLE* register is used to control tracing functions of the Coherence Manager performance counters. The *TCBCONTROLE* register is written by an EJTAG TAP controller instruction, *TCBCONTROLE* (0x16). This register is also mapped to offset 0x0020 in the Global Debug Block of the CM GCRs. The format of the *TCBCONTROLE* register is shown below, and the fields are described in [Table 17.24](#).

**Figure 17.9 TCBCONTROLE Register Format**




**Table 17.24 TCBCONTROLE Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:9	Reserved for future use. Must be written as zero; returns zero on read.	0	0
<i>TrIdle</i>	8	Trace Unit Idle. This bit indicates if the trace hardware is currently idle (not processing any data). This can be useful when switching control of trace from hardware to software and vice versa. The bit is read-only and updated by the trace hardware. TrIdle is set when the all cores and the CM have disabled PDTrace and the trace funnels has written all outstanding trace information to the off-chip or on-chip memory.	R	1
0	7:1	Reserved for future use; Must be written as zero; returns zero on read. (Hint to architect, Reserved for future expansion of performance counter trace events).	0	0
<i>PeC</i>	0	Performance counter tracing is not implemented.	R	0

### TCBCONFIG Register (Reg 0)

The *TCBCONFIG* register holds information about the hardware configuration of the TCB. This register is also mapped to offset 0x0028 in the Global Debug Block of the CM GCRs. The format of the *TCBCONFIG* register is shown below, and the field is described in [Table 17.25](#).

**Figure 17.10 TCBCONFIG Register Format**



**Table 17.25 TCBCONFIG Register Field Descriptions**

Fields		Description	Read / Write	Reset State										
Name	Bits													
CF1	31	This bit is set if a <i>TCBCONFIG1</i> register exists. In this revision, <i>TCBCONFIG1</i> does not exist and this bit always reads zero.	R	0										
0	30:21	Reserved. Must be written as zero; returns zero on read.	R	0										
SZ	20:17	On-chip trace memory size. This field holds the encoded size of the on-chip trace memory. The size in bytes is given by $2^{(SZ+8)}$ , implying that the minimum size is 256 bytes and the largest is 8Mb. This bit is reserved if on-chip memory is not implemented.	R	Preset										
CRMax	16:14	Off-chip Maximum Clock Ratio. This field indicates the maximum ratio of the core clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 17.20</a> . This bit is reserved if off-chip trace option is not implemented.	R	Preset										
CRMin	13:11	Off-chip Minimum Clock Ratio. This field indicates the minimum ratio of the core clock to the off-chip trace memory interface clock. The clock-ratio encoding is shown in <a href="#">Table 17.20</a> . This bit is reserved if off-chip trace option is not implemented.	R	Preset										
PW	10:9	Probe Width: Number of bits available on the off-chip trace interface <i>TR_DATA</i> pins. The number of <i>TR_DATA</i> pins is encoded, as shown in the table. <table border="1" data-bbox="500 1037 1049 1228"> <thead> <tr> <th>PW</th> <th>Number of bits used on <i>TR_DATA</i></th> </tr> </thead> <tbody> <tr> <td>00</td> <td>4 bits</td> </tr> <tr> <td>01</td> <td>8 bits</td> </tr> <tr> <td>10</td> <td>16 bits</td> </tr> <tr> <td>11</td> <td>reserved</td> </tr> </tbody> </table> This field is preset based on input signals to the TCB and the actual capability of the TCB. This bit is reserved if off-chip trace option is not implemented.	PW	Number of bits used on <i>TR_DATA</i>	00	4 bits	01	8 bits	10	16 bits	11	reserved	R	Preset
PW	Number of bits used on <i>TR_DATA</i>													
00	4 bits													
01	8 bits													
10	16 bits													
11	reserved													
PiN	8:6	Pipe number. Indicates the number of execution pipelines.	R	0										
OnT	5	When set, this bit indicates that on-chip trace memory is present. This bit is preset based on the selected option when the TCB is implemented.	R	Preset										
OfT	4	When set, this bit indicates that off-chip trace interface is present. This bit is preset based on the selected option when the TCB is implemented, and on the existence of a PIB module ( <i>TC_PibPresent</i> asserted).	R	Preset										
REV	3:0	Revision of TCB.	R	0										




### TCBTW Register (Reg 4)

The *TCBTW* register is used to read Trace Words from the on-chip trace memory. The TW read is the one pointed to by the *TCBRDP* register. A side effect of reading the *TCBTW* register is that the *TCBRDP* register increments to the next TW in the on-chip trace memory. If *TCBRDP* is at the max size of the on-chip trace memory, the increment wraps back to address zero.

This register is also mapped to offset 0x0200 (lower 32 bits) and 0x0208 (upper 32 bits) in the Global Debug Block of the CM GCRs.

The format of the *TCBTW* register is shown below, and the field is described in [Table 17.26](#).

**Figure 17.11 TCBTW Register Format**



**Table 17.26 TCBTW Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
Data	63:0	Trace Word	R/W	0


### TCBRDP Register (Reg 5)

The *TCBRDP* register is the address pointer to on-chip trace memory. It points to the TW read when reading the *TCBTW* register. When writing the *TCBCONTROLB<sub>RM</sub>* bit to 1, this pointer is reset to the current value of *TCBSTP*.

This register is also mapped to offset 0x0108 in the Global Debug Block of the CM GCRs.

The format of the *TCBRDP* register is shown below, and the field is described in [Table 17.27](#). The value of n depends on the size of the on-chip trace memory. As the address points to a 64-bit TW, lower three bits are always zero.

**Figure 17.12 TCBRDP Register Format**



**Table 17.27 TCBRDP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
Data	31:(n+1)	Reserved. Must be written zero, reads back zero.	0	0
Address	n:0	Byte address of on-chip trace memory word.	R/W	0


### TCBWRP Register (Reg 6)

The *TCBWRP* register is the address pointer to on-chip trace memory. It points to the location where the next new TW for on-chip trace will be written.

This register is also mapped to offset 0x0110 in the Global Debug Block of the CM GCRs.

The format of the *TCBWRP* register is shown below, and the fields are described in [Table 17.28](#). The value of *n* depends on the size of the on-chip trace memory. As the address points to a 64-bit TW, the lower three bits are always zero.

**Figure 17.13 TCBWRP Register Format**



**Table 17.28 TCBWRP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
Data	31:(n+1)	Reserved. Must be written zero, reads back zero.	0	0
Address	n:0	Byte address of on-chip trace memory word.	R/W	0


***TCBSTP Register (Reg 7)***

The *TCBSTP* register is the start pointer register. This pointer is used to determine when all entries in the trace buffer have been filled (when *TCBWRP* has the same value as *TCBSTP*). This pointer is reset to zero when the *TCBCONTROLB<sub>TR</sub>* bit is written to 1. If a continuous trace to on-chip memory wraps around the on-chip memory, *TSBSTP* will have the same value as *TCBWRP*.

This register is also mapped to offset 0x0118 in the Global Debug Block of the CM GCRs.

The format of the *TCBSTP* register is shown below, and the fields are described in [Table 17.29](#). The value of *n* depends on the size of the on-chip trace memory. As the address points to a 64-bit TW, lower three bits are always zero.

**Figure 17.14 TCBSTP Register Format**




**Table 17.29 TCBSTP Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Names	Bits			
Data	31:(n+1)	Reserved. Must be written zero, reads back zero.	0	0
Address	n:0	Byte address of on-chip trace memory word.	R/W	0

***TCBSYS Register (Reg 30)***

The *TCBSYS* register contents are driven to the *SI\_TC\_Sys\_UserCtl[31:0]* output signals. This register is also mapped to offset 0x0040 in the Global Debug Block of the CM GCRs. Thus, any change to this register will be reflected in these output signals. The format of the *TCBSYS* register is shown below, and the fields are described in [Table 17.30](#).

**Figure 17.15 TCBSYS Register Format**



**Table 17.30 TCBSYS Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
STA	31	System Trace Available. Set to 1 if the System Trace Interface is present. Otherwise it is set to 0.	R	present: 1 not present: 0
UsrCtl	30:0	User-defined Control.	R/W	0

**Register Reset State**

Reset state for all register fields is entered when either of the following occur:

1. TAP controller enters/is in Test-Logic-Reset state.
2. *EJ\_TRST\_N* input is asserted low.

**17.3.4 MIPS Trace Capability**

There are several build-time options for trace support within the interAptiv Multiprocessing System (MPS):

1. No trace logic included.
2. Trace logic to support an on-chip trace memory (embedded within the MPS).
3. Trace logic to support an off-chip trace probe (with off-chip trace memory).
4. Combination of options 2 and 3.

**17.3.5 Memory-Mapped Access to PDtrace™ Control and On-Chip Trace RAM**

PDtrace can be controlled entirely through software and the on-chip trace memory can be accessed directly by software using load and store instructions.

**17.3.6 On-Chip Trace Buffer Usage**

In order to direct trace data to the on-chip buffer instead of the off-chip interface, the OfC bit in the TCBControlB register of the trace master must be cleared. Once this is done, the trace funnel will combine trace data it receives from the CM and CPUs and write it to the on-chip memory. Tracing can be enabled or disabled on a per CM/CPU basis by setting or clearing the EN bits in the corresponding TCBControlB registers.

To initialize the on-chip trace buffer, the TR bit of the TCBControlB register of the trace master is set by software. This will initialize TCBRDP, TCBWRP and TCBSTP pointers to zero. These pointers do not have to explicitly written by software for initialization, the reset function that is caused by setting the TR bit is sufficient.

When it is desired to read out the Trace Words from the on-chip buffer, software first sets the RM bit within TCBControlB. This will load the TCBRDP register with the value held in the TCBSTP register. The TraceWord pointed by

TCBRDP can be then read out through the TCBTW register. The read will automatically update the TCBRDP value to point to the next newer entry. A subsequent read from TCBTW register will thus read out the next newer Trace-Word. Software does not have to explicitly update the TCBRDP register.

If the TM field of TCBControlB register is set to Trace-From mode, the trace-buffer contents stop being updated when the trace-buffer is full (when TCBWRP points to the same entry as TCBSTP). This event is denoted by the BF bit of TCBControlB register. The BF bit can be polled by software to decide when to read out the trace buffer contents.

For production testing, such as stuck-at testing of memory cells within the trace buffer, the TCBRDP and TCBWRP registers can be explicitly written by software to write and read specific entries within the trace buffer. As previously stated, for normal usage these pointer registers do not have to be explicitly written by software.

## o32 Binary Interface and General Purpose Registers

The MIPS32 ABI provides the application binary interface to the interAptiv Multiprocessing System. The 32 general purpose registers (GPR) are used by the compiler and provide a set of temporary storage locations used for function calls and other software functions. Some registers have unique functions, while others can be used for any purpose.

This chapter contains the following sections:

- [Section 18.1 “o32 ABI Parameters”](#)
- [Section 18.2 “Compiler Goals”](#)
- [Section 18.3 “Register Naming Conventions and Usage”](#)
- [Section 18.4 “Conventional Naming and General Purpose Register Usage”](#)
- [Section 18.5 “Floating Point Register Conventions”](#)
- [Section 18.6 “Virtual Memory Layout Overview”](#)
- [Section 18.7 “Mapping Data Types into Memory”](#)
- [Section 18.8 “Calling Conventions”](#)

### 18.1 o32 ABI Parameters

This section contains information on the o32 binary interface. The o32 ABI contains calling and linkage conventions for the MIPS32 architecture.

[Table 18.1](#) shows the o32 ABI parameters.

**Table 18.1 MIPS o32 ABI Parameters**

Parameter	Value
Registers saved and restored as	32-bit
Argument structure	4-byte slots 8-byte alignment at least 4 slots long
Argument registers	4 integer, 2 FP
Arguments in FP registers?	Leading FP arguments only (doesn't need correct function prototypes)
Return values	Only scalars are ever returned in registers; v1 is used only for <b>long long</b> data.
<b>long long</b> type	Implemented with register pairs and hardware type library calls
gp register in PIC code	Not preserved over calls

## 18.2 Compiler Goals

Compilation systems should adhere to following standards (“ABIs”) in order to achieve the following goals:

- **Inter-calling** : a binary program built with one compiler should be able to call a subroutine defined in another (so long as address resolution problems are solved). The standards relevant to this are called the “calling conventions”; they describe how subroutines pass parameters, return values, and cooperate to share the register set and stack resources. Refer to [Section 18.8 “Calling Conventions”](#).
- **Interlinkable**: object files built with one compiler can be linked successfully with those produced by another. The standard relevant to this is the object code definition, in particular the definition of symbols and relocation mechanisms.
- **Runnable**: a binary produced with a compliant tool kit can be successfully executed on a compliant OS (Linux in particular). For more information, refer to [Section 18.6.3 “Linux Applications”](#).
- **Debuggable**: more conventions and standards are required before a program built with a tool kit can be successfully debugged.
- **Profilable**: where available, code profilers have their own requirements - related to but not identical to those of debuggers.

## 18.3 Register Naming Conventions and Usage

Each interAptiv core offers 32 general purpose registers available for program use. They are numbered \$0 to \$31.

Two, of these registers behave as follows::

- \$0 always returns zero, no matter what is stored in it.
- \$31 is always used by the normal subroutine-calling instruction (`jal`) for the return address. Note that the call-by register version (`jalr`) can use any register for the return address, though use of anything except \$31 is not recommended .

In all other respects the general purpose registers are identical and can be used in any instruction (it is legal to use \$0 as the destination of instructions, though the result data will not be saved).

The return address of a `jal` is the next instruction, but one in sequence:

```
...
jal printf
move $4, $6
xxx # return here after call
```

The floating point math coprocessor (called FPA for floating point accelerator), if included, adds 32 floating point registers with their own conventions: see [Section 18.5 “Floating Point Register Conventions”](#).

## 18.4 Conventional Naming and General Purpose Register Usage

Although the hardware makes few rules about the use of registers, their practical use is governed by a number of conventions. As part of those conventions, the registers are referred to by conventional names — typically defined in a header file<sup>2</sup> and implemented by using the C preprocessor on assembler files.

MIPS hardware ignores these conventions, but all the benefits of software by register name and associated common name. These common name functions are described in the following subsections.

**Table 18.2 CP0 Registers Grouped by Number**

Register Name	Common Name	Description
\$0	zero	Always has the value 0. Any writes to this register are ignored.
\$1	AT	Assembler temporary.
\$2	v0	Function result register. Functions return integer results in v0, and 64-bit integer results in v0 and v1 when using 32-bit registers. In cases where floating-point hardware is not present, or when compiler options enable floating-point emulation, functions return single precision floating-point results in v0 and double precision floating-point results in v0 and v1 when using 32-bit registers. v0 and v1 can be temporary registers. Not preserved across function calls.
\$3	v1	
\$4	a0	Function argument registers that hold the first four words of integer type arguments. Functions use these registers to hold floating-point arguments. When floating-point hardware is not present, or compiler options enable floating-point emulation, functions use a0 to hold the first single precision floating-point argument and a1 to hold the second single precision floating-point argument. Functions use a0-a1 for the first double precision floating-point argument, and a2-a3 to hold the second double precision floating-point argument. Not preserved across function calls.
\$5	a1	
\$6	a2	
\$7	a3	
\$8	t0	Temporary registers. Can be used for any purpose. Not preserved across function calls.
\$9	t1	
\$10	t2	
\$11	t3	
\$12	t4	
\$13	t5	
\$14	t6	
\$15	t7	
\$24	t8	
\$25	t9	
\$16	s0	Saved registers to use freely. Preserved across function calls. These registers must be saved before use by the called function.
\$17	s1	
\$18	s2	
\$19	s3	
\$20	s4	
\$21	s5	
\$22	s6	
\$23	s7	
\$30	s8	
\$26	k0	
\$27	k1	
\$28	gp	Global pointer. May be used as save register for called functions.
\$29	sp	Stack pointer.
\$31	ra	Return address register, saved by the calling function. Available for use after saving.
\$f0		Function return register used to return float and double values from function calls.

**Table 18.2 CP0 Registers Grouped by Number (continued)**

Register Name	Common Name	Description
(\$f12, \$f13) and (\$f14 and \$f15)		Two pairs of registers used to pass float and double valued parameters to functions. Pairs of registers are parenthesized because they have to pass double values. To pass float values, only \$f12 and \$f14 are used.

### 18.4.1 Common Name — AT (GPR \$1)

GPR \$1 is reserved for the synthetic instructions generated by the assembler. When using this register, such as when saving or restoring registers in an exception handler, there is an assembler directive to stop the assembler from using it without permission (but then some of the assembler’s macro instructions won’t be available.) The assembler directive’s existence is the reason why this name is traditionally used in upper case.

### 18.4.2 Common Name — v0, v1 (GPR \$2, \$3)

GPR \$2 and \$3 are used when returning non-floating-point values from a subroutine. If the value returned is too large to fit in two registers, the compiler will allocate a memory buffer whose address will be passed as an invisible first argument.

While a function is running v0-v1 can be freely used as temporaries. Integer values are returned in these registers. Structure or array types (even if small enough to fit in the two registers) are always returned through a data area defined by the caller and whose address is an invisible first argument.

### 18.4.3 Common Name — a0 - a3 (GPR \$4 - \$7)

GPR \$4 - \$7 are used to pass the first four non-FP parameters to a subroutine. Argument registers which are unused or whose value is no longer needed can be freely used as temporaries. For more information, refer to [Section 18.8 “Calling Conventions”](#)

### 18.4.4 Common Name — t0 - t9 (GPR \$8 - \$15, GPR \$24 - \$25)

By convention, subroutines may use these registers without doing anything to preserve their previous contents. This makes them a good choice for “temporaries” when evaluating expressions. However, the compiler/programmer must remember that values stored in them may be destroyed by a subroutine call.

### 18.4.5 Common Name — s0 - s8 (GPR \$16 - \$23, GPR \$30)

By convention, subroutines must guarantee that the values of these registers on exit are the same as they were on entry. This can be accomplished by either not using them, or by saving them on the stack and restoring before exit.

This makes them eminently suitable for use as “register variables” or for storing any value which must be preserved over a subroutine call.

### 18.4.6 Common Name — k0, k1 (GPR \$26, \$27)

GPR \$26 and GPR \$27 are reserved for use by the trap/interrupt handlers in an operating system, which uses them and does not restore their original value. Hence they are of little use to other code. These registers are not used at all by application code.



## 18.4.7 Common Name — gp (GPR \$28)

GPR \$28 has two quite different roles. These registers are used in position-independent (PIC) code and non-PIC code. Each case is described below.

### 18.4.7.1 Position Independent Code (PIC)

In PIC code, typically used only for application and library code in a large OS, it is used in the double-indirection used to reach variables and functions whose location is not known until the program and its libraries are loaded.

In position-independent code, *gp* acts as a pointer to the GOT (“global offset table”), as described in [Section 18.6.4 “PIC Code and the Global Offset Table”](#). The GOT pointer is loaded by code in the prologue of every function which makes a reference through the GOT. A function call may overwrite the value in *gp* and the compiler must ensure it’s reloaded after any such call.

### 18.4.7.2 Non-Position Independent Code (Non-PIC)

In non-PIC code, typically used for all non-Linux embedded applications, this register is sometimes used to provide efficient access to C static/extern data.

If used, the *gp* register is initialized to point to a load-time-determined location of the static data. This means that loads and stores to data lying within 32 KBytes on either side of the *gp* value can be performed in a single instruction using *gp* as the base register. Note that the pointer in *gp* is a constant. No application code ever writes to the register once it has been initialised.

Without the global pointer, loading data from a static memory area takes two instructions: one to load the most significant bits of the 32-bit constant address computed by the compiler and loader, and one to do the data load.

To use *gp*, a compiler must know at compile time that a datum will end up linked within a 64 KByte range of memory locations. In practice this cannot be known, only estimated. The usual practice is to put “small” global data items (8 bytes and less in size) in the *gp* area, and to get the linker to complain if it still gets too big. The compiler `-Gnn` flag can be used to adjust the threshold of what is considered “small”.

## 18.4.8 Common Name — sp (GPR #29)

The GPR #29 register functions as a stack pointer. It takes explicit instructions to raise and lower the stack pointer, so MIPS code usually adjusts the stack only on subroutine entry and exit; and it is the responsibility of the subroutine being called to do this.

*sp* is normally adjusted, on entry, to the lowest point that the stack will need to reach at any point in the subroutine. Now the compiler can access stack variables by a constant offset from *sp*.

## 18.4.9 Common Name — fp (GPR #30)

A subroutine will use a “frame pointer” to keep track of the stack if it wants to do things which involve extending the stack by an amount which is determined at run-time. Some languages, including C++, may do this implicitly.

If the stack bottom can’t be computed at compile time, the stack variables cannot be accessed from *sp*, so *fp* is initialized by the function prologue to a constant position relative to the function’s stack frame. Efficient use of register conventions means that this behavior is local to the function, and doesn’t affect either the calling code, or any nested function calls.

### 18.4.10 Common Name — ra (GPR #31)

The GPR #30 register is used to store the return address. On entry to any subroutine, ra holds the address to which control should be returned. Therefore, a subroutine typically ends with the instruction: jr ra. Subroutines which themselves call subroutines must first save ra, usually on the stack.

## 18.5 Floating Point Register Conventions

In addition to the 31 general purpose registers, the MIPS32 architecture also includes a corresponding set of standard floating point registers.

### 18.5.1 MIPS Floating Point Registers

The interAptiv core contains 32 floating point registers, whose assembler names are \$f0 - \$f31. The architecture also supports the 64-bit IEEE double-precision format.

The o32 ABI generates code that is compatible with not only previous generation 32-bit MIPS I and MIPS II CPUs, but also the MIPS III CPU's such as the interAptiv Multiprocessing System. All of these architectures do floating point arithmetic using the 16 even-numbered registers \$f0 - \$f30.

The odd-numbered registers are referred to in move and load/store instructions; but the assembler provides synthetic “macro” instructions for move and load/store double, so the odd-numbered registers will normally not be required when writing o32 code.

### 18.5.2 Floating Point Register Software Use and Calling Conventions

Like the general-purpose registers, the MIPS calling conventions contains restrictions about register use that have nothing to do with the hardware. These restrictions include which FP registers are used for passing arguments, which ones' values are expected to be preserved over function calls, and so on. The division of functions is much as for the integer registers, less the special cases.

[Table 18.3](#) shows the floating point register usage conventions. Note that the o32 ABI assumes that the CPU either has a MIPS II or earlier FP unit or that the CPU has the SR[FR] compatibility bit cleared to zero; in either case only 16 registers are usable for arithmetic, so there are no odd-numbered registers in the table.

**Table 18.3 Floating Point Register Usage Conventions**

Parameter	Value
Function return values	\$f0, \$f2
Argument registers	\$f12, \$f14
Saved over function call (suitable for register variables)	Events \$f20 - \$f30
Temporaries (not saved over function call, or “caller-saved”)	Evens \$f4 - \$f10, \$f16, \$f18

## 18.6 Virtual Memory Layout Overview

Although this section is not technically part of the ABI definition, it is useful to examine the interaction between the various pieces of code and data which might make up an ABI-compliant application. The following subsections describe some aspects of how the virtual address may be used.

### 18.6.1 Memory Regions and Object Code Naming Conventions

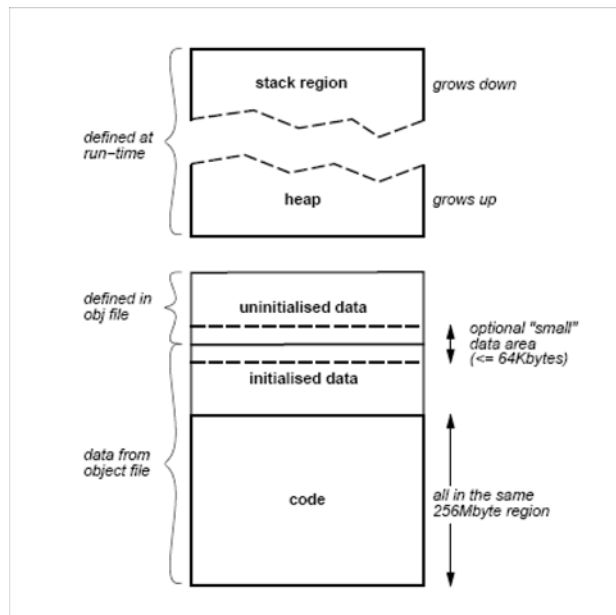
Some of the memory region and object code naming conventions are as follows:

- **Module:** A compilation unit (the assembler is seen as just another compiler...), and also used for an object file generated from one compilation unit.
- **Program:** all the addressable data and code associated with an application. Strictly speaking, that associated with an instance of an application; in Linux there may be many copies of the shell running, and they're distinct programs in this sense. For Linux applications, this excludes the kernel and other parts of the memory map which are not accessible to the application.
- **Link unit:** a part of a program which has been bound together so that its components are at fixed offsets from each other.
- **Segment:** a part of a program which is contiguous in the memory image of the running program, and which is distinguished for link/build purposes. By convention, segment names begin with a dot, and are called things like `.text` and `.bss`. When several modules are being combined into a single link unit during the build process, sections of the same name in different modules are brought together and various sections concatenated to make a segment.
- **`_main`:** C programmers often think execution begins with `main()`; but in reality there's always a more primitive, machine-dependent startup routine supplied by the build environment. This does things like initializing the `sp` register to mark the stack region, and zero-ing the memory region which contains the "uninitialised" C variables. If you write C++, this will also arrange for initialization routines to be run.

### 18.6.2 Simple Standalone Application

[Figure 18.1](#) describes the memory map of a statically linked application.

**Figure 18.1 Memory Map of a Statically Linked Application**



The memory segments used in [Figure 18.1](#) are described below.

- **Code** (including initialiaed and uninitialiaed data): This space forms a single link unit; and their relative positions are fixed when the software is built. In fact, their absolute locations in program memory are also typically fixed at build-time. This code is position-dependent.
- **Stack**: The stack is assigned by the start-up program in accordance with OS and tool chain conventions. It grows down, so is typically placed at the top of the program’s memory space.
- **Heap**: The heap defines the data space allocated by the program through C setbrk() or (slightly higher level) malloc() calls. The heap usually starts at the lowest suitably aligned location available after allowing for the linked code and data.
- **Small data area**: For this space, it takes two MIPS instructions to load from or store to a C location declared at module level or as static. When the “small” data area is used, the gp register is set to point to the middle of it by \_\_main(). This allows load and store to variables in that area to be accessed with a single instruction.

Note that for some programs, the entire data set will not fit within the 64 KByte address range limit imposed by the MIPS load/store instruction’s 16-bit offset. Therefore, during compilation and build, only data items below a certain size are considered as candidates for this area. That’s why it’s called “small” data. The small data area, if provided, overlaps both the initialized and uninitialized segments (and is implemented as a pair of sub-segments).

- **Common segment names:**

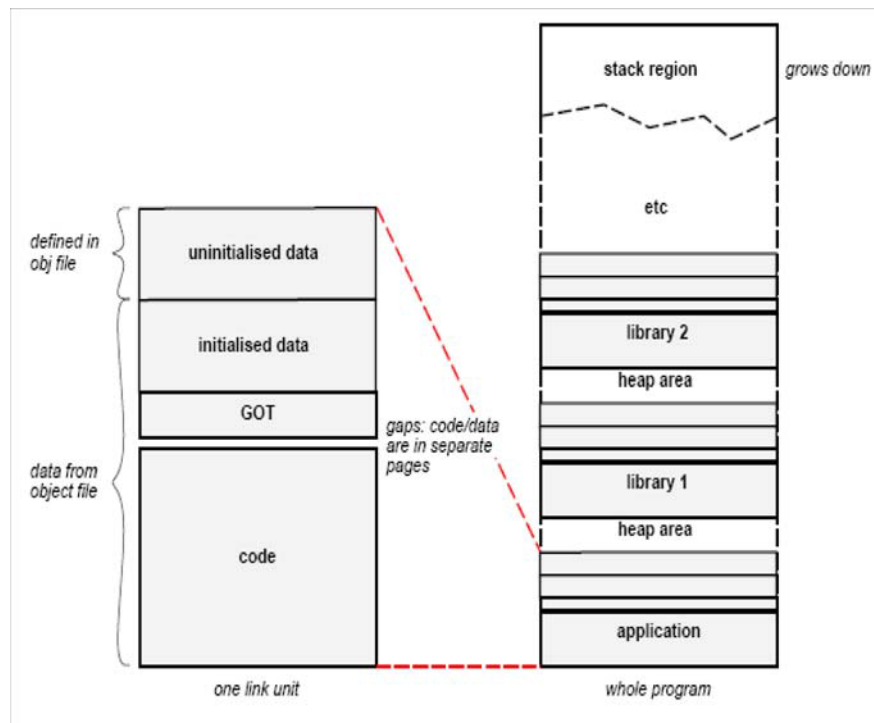
```
.text all the code
.data initialised data possibly excluding...
.sdata initialised data for the “small” area
.bss uninitialised data possibly excluding
.sbss uninitialised data for the “small” area.
```

Note that in some embedded operating systems, all the software runs in the same address space and does not take advantage of the MIPS memory management facilities. Such systems have complicated memory maps which are deeply OS dependent.

### 18.6.3 Linux Applications

Many Linux applications are built without their library functions. In this case, the library routines are linked in as the program is loaded into memory. The library routine may have been updated since the application was built, and it should still work. The result is a much more complicated memory map with a number (perhaps quite a large number) of separately linked pieces. This complexity is illustrated in [Figure 18.2](#).

**Figure 18.2 Memory Map for a Typical Linux Application**



In the memory map shown in [Figure 18.2](#) above, all the link units except the base application are shared libraries of some kind, either built-in shared libraries or dynamically loaded by explicit programming. They are loaded into program memory working upward on a first-come first-served basis.

While the base application runs at program addresses which were known at build time, the libraries must be able to run at arbitrary memory addresses. The requirement that library modules should link in just anywhere and still work (“position-independent code” or PIC) forces considerable changes to the way code is generated.

In the MIPS architecture, the preferred subroutine call instruction is `jal`, and that instruction is not PC-relative; it encodes (most of) the absolute virtual address of the subroutine entry point. Moreover, Linux standards require that the shared libraries should also be able to share extern data.

## 18.6.4 PIC Code and the Global Offset Table

An application's binary code is built before it is known where data or subroutines in other link units will reside in the program's memory map. Both the absolute and relative position of each link unit depends on what versions of what libraries get loaded in what order. It is not possible for the run-time loader to fix up these addresses in the code itself, because the code itself must be shared between different instances of the application program (and each instance may have a different library layout). As a result, an application's or library's binary will contain the address information needed to reference functions and data in a different link unit.


### 18.6.4.1 Global Offset Table (GOT)

Instead, the compiler generates code which makes every function call and every reference to static/extern data indirect, via a table of pointers. The table of pointers are the Global Offset Table or "GOT", which reside in a data segment. Separate copies are kept for each instance of the application, so it can be and is fixed up by the loader.

Figure 18.3 shows an example of a global offset table. The GOT contains an entry for each function or data item that is accessed by any code in the link unit (the loader finds each item by its name, so there is an entry for each symbol). The table offset for a particular symbol is known at build time, and is a constant in the binary code.

For MIPS code, the *gp* register is maintained as a pointer to the GOT of the link unit.

Figure 18.3 Example of a Global Offset Table



The the *gp* register is set to point to the GOT by code included as part of the prologue of each function (at least, each function which makes any use of the GOT.) This is suboptimal, since for intra-link-unit calls it will already hold the right value. However, the compiler cannot (in general) distinguish intra- and inter-link-unit calls.

In the o32 ABI, the calling code must be aware that a function call might overwrite the value in the *gp* register, and the caller must preserve or recalculate the value after the call if required.

#### 18.6.4.2 Loading a PIC Application and its Libraries

The program loader is the Linux application, which runs when any binary is loaded which uses shared libraries. The program loader maps the application code and data and any libraries it needs into the program's address space. The build system leaves a list of required library names in the application's object file, and the program loader finds the library files via a series of search path mechanisms. Conventions about file names (if followed correctly) make sure the program finds the "right" library.

The program loader maintains symbol tables for the data items and subroutine entry points which are exported by the application and each library, so it can tie up references between separate link units.

#### 18.6.4.3 Loading and Binding of Libraries

While it is not necessary to read in all the code of the libraries required by an application (the ordinary virtual memory paging system takes care of that), the process of binding in a link unit, fixing up its GOT and getting it ready for use, is relatively time-consuming. This penalty is paid even for libraries which provide facilities which the application rarely uses. That can slow the application startup.

To optimize the initialization process, Linux defers loading and fixing up libraries until they are first used. By the nature of the PIC code the unresolved references are all in the GOT. Where the first reference to the new library is a function call this is relatively straightforward; the GOT entry for an unresolved subroutine reference is set to point to a function in the run-time loader which then loads the library, patches the GOT so that future calls will go direct, and calls the library function.

There are other more subtle issues. For example, when the same symbol is provided by two different libraries, this can make loading problematic. As a result, the build system is charged with identifying which libraries are safe to load, and to identify them in the application binary. The loader can then load unsafe libraries at startup.

#### 18.6.4.4 Dynamic (explicit) Loading of Libraries - `dlopen()`

It is also possible to get software to pick its own shared library and then build an explicit software-visible table of calls to it. This mechanism fits naturally onto the object/class concepts of C++, and libraries loaded like this are referred to as "dynamic shared objects".

It is not necessary to build a Linux shared library in a special way to make it fit for `dlopen()`, any library will do. At the lowest level, a call to `dlopen()` to grab the library and `dlsym()` calls to obtain pointers to named data items or functions in the dynamic shared object. But because dynamic libraries are just shared libraries, some unexpected "bonus" semantics may be observed.

Firstly, the explicitly-loaded library will gain access to any public symbols in the application (or its pre-loaded libraries). Perhaps more unexpectedly, a straightforward extern function pointer reference in the application can bind to a symbol from a library which wasn't mentioned at all at build time, but only brought in with `dlopen()`.

#### 18.6.4.5 PIC and GOT Constraints

There are some PIC and GOT constraints that are worth mentioning.

1. **What to do when your GOT overflows:** On MIPS, GOT pointer loads are usually compiled to a single load relative to the `gp` register; but this can only span a table 64 Kbytes in size (16K pointer entries). Large applications and libraries can use more symbols than that.

There are two approaches. One is to just let the GOT grow above 64 Kbytes, and require the compiler to generate code which can load/store arbitrary entries in it. This generally uses the `gcc -PIC` option - it's trouble-free and

portable but generates truly awful code.

Some compilers support an option that generates one GOT to each module in a link unit (`gcc -multigot`). Done properly, this is no trouble, but the dynamic loader has to know about it.

2. **Managing the overheads of PIC code:** Nothing in the ABI obliges the compiler to go through the GOT when accessing data or calling subroutines which are in the same link unit; neither is it strictly necessary for a function to reset the `gp` register on an intra-link-unit call.

However, there are several reasons why this hasn't been done:

- Even within the link unit only relative addresses are known; the MIPS architecture lacks efficient PCrelative call and load instructions.
- The compiler doesn't know which references are in the link unit. While it's possible to get the linker to do some instruction re-writing to simplify intra-link-unit calls and references, it's bad practice.
- The PIC calling convention for MIPS requires that on entry to a function the `t9` register holds the address of the function's entry point. Since calls made through the GOT mean the address may be in some register, this seems acceptable — but this requirement is burdensome to any possible future intra-link-unit (or even intra-module) call mechanism.

## 18.7 Mapping Data Types into Memory

Memory is taken as an array of unsigned 8-bit quantities, whose index is the virtual address. For all MIPS CPU's, this corresponds to a C definition unsigned `char` [].

MIPS uses 2s-complement representation for signed integers - so in any data size “-1” is represented by binary all-ones. The overwhelming advantage of 2s-complement numbers is that the basic arithmetic operations (add, subtract, multiply, divide) have the same implementation for signed and unsigned data types.

### 18.7.1 Sizes of C Data Types

Table 18.4 lists fundamental C data types and how they're implemented for MIPS architecture CPUs.

**Table 18.4 Data Types and Memory Representations**

C Type	MIPS ASM Name	Size in Bytes
char	byte	1
short	half	2
int	word	4
long long	dword	8
float	word	4
double	dword	8


Note that all of the primitive data types shown in Table 18.4 can only be directly handled by standard MIPS instructions if they are naturally aligned: that is, a 2-byte datum starts at an address which is even (zero modulo 2), a 4-byte datum starts at an address which is zero modulo 4, and an 8-byte datum starts at an address which is zero modulo 8.



## 18.7.2 Basic Type Memory Layout and Endianness

Figure 18.4 shows how each basic type is laid out in byte-addressed memory. Note that the arrangement is different for bigendian and little-endian software.

Figure 18.4 C Data Types in Memory — Big Endian and Little Endian



In Figure 18.4, I've given in to the temptation to reverse the bit-numbering within each byte for the big-endian layouts. For memory addressing purposes this is meaningless; bytes are indivisible 8-bit objects. However, reversing the bit numbers as above makes the bitwise depiction of the fields of floating point numbers easier to absorb (and prettier). Each of these data types is naturally aligned, as described above. “Endianness” can be a troubling subject. If you are uneasy about it, read it up in [SMR].

## 18.7.3 Memory Layout and Alignment of Structure and Array Types

Complex types are built by concatenating simple types, but inserting unused (“padding”) bytes between items so as to respect the alignment rules.

The following shows the byte offsets of data items in a struct mixed:

```
struct mixed {
    char c; /* byte 0 */
    /* bytes 1-14 are “padding” */
    double d; /* bytes 8-15 */
    short s; /* bytes 16-17 */
};
```

Note that the byte offsets of the fields of constructed data types (other than those using C bitfields, see [Section 18.7.3 “Memory Layout and Alignment of Structure and Array Types”](#) below) are unaffected by endianness.

Constructed data types are aligned in memory to the largest alignment boundary required by the data type defined inside them. So a struct mixed will start on an 8-byte boundary; which means that if an array is built using these structures, padding will be required between each array element. C compilers provide for this by “tail padding” the structure to make it usable for an array, so `sizeof(struct mixed) == 24` and the structure should really be annotated:

```
struct mixed {
    char c; /* byte 0 */
    /* bytes 1-14 are “padding” */
    double d; /* bytes 8-15 */
    short s; /* bytes 16-17 */
    /* bytes 18-23 are “tail padding” */
};
```

## 18.7.4 Bit Fields in Structures

C allows the user to define structures which pack several short “bit field” members into one or more locations of a standard integer type. This is a useful feature for emulation, hardware interfacing, and perhaps for defining dense data structures, but is fairly incomplete. Bitfield definitions are nominally CPU-dependent but also genuinely endianness-dependent.

One can, for example, define a data structure which permits access to the various fields of a MIPS single-precision floating point number as shown below:

```
#if BYTE_ORDER == BIG_ENDIAN

struct ifloat {
    unsigned int sign:1;
    unsigned int bexp:8;
    unsigned int mant:23;
};

#else /* little-endian */

struct ifloat {
    unsigned int mant:23;
    unsigned int bexp:8;
    unsigned int sign:1;
};

#endif
```

In this case the three fields are packed into one 32-bit int storage unit. These two cases are the same in that, for both Endian formats, the bitfields are allocated with the first-defined field occupying the lowest byte-addressed part of the int.

These two examples differ as follows: For big-endian, the high-order bits are occupied first. For little-endian, the low-order bits are occupied first.

If one tries to implement bitfields in a less endianness-dependent way, then in the following example struct fourbytes would have a different memory layout from struct fouroctets as shown:

```

struct fourbytes {
    signed char a; signed char b; signed char c; signed char d;
}

struct fouroctets {
    int a:8; int b:8; int c:8; int d:8;
}

```

A field can only be packed inside one storage unit of its defined type. When trying to define a structure for a MIPS double-precision floating point number, the mantissa field contains part of two 32-bit int storage units and can't be defined in one attempt. The best that can be done is something similar to the following:

```

struct ieee754dp_konst {
    unsigned sign:1;
    unsigned bexp:11;
    unsigned manthi:20; /* cannot get 52 bits into... */
    unsigned mantlo:32; /* .. a regular C bitfield */
};

```

### 18.7.5 Alignment Rules

The full alignment rules for bit-fields are as follows:

- A bit-field must reside entirely in a storage unit that is appropriate for its declared type. Thus a bit-field never crosses its unit boundary.
- Bit-fields can share a storage unit with other struct/union members, including members that are not bit-fields (to pack together, the adjacent structure member must be of a smaller integer type).
- Structures generally inherit their own alignment requirement from the alignment requirement of their most demanding type. Named bit-fields will cause the structure to be aligned (at least) as well as the type requires.

Unnamed fields - regardless of their defined type - only force the storage unit or overall structure alignment to that of the smallest integer type which can accommodate that many bits.

- It may be advantageous to force subsequent structure members to occupy a new storage unit. In some compilers this can be done with an unnamed zero-width field. Zero-width fields are otherwise illegal.

## 18.8 Calling Conventions

The calling convention describes how arguments are passed to functions, and how values are returned. It's also a convenient place to describe the stack frame structure which builds up to represent the current function nest.

ANSI C permits pretty much any value - structures and arrays as well as scalars - to be passed as arguments or returned by a function.

### 18.8.1 Stack Maintenance and Alignment

When the stack is adjusted by functions to make space for local variables, register saves and argument passing, it is always adjusted by a multiple of 8 bytes so that the stack base is aligned to the greatest extent required by any variable.

## 18.8.2 Registers and the Argument Structure

For efficiency, it is recommended to pass arguments in registers and avoid data loads/stores. But C permits pretty much any non-array data type - no matter how large or complex - to be passed as an argument. It is not “obvious” how such arguments should be passed. To make sure the corner cases are handled correctly, the set of arguments passed to a function is mapped as it would be to a memory-based argument structure, and then as much of that structure as will fit is pasted into the available registers. For any arguments left over after all available argument registers have been used up, a copy of that part of the argument structure is placed onto the stack.

The rules are as follows:

1. Each argument is aligned to the start of a new argument slot within the argument structure. These slots are 4 bytes in size, chosen to match the size of the general-purpose registers. If the next slot doesn't have the correct alignment for a value (for example, a double on o32 requires 8-byte alignment), it is skipped to find a slot which is correctly aligned. Skipped slots remain unused. Large arguments may spill over into more than one slot.
2. Integer values are first converted to the type of the argument (if there is a function prototype) using standard C rules. Where there is no function prototype, the rules are that integer and floating point values are coerced to signed int and double respectively.
3. Integers smaller than int are expanded to int by zero- or sign-extending them in accordance with C rules.
4. Non-integer arguments smaller than a register-sized slot are aligned to the lowest addressed part of the slot.
5. Float arguments are 8-byte aligned and occupy two slots (even though there's nothing useful in the second four bytes).
6. The argument registers are identified with a particular slot in the argument structure. If for alignment or other reasons a slot cannot be used, then the corresponding register won't be used to pass an argument.
7. The caller will always build an argument data structure, even though it may remain unused in whole or part. Moreover, the data structure is always a minimum of 16 bytes (four register-sized slots) in size.
8. The first 4 x register-sized (ie 4 byte) slots of the structure are mapped to registers a0-3.
9. o32 does not assume the existence of function prototypes. For reasons to do with the implementation of functions with variable numbers of arguments, it is difficult to ensure that the caller and the called function always agree when to use a floating point rather than a general-purpose register for an argument.

o32's rule is that up to two leading floating point arguments will be passed in FP registers, but if the first argument is not an FP a second FP argument will not be put in an FP register. In functions like printf() the first argument is a pointer, so floating point values will be passed in integer registers or on the stack.

## 18.8.3 Returning Values from a Function

In the o32 ABI, a simple scalar value is returned in a register; v0 for integers, and fv0 for floating point values. A second integer register is defined for returning larger values, and is used when returning a long long value in o32.

For all other structures or larger values which are not accommodated in the registers, the caller must provide a pointer to a memory buffer (usually on the stack, but that's not mandatory). The caller prepends a pointer to the memory buffer as an implicit first argument, followed by its explicit arguments. The called function should copy the return value to the supplied address.

## 18.8.4 Calling Conventions Extended for Linux (“MIPS ABI”) PIC Code.

In PIC code functions are not called directly; instead the compiler/assembler generate code which loads the function address from the GOT table (see [Section 18.6.4 “PIC Code and the Global Offset Table”](#) above). The disassembled code looks something like this:

```
/* (caller) */
    lw t9, <function symbol offset in GOT>(gp)
    # nop
    jalr t9
    # nop
    ...

/* function */
    /* _gp_disp is magic symbol for offset between start of
    function and gp pointer into GOT */
    li gp, _gp_disp
    addu gp, gp, t9
    ...
```

It’s mandatory that the *t9* register should be used to compute the function address; the function itself depends on it to recalculate the GOT base register *gp*. *\_gp\_disp* is calculated so as to place *gp* 32 KBytes on from the start of the GOT, to maximise the amount of the table which is in reach of a MIPS load instruction (which has a  $\pm 32\text{K}$  offset range).



# Programming Concepts

This chapter describes some programming concepts that can be followed when programming in either User mode or Kernel mode.

## 19.1 Ordering and Synchronization

### 19.1.1 Consistency Model

The interAptiv core uses weak memory ordering, that is, cacheable loads and stores on a processor can be executed out of program order (for example, for hit-under-miss). Software must include SYNC instructions to enforce ordering in the cases where it is required.

### 19.1.2 LL/SC

The Load Linked (LL) and Store Conditional (SC) instructions provide a mechanism that ensures atomic access to a memory location.

An LL instruction reads a memory location and sets an internal (per-TC) state bit called the LL bit. The address read by the LL instruction is stored in the *LLAddr* register. The LL bit can be cleared because of actions on the processor, such as an ERET instruction or a write to TCstatus. If the LL bit is cleared before the SC completes, the SC fails and does not update memory.

On the interAptiv core, the value in *LLAddr* is also checked on interventions. If another CPU requests write access to the cache line, the LL bit will be cleared. LL instructions always request the line in a Shared state, so that an LL itself does not clear the LL bit in another CPU. If the line is installed as Shared, when the SC is executed, it must make a CohUpgrade request to obtain write access to the line. If multiple cores are trying to access the same location, there can be a race, and the first Upgrade request to be serialized in the CM will win. This will cause the other SCs to fail, and the other cores must retry the sequence from the LL.

These actions allow the memory location to appear as though it were atomically updated—the SC will not write the location unless the update will appear atomic.

### 19.1.3 Memory Barriers

The SYNC instruction is used to enforce the ordering of loads and stores. Because the core processes instructions in order and generates memory requests in order, these SYNCs can complete with much less delay than the traditional heavyweight SYNC. All of the lightweight stypes (0x4, 0x10-0x13) are treated identically by the CPU as follows:

1. The LSU forces any pending evictions to complete their cache reads and send the writes to the BIU.
2. The BIU flushes the write-back buffer.
3. The BIU indicates that it is complete and allows the LSU to resume processing instructions.

4. No external SYNC request is generated.

Additionally, the CPU supports two implementation-specific `stype` values as well as the standard `stype 0x0`. These are used to explicitly set the ‘level’, which controls how far into the system SYNCs are propagated: If Coherence is enabled:

- `stype 0x2`: A coherent SYNC is sent to the CM. The CM responds when all older coherent requests have completed their interventions.
- `stype 0x3`: A Coherent SYNC transaction is sent to the CM. If `SI_CM_SyncTxEn` is 0 or `CM_SYNC_TX_DISABLE` is 1 then the CM responds when all previous coherent requests have completed their interventions and all previous requests have been accepted on the L2/Memory interface. If `SI_CM_SyncTxEn` is 1 and `CM_SYNC_TX_DISABLE` is 0, then the CM waits until all previous coherent requests have been completed and before issuing a Legacy SYNC transaction to L2/memory (behind all previous coherent and non-coherent requests from this CPU) to enforce ordering throughout the system. In this case, the CM responds when it has received a response from L2/Memory.
- `stype 0x0`: The level that normal SYNCs use can be controlled by the SYNCCTL bit in the *Global CM2 Control* register located at offset address 0x0010. Refer to the *Global CM2 Control* register in the CM2 Registers chapter for more information.

All other `stypes` are reserved and currently default to type 0x0.

If Coherence is disabled:

- `stype 0x0, 0x2, 0x3`: A legacy SYNC transaction is issued to the CM. If `SI_CM_SyncTxEn` is 0 or `CM_SYNC_TX_DISABLE` is 1 then CM responds when all previous requests have been accepted on the L2/Memory interface. If `SI_CM_SyncTxEn` is 1 and `CM_SYNC_TX_DISABLE` is 0 then the CM issues a Legacy SYNC transaction to L2/Memory (behind all previous non-coherent requests from this CPU) to enforce ordering throughout the system. In this case, the CM responds when it has received a response from L2/Memory.

All other `stypes` are reserved and currently default to type 0x0.

**Table 19.1 Supported SYNC stypes**

Coherence Enabled?	stype	Behavior
Yes	0x0	Can be configured as level0 or level1, as defined below.
	0x2	Level0 - SYNC transaction is sent to the Coherence Manager and waits for all previous coherent transactions to finish their intervention stage.
	0x3	Level1 - After completing level0 steps, memory accesses are also completed. Depending on the setting of the <i>SyncTxEn</i> and <i>CM_SYNC_TX_DISABLE</i> bits, an external SYNC transaction may also be generated to flush external devices.
No	0x0, 0x2, 0x3	Core issues Legacy SYNC
-	0x4, 0x10-0x13	Lightweight SYNC - handled entirely within the CPU, completes evictions and flushes WBB.
	All others	Reserved. Default to type 0x0.



## 19.1.4 CACHE and SYNCI Instructions

Coordinating software maintenance of the caches across multiple cores can be rather challenging and involve a lot of overhead. To simplify the task of maintaining cache Coherence via software, the interAptiv core includes hardware support for the globalization of a number of cache maintenance operations. When a cache operation is globalized, it becomes a coherent request and is sent through the Coherence Manager to be performed on all of the cores. The decision of whether or not to globalize an operation is based on whether or not the target address for the operation is coherent. The operations that are globalized are Hit-type L1 CACHE instructions and SYNCI instructions.

Several special cases deserve additional consideration and are discussed below.

### 19.1.4.1 Cache Line Locking

Locking lines into a cache is somewhat counter to the idea of coherence. If a line is locked into a particular cache, it is expected that any processes utilizing that data will be locked to that processor and coherence is not needed. Based on this usage model, locking coherent lines into the cache is not recommended. If it is done, the cores will use the following rules:

- SYNCI instructions are user-mode instructions. Because locking is a kernel-mode feature (it requires the CACHE instruction), SYNCI is not allowed to unlock cache lines. This applies to both local and globalized SYNCI instructions.
- Locking overrides coherence. Intervention requests from other cores and I/O devices that match on a locked line will be treated as misses.
- Self-intervention requests for globalized CACHE instructions will be allowed to affect a locked line. This is done primarily for handling lock and unlock requests for kseg0 addresses when kseg0 is being treated coherently.

### 19.1.4.2 Index Type and Optimized Routines

Index-type CACHE instructions are not globalized. Because they refer to a specific cache location, it does not make sense to apply them to other caches, particularly if the cache configurations are not homogeneous.

One case where software may attempt to use index-type CACHE instructions is an optimization used when flushing large blocks of memory. If the region to be flushed is larger than the size of the cache, flushing the entire cache could be faster than walking through the region and flushing each cache line individually (though the flushing of unrelated cache lines may mitigate the benefit of this optimization). Because indexed operations are not globalized, this sequence only flushes the local cache. If flushing of the remote caches is also required, the code sequence must also run on the remote cores. It is probably better to disable this software optimization and make use of the efficiency of the globalized hit-type CACHE instructions.

### 19.1.4.3 Completion

Globalizing a cache operation changes its timing, compared to a local operation. The external request must be made, serialized in the Coherence Manager, and then sent to the cores on the intervention port. This is not a blocking action, and subsequent instructions on the requesting CPU will continue to execute. In order to guarantee that the operation has been completed, a SYNC instruction must be executed prior to any instruction that requires the updated state. This can be a single SYNC after a series of cache operations. This SYNC should also be used on non-coherent cores in the Cluster to ensure maximum compatibility moving forward.

### 19.1.4.4 L2 CACHE Instructions

It is important to note that L2 CACHE instructions only impact the L2 cache and do not affect the L1 data or instruction caches.

## 19.1.5 PREF Instructions

Prefetch instructions are also impacted by coherence. The different types of PREF react differently, as described below.

- *Normal*: Load/store(*\_\**) type hint values will cause the appropriate type of request to be issued when a coherent CCA is used—a store hint will request Exclusive ownership, and a load hint will request either Shared or Exclusive, depending on the CCA. However, a store-type PREF that hits on a Shared line will not make an Upgrade request.
- *Writeback\_invalidate* (also called *nudge*): This operation behaves the same for both coherent and non-coherent CCAs and in both cases will only force a writeback (if needed) from the local cache.
- *Prepare for Store*: This operation is intended to avoid the memory read when software is going to be writing an entire cache line. When a coherent address is used, an Invalidate request is generated to clear the lines of any other data caches in the system and acquire Exclusive ownership on the local processor.  
Note: This operation changes the state of memory, and the data values are unpredictable until the series of stores has completed. If other software (running on other processors, TCs, or even the same TC) accesses the line before the series of stores has completed, this unpredictable intermediate state can be observed.

## 19.2 User Mode Programming

This section contains the following programming concepts relative to user mode programming:

- [Section 19.2.1, "User Mode Accessible CP0 Registers"](#)
- [Section 19.2.2, "Prefetching Data Using the pref and prefix Instructions"](#): how it works.
- [Section 19.2.3, "Using "SYNCl" When Writing Instructions"](#): writing instructions without needing to use privileged cache management instructions.
- [Section 19.2.4, "Integer Multiply and Divide"](#): multiply, multiply/accumulate and divide timings.
- [Section 19.2.5, "Tuning Software for the Pipeline"](#): for determined programmers, and for compiler writers. It includes information about the timing of the DSP ASE instructions.
- [Section 19.2.6, "Branch Misprediction Delays"](#): the floating-point unit often runs at half speed, and some of its interactions (particularly about potential exceptions) are complicated. This section offers some guidance about the timing issues you'll encounter.
- [Section 19.2.7, "Load Delayed by \(Unrelated\) Recent Store"](#)
- [Section 19.2.8, "Minimum Load-miss Penalty"](#)
- [Section 19.2.9, "Data Dependency Delays"](#)
- [Section 19.2.10, "Advice on Tuning Instruction Sequences \(particularly DSP\)"](#)
- [Section 19.2.11, "Multiply/Divide Unit and Timings"](#)

### 19.2.1 User Mode Accessible CP0 Registers

In the interAptiv architecture, privileged code executed in kernel mode can access any CP0 register. Conversely, unprivileged user mode code does not have access to any CP0 register. However, there are instances where unprivileged user mode programs may need information from some of the CPU registers, normally to share information which is worth making accessible to programs without the overhead of a system call.


This can be accomplished by setting selected bits in the CP0 Hardware Enable (HWREna) register located at Register 7, Select 0. See [Section 19.2.1.1 "Programming the HWREna Register"](#) for more information.

#### 19.2.1.1 Programming the HWREna Register

To facilitate non-privileged user mode accesses to selected CP0 registers, the interAptiv core allows selected information from the CP0 register set to be accessed via the `rdhwr` instruction. The operating system can control access to each register individually, through a bitmask in the CP0 register *HWREna* - (set bit 0 to enable register 0 etc) register. *HWREna* is cleared to all-zeroes on reset, so software has to explicitly enable user access.

[Figure 19.1](#) shows the bit assignments for the *HWREna* register. Note that the entire register is cleared to zero on reset, so that no hardware register is accessible without positive OS clearance.

**Figure 19.1 Fields in the HWREna Register**



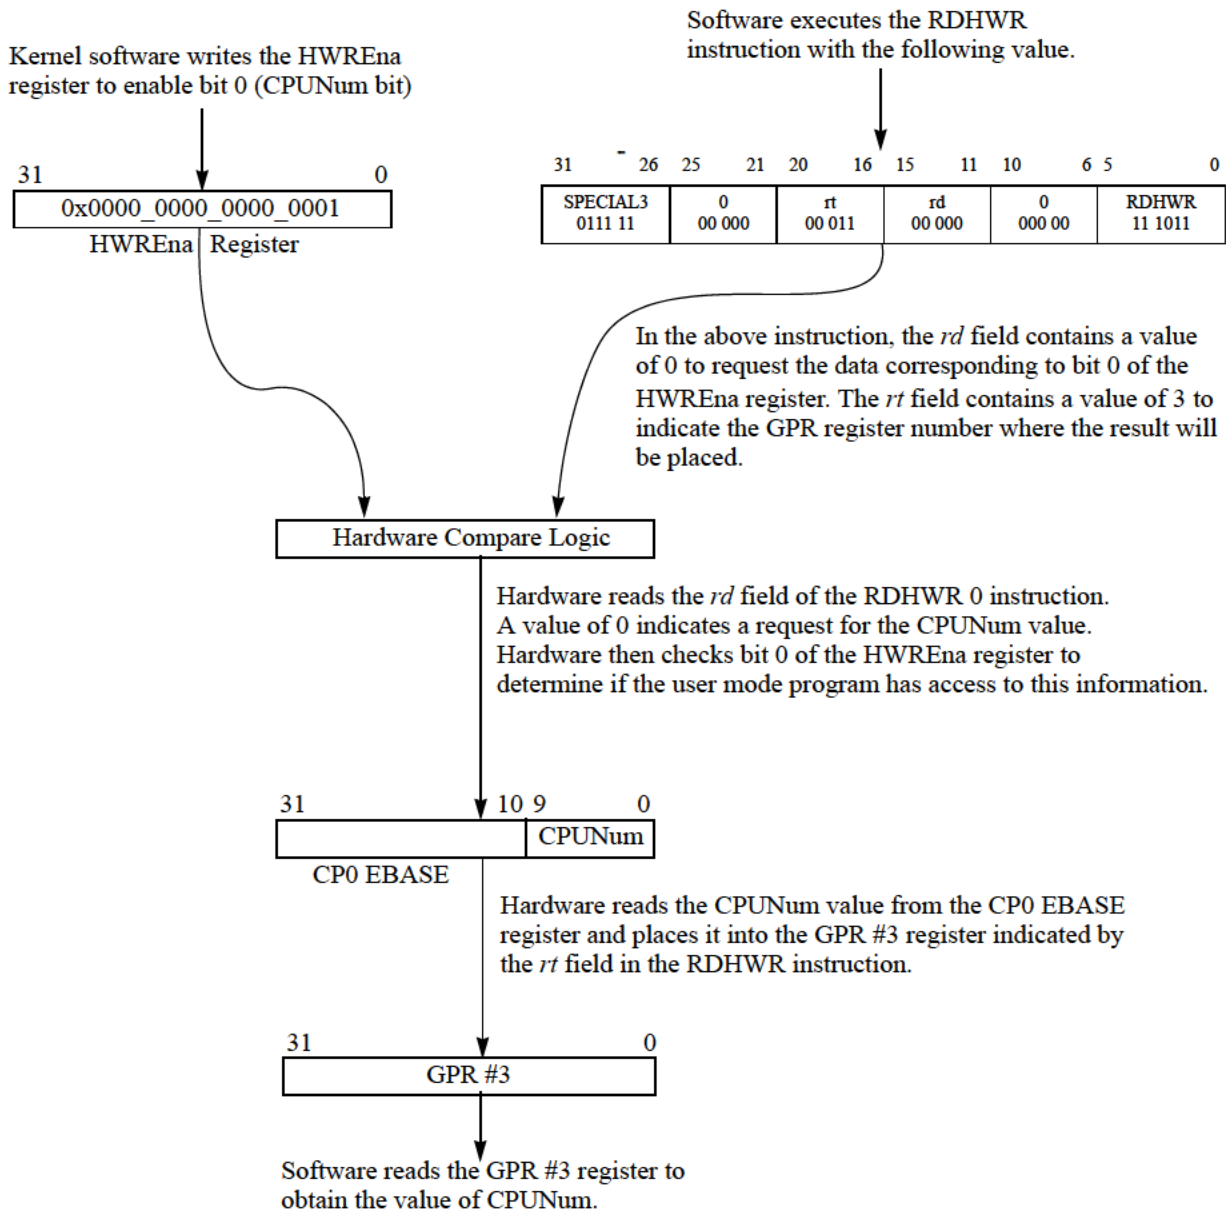
The *HWREna* register contains five bit-fields that allow access to the following information:

- *HWREna<sub>CPUNum</sub>* (bit 0): Software sets this bit to enable hardware fetch the number of the CPU on which the program is currently running. Upon execution of a **rdhwr 0** instruction, hardware fetches this information directly from the CP0 *EBase<sub>CPUNum</sub>* field and places it into the GPR register designated by the *rt* field of the **rdhwr 0** instruction. The *CPUNum* field in the *EBase* register is initially set by hardware based on the setting of external pins at reset.
- *HWREna<sub>SYNCL\_Step</sub>* (bit 1): When this bit is set, execution of a **rdhwr 1** instruction returns the effective size of an L1 cache line. This information is important to user programs because they can now do things to the caches using the **synci** instruction to make recently written instructions visible for execution. The information returned indicates the “step size” — the address increment between successive **synci** instructions required to cover all the instructions in a range. In the interAptiv core, the line size is always 32-bytes.
- *HWREna<sub>CC</sub>* (bit 2): When this bit is set, the interAptiv hardware allows user mode read-only access to the CP0 *Count* register, for high-resolution counting. Execution of a **rdhwr 2** instruction returns the current value in the *Count* register and places it into the GPR register designated by the *rt* field of the **rdhwr 2** instruction.
- *HWREna<sub>CCRes</sub>* (bit 3): When this bit is set, which tells you how fast *Count* counts. It’s a divider from the pipeline clock. If the **rdhwr 3** instruction reads a value of “2”, then the *Count* register increments every 2 cycles, at half the pipeline clock rate. In the interAptiv core, the *CCRes* value is always 2 to indicate that the *CC* register increments every second core cycle.
- *HWREna<sub>UL</sub>* (bit 29): When this bit is set, hardware allows user mode read-only access to the CP0 *UserLocal* register. The execution of a **rdhwr 29** instruction provides a core identifier to user mode programs.

### 19.2.1.2 Programming Example

The following example shows the flow of information through the interAptiv core when user code accesses the CPU number from the CP0 register set.

**Figure 19.2 Obtaining the CPU Number in User Mode**



### 19.2.2 Prefetching Data Using the `pref` and `prefx` Instructions

MIPS32 CPUs are being increasingly used for computations which feature loops accessing large arrays, and the runtime is often dominated by cache misses.

These are excellent candidates for using the `pref/prefx` instructions, which gets data into the cache without affecting the CPU's other state. In a well-optimized loop with prefetch, data for the next iteration can be fetched into the cache in parallel with computation for the last iteration. The `pref/prefx` instructions are logically a no-op<sup>1</sup> and should have *no software-visible effect* other than to make things go faster.

The `pref/prefx` instructions come with various possible “hints” which allow the program to express its best guess about the likely fate of the cache line. The “load” and “store” variants of the hints for this instruction perform differently. Specifically, when dealing with coherent addresses, the “store” types will request Exclusive ownership so that an eventual store does not need to make an upgrade request before modifying the line.

The interAptiv core acts on hints as summarized in [Table 19.2](#).

**Table 19.2 Hints for pref and prefx Instructions**

<i>Hint</i>		<i>Action Taken by the Core</i>	<i>Usage</i>
<i>Number</i>	<i>Name</i>		
0	load	Read the cache line into the D-cache if not present.	When you expect to read the data soon. Use “store” hint if you also expect to modify it.
1	store		
4	load_streamed	Fetch data, but always use cache way zero - so a large sequence of “streamed” prefetches will only ever use a quarter of the cache.	For data you expect to process sequentially, and can afford to discard from the cache once processed
5	store_streamed		
6	load_retained	Fetch data, but never use cache way zero. That means if you do a mixture of “streamed” and “retained” operations, they will not displace each other from the cache.	For data you expect to use more than once, and which may be subject to competition from “streamed” data.
7	store_retained		
25	writeback_invalidate/nudge	If the line is in the cache, invalidate it (writing it back first if it was dirty). Otherwise do nothing. However (with the interAptiv core only): if this line is in a region marked for “uncached accelerated write” behavior, then write-back this line.	When you know you’ve finished with the data, and want to make sure it loses in any future competition for cache resources.
30	PrepareForStore	If the line is not in the cache, create a cache line - but instead of reading it from memory, fill it with zeroes and mark it as “dirty”. If the line is already in the cache do nothing - <i>this operation cannot be relied upon to zero the line.</i>	When you know you will overwrite the whole line, so reading the old data from memory is unnecessary. A recycled line is zero-filled only because its former contents could have belonged to a sensitive application - allowing them to be visible to the new owner would be a security breach.

### 19.2.3 Using “SYNCI” When Writing Instructions

The `synci` instruction (introduced with Revision 2 of the MIPS32 architecture specification) ensures that instructions written by a program (necessarily through the D-cache) get written back from the D-cache and corresponding I-cache locations invalidated, so that any future execution at the address will reliably execute the new instructions. `synci` takes an address argument, and it takes effect on a whole enclosing cache-line sized piece of memory. User-level programs can discover the cache line size because it’s available in a “hardware registers” accessed by the `rdhwr` instruction, as described in [Section 19.2.1, “User Mode Accessible CP0 Registers”](#) above.

Since `synci` is modifying the program’s own instruction stream, it’s inherently an “instruction hazard”. Therefore, once the last instruction has been written and the last `synci` has been issues, programmer’s should use a `jr.hb` or equivalent to call the new instructions.

The following code example shows how the `synci` can be used.

1. Note that the `pref` instruction with the “PrepareForStore” hint can zero out some data which wasn’t previously zero.

```

/*
* This routine makes changes to the instruction stream effective to the hardware. It should be called after the instruc-
* tion stream is written. On return, the new instructions are effective.
*
* Inputs:
* a0 = Start address of new instruction stream
* a1 = Size, in bytes, of new instruction stream
*/

        beq    a1, zero, 20f                /* If size==0, */
        nop                                /* branch around */
        addu   a1, a0, a1                  /* Calculate end address + 1 */
        rdhwr  v0, HW_SYNCI_Step          /* Get step size for SYNCI from new */
                                                /* Release 2 instruction */

        beq    v0, zero, 20f                /* If no caches require synchronization, */
        nop                                /* branch around */
10:     synci  0(a0)                        /* Synchronize all caches around address */
        addu   a0, a0, v0                  /* Add step size in delay slot */
        sltu   v1, a0, a1                  /* Compare current with end address */
        bne   v1, zero, 10b                /* Branch if more to do */
        nop                                /* branch around */
        sync                                       /* Clear memory hazards */
20:     jr    hb    ra                        /* Return, clearing instruction hazards */
        nop

```

## 19.2.4 Integer Multiply and Divide

As is traditional with MIPS CPUs, the integer multiplier is a semi-detached unit with its own pipeline. All MIPS32 CPUs implement:

- **mult/multu**: a 32×32 multiply of two GPRs (signed and unsigned versions) with a 64-bit result delivered in the multiply unit's pseudo-registers *hi* and *lo* (readable only using the special instructions **mflhi** and **mfllo**, which are interlocked and stall until the result is available).
- **madd, maddu, msub, msubu**: multiply/accumulate instructions collecting their result in *hi/lo*.
- **mul/mulu**: simple 3-operand multiply as a single instruction.
- **div/divu**: divide - the quotient goes into *lo* and the remainder into *hi*.

No multiply/divide operation ever produces an exception - even divide-by-zero is silent - so compilers typically insert explicit check code where it's required.

The interAptiv core multiplier is high performance and pipelined. *Multiply/accumulate* instructions can run at a rate of 1 per clock, but a 32×32 3-operand multiply takes six clocks longer than a simple ALU operation. Divides use a bit-per-clock algorithm, which is short-cut for smaller dividends. Multiply/divide instructions are generally slow enough that it is difficult to arrange programs so that their results will be ready when needed.

## 19.2.5 Tuning Software for the Pipeline

This section is addressed to low-level programmers who are tuning software by hand and to those working on efficient compilers or code translators.

### 19.2.5.1 Cache Delays

In a typical CPU implementation a cache miss which has to be refilled from DRAM memory will be delayed by a some period of time, perhaps long enough to run 50-200 instructions. In addition, a miss or uncached read may easily be several times slower.

Because these delays are so large, there is little that can be done when a cache miss except wait for it to be resolved. To mitigate cache misses, the interAptiv core supports non-blocking loads. Therefore, if the programmer can provide separation in the code stream between a load instruction producer and its consumer, the memory delay will not begin until the consuming instruction is executed.

Compilers and programmers may find it difficult to move fragments of an algorithm around like this, so the interAptiv core also provides *prefetch* instructions, such as **pref** and **prefx**, which fetch designated data into the D-cache, but do nothing else. Any loop which walks predictably through a large array is a candidate for prefetch instructions, which are conveniently placed within one iteration to prefetch data for the next.

The **pref PrepareForStore** prefetch saves a cache refill read, for cache lines which are intended to be overwritten in their entirety. Read more about prefetch in [Section 19.2.2, "Prefetching Data Using the pref and prefx Instructions"](#) above.

#### *Tuning Data-Intensive Common Functions*

Bulk operations like **bcopy** () and **bzero** () can benefit from CPU-specific tuning. To get excellent performance for in-cache data, it's only necessary to reorganize the software enough to cover the address-to-store and load-to-use delays. To get the loop to achieve the best performance when the cache misses, **pref** instructions can be used.

## 19.2.6 Branch Misprediction Delays

In a pipelined design with multiple stages, branch delays would be lengthy if software waited until the branch was executed before fetching any more instructions. In general, the amount of delay depends on the type of branch. For example, a conditional branch which closes a tight loop will almost always be predicted correctly after the first time around.

However, too many branches in too short a period of time can overwhelm the ability of the instruction fetch logic to keep ahead with its predictions, even if the predictions are almost always right. Three empty cycles occur between the delivery of the branch delay slot instruction and the first instruction(s) from the branch target location. To mitigate the effects of 'branchy' code, the code can be replaced by conditional moves or tight loops "unrolled" to get at least 6-8 instructions between branches. This should provide a significant performance benefit.

The branch-likely instructions deprecated by the MIPS32 architecture document are predicted just like any other branch. The misprediction of branch-likely instructions costs an extra cycle or two, because the branch and the delay slot instruction needs to be re-executed after a mispredict. Branch-likely instructions sometimes improve the performance of small loops on the interAptiv core, but they set problems for the designers of complex CPUs, and may one day disappear from the standard. Good compilers for the MIPS32 architecture should provide an option to avoid these instructions.

## 19.2.7 Load Delayed by (Unrelated) Recent Store

Load instructions are handled within the execution unit with "standard" timing, just so long as they hit in the cache. When a load misses (or turns out to be uncached) then a dependent operation which has already been issued will have to be reissued if the dependent instruction has been dispatched, which can generate additional delay. If the dependent instruction has not been dispatched, it remains in the data queue (DDQ) until the load data becomes available.



Conversely, store instructions are graduated before they are completed. This is because store instructions cannot be written to the cache (or commit a write to real memory) until they graduate and cease to be speculative. This can present a problem in that a programmer may write code which stores a value in memory, then immediately loads the same value. The CPU pipeline detects circumstances where instructions are dependent for register values, but cannot do the same for addresses. As such, the load can get the right data from an incomplete store as a side-effect of checking whether the requested data might be in the FSB (fill/store buffer) attached to the D-cache. In addition, the store data can also be in intermediate stages/queues before being written into the FSB. Any data that matches stores in such intermediate queues will also be bypassed back to the pipeline as if the load hit in the cache.

## 19.2.8 Minimum Load-miss Penalty

The interAptiv core runs at high frequencies, so any load that misses in the L1 D-cache is likely to be substantially delayed, waiting for the data to come back from the L2 cache. If the load misses in the L2 cache, a much greater delay is incurred.

If an instruction that consumes the loaded data issues before it is determined that the load has missed, then that instruction will have to be re-executed by stopping execution and starting again on the consuming instruction. This is likely to occur if the consuming instruction is only a few places behind in the instruction sequence. That means it has to be re-fetched from the I-cache, and this involves a delay of approximately 15 cycles.

## 19.2.9 Data Dependency Delays

The out-of-order pipeline in the interAptiv core allows dependent instructions to be executed as soon as possible, in hardware. So to some extent the out-of-order pipeline makes it unnecessary to manage data delays by moving instructions around in the program sequence.

Compilers might reasonably try to schedule code to create opportunities for dual-issue and so that instructions might be issued at full speed despite dependencies, but should rarely do so if the cost is significant — the hardware is already gaining much of this advantage within its out-of-order window, and compiler scheduling will not be worth many extra instructions or significant code bloat unless it reaches beyond such a window. Loop unrolling will often help, but local scheduling will be unlikely to make a lot of difference.

In the MIPS instruction set, most dependent instructions can run nose-to-tail just one clock apart. Each register has a “standard” place in the pipeline where the producer should deliver its value and another place in the pipeline where the consumer picks it up: where those places are 1 cycle apart, the dependent instructions to run in successive cycles. Producer/consumer delays happen when either the producer is late delivering a result to the register, or the consumer insists on obtaining its operand early. If either of these conditions occurs, the delays can add up.

Most of these delays are hidden by out-of-order execution.

Different register classes are read/written in different “standard” pipeline slots, so it’s important to be clear what class of registers is involved in any of these delays. For non-floating-point user-level code, there are just three:

- General purpose registers (“GPR”).
- The multiply unit’s *hi/lo* pair together with the three additional multiply-unit accumulators defined by the MIPS DSP ASE (“ACC”).

The MIPS architecture encourages implementations to provide integer multiply and divide operations in a separate pipelined unit capable of doing multiply-accumulate operations at a rate of one per clock. No multiply unit operation ever causes an exception, which makes the longer multiply-unit pipeline rather invisible. It shows up in late delivery of GPR values by those few multiply-unit instructions which deliver GPR results.

- The fields of the *DSPControl* register, used for condition codes and exceptional conditions resulting from DSP ASE operations.

## 19.2.10 Advice on Tuning Instruction Sequences (particularly DSP)

DSP algorithm functions are often the subject of intense tuning. There are four basic classes of DSP instructions:

- A group of specially-simple ALU instructions run in one cycle. This includes bitwise logical instructions, `mov` (an alias for `addu` with `$0`), shifts up to 8 positions down or up, test-and-set instructions, and sign-extend instructions.
- Simple DSP ASE operations (no multiply, no saturation) have 2-cycle latency, the same as most regular MIPS32 arithmetic.
- Non-multiply DSP instructions which feature saturation or rounding have 3-cycle latency.
- Special DSP multiply operations (or any other access to the multiply unit accumulators): these have timings like standard multiply and multiply-accumulate instructions, so they're in with the multiply operations under the next heading.
- Instruction dependencies relating to different fields in the *DSPControl* register are tracked separately, and efficiently, as if they were separate registers. But any `rddsp` or `wrdsp` instruction which reads/writes multiple fields at once is dependent on multiple fields, and that can't be tracked through the CB system. Such a `rddsp` is not issued until all predecessors have graduated, and such a `wrdsp` must graduate before its successors can issue. You can often avoid this by using the "masked" versions of these instructions to read or write only a particular field.

## 19.2.11 Multiply/Divide Unit and Timings

As is traditional with MIPS CPUs, the integer multiplier is a semi-detached unit with its own pipeline. This pipeline implements:

- `mult/multu`: multiply two 32-bit numbers from GPRs (signed and unsigned versions) with a 64-bit result delivered in the multiply unit's accumulator. The accumulator was traditionally seen as pseudo-registers *hi* and *lo*, readable only using the special instructions `mghi` and `mflo`. Operations into the accumulator do not hold up the main CPU and run independently, but `mghi/mflo` are interlocked and delay execution as required until the result is available.
- `madd, maddu, msub, msubu`: multiply/accumulate instructions collecting their result in the accumulator.
- `mul/mulu`: simple 3-operand multiply as a single instruction.
- `div/divu`: divide - the quotient goes into *lo* and the remainder into *hi*.

Many of the most powerful instructions in the MIPS DSP ASE are variants of multiply or multiply-accumulate operations, and are described in the MIPS DSP chapter of this manual. The DSP ASE also provides three additional "accumulators" which behave like the *hi/lo* pair: the now four accumulators are called *ac0-3*. When we talk about the "multiply/divide" group of instructions we include any instruction which reads or writes any accumulator.

No multiply/divide operation ever produces an exception - even divide-by-zero is silent — compilers typically insert explicit check code where it's required.

Timing varies. Multiply-accumulate instructions (there are many different flavors of MAC in the DSP ASE) have been pipelined and tuned to achieve a 1-instruction-per-clock repeat rate, even for sequences of instructions targeting the same accumulator. But because that requires a relatively long pipeline, multiply/divide unit instructions which

produce a result in a GP register are relatively “slow”: for example, an instruction consuming the register value from a `mflo` will not be issued until at least 7 cycles after the `mflo`.

What that means is that in an instruction sequence like:

```
mult $1, $2
mflo $3
addu $2, $3, 1
```

The `mflo` will be issued 4 cycles after the `mult`, and the `addu` will go at least 2 cycles after the `mflo`. The execution unit may (or may not) be able to find other instructions to keep it busy, but each trip through that code sequence will take a minimum of 9 cycles.

## 19.3 Kernel Mode Programming

This section covers the following topics:

- [Section 19.3.1, "Hazard Barrier Instructions"](#)
- [Section 19.3.2, "Enhanced Interrupt System"](#)
- [Section 19.3.3 "External Interrupt Controller \(EIC\) Mode"](#)

### 19.3.1 Hazard Barrier Instructions

When privileged “CP0” instructions change the machine state, unexpected behavior can occur if an instruction is deferred out of its normal instruction sequence. But that can happen because the relevant control register only gets written some way down the pipeline, or because the changes it makes are sensed by other instructions early in their pipeline sequence: this is called a CP0 *hazard*.

Traditionally, MIPS CPUs left the kernel/low-level software engineer with the job of designing sequences which are guaranteed to run correctly, usually by padding the dangerous operation with enough `nop` or `ssnop` instructions.

To help manage pipeline hazards, the interAptiv core implements explicit *hazard barrier* instructions. If a hazard barrier instruction is executed between the instruction which makes the change (the “producer”) and the instruction which is sensitive to it (the “consumer”), you are guaranteed that the change will be seen as complete. Hazards can appear when the producer affects even the instruction fetch of the consumer - that’s an “instruction hazard” - or only affecting the operation of the consuming instruction (an “execution hazard”). Hazard barriers come in two strengths: `ehb` deals only with execution hazards, while `eret`, `jr.hb` and `jalr.hb` are barriers to both kinds of hazard.

In most implementations the strong hazard barrier instructions are quite costly, often discarding most or all of the pipeline contents: they should not be used indiscriminately. For efficiency you should use the weaker `ehb` where it is enough. Since some implementations work by holding up execution of all instructions after the barrier, it’s preferable to place the barrier just before the consumer, not just after the producer.

The following tables list the execution hazards and the instruction hazards for the interAptiv core.

#### 19.3.1.1 Execution Hazards

Execution hazards are those created by the execution of one instruction, and seen by the execution of another instruction. The following table lists possible execution hazards .

**Table 19.3 Execution Hazards**

Producer	→	Consumer	Hazard On
TLBWR, TLBWI, TLBINV, TLBINVF	→	Load/store using new TLB entry	TLB entry
MTC0	→	Load/store affected by new state	<i>WatchHi</i> <i>WatchLo</i>
MTC0	→	MFC0	Any CP0 register
MTC0	→	EI/DI	<i>Status</i>
MTC0	→	RDHWR \$3	<i>Count</i>
MTC0	→	Coprocessor instruction execution depends on the new value of Status <sub>CU</sub>	<i>Status<sub>CU</sub></i>

**Table 19.3 Execution Hazards (continued)**

<b>Producer</b>	<b>→</b>	<b>Consumer</b>	<b>Hazard On</b>
MTC0	→	ERET	<i>EPC</i> <i>DEPC</i> <i>ErrorEPC</i>
MTC0	→	ERET	<i>Status</i>
EI, DI	→	Interrupted instruction	<i>Status<sub>IE</sub></i>
MTC0	→	Interrupted instruction	<i>Status</i>
MTC0	→	User-defined instruction	<i>Status<sub>ERL</sub></i> <i>Status<sub>EXL</sub></i>
MTC0	→	Interrupted Instruction	<i>Status<sub>IM</sub></i> <i>(Cause<sub>IP</sub>)</i>
TLBR	→	MFC0	<i>EntryHi,</i> <i>EntryLo0,</i> <i>EntryLo1, PageMask</i>
TLBP	→	MFC0	<i>Index</i>
MTC0	→	RDPGPR WRPGPR	<i>SRSCtl<sub>PSS</sub></i>
MTC0	→	Instruction not seeing a Timer Interrupt	Compare update that clears Timer Interrupt
MTC0	→	Instruction affected by change	Any other CP0 register
CACHE	→	MFC0	<i>TagHi, TagLo, DataHi,</i> <i>DataLo</i>

### 19.3.1.2 Instruction Hazards

Instruction hazards are those created by the execution of one instruction, and seen by the instruction fetch of another instruction. Table 19.4 lists the instruction hazards. Because the fetch unit is decoupled from the execution unit, these hazards are rather large. The use of a hazard barrier instructions is required for reliable clearing of instruction hazards.

**Table 19.4 Instruction Hazards**

<b>Producer</b>	<b>→</b>	<b>Consumer</b>	<b>Hazard On</b>
TLBWR, TLBWI, TLBINV, TLBINVF	→	Instruction fetch using new TLB entry	TLB entry
MTC0	→	Instruction fetch seeing the new value including: <ul style="list-style-type: none"> <li>• change to ERL followed by an instruction fetch from the useg segment and</li> <li>• change to ERL or EXL followed by a Watch exception</li> </ul>	<i>Status</i>
MTC0	→	Instruction fetch seeing the new value	<i>EntryHi<sub>ASID</sub></i>
MTC0	→	Instruction fetch seeing the new value	<i>WatchHi</i> <i>WatchLo</i>
Instruction stream write via CACHE	→	Instruction fetch seeing the new instruction stream	Cache entries

**Table 19.4 Instruction Hazards (continued)**

Producer	→	Consumer	Hazard On
Instruction stream write via store	→	Instruction fetch seeing the new instruction stream	Cache entries

## 19.3.2 Enhanced Interrupt System

The features for handling interrupts include:

- **Vectored Interrupt (VI) mode** offers multiple entry points (one for each of the interrupt sources), instead of the single general exception entry point.


**External Interrupt Controller (EIC) mode** goes further, and reinterprets the six CPU interrupt input signals as a 64-value field - potentially 63 distinguished interrupts each with their own entry point (the zero code, of course, is reserved to mean “no interrupt active”).

Both these modes need to be explicitly enabled by setting bits in the *Config3* register; if you don’t do that, the CPU behaves just as the original (release 1) MIPS32 specification required.

- Shadow registers - alternate sets of registers, often reserved for interrupt handlers, are described in [Section 19.5, "Saving Power"](#). Interrupt handlers using shadow registers avoid the overhead of saving and restoring user GPR values.
- The *Cause[TI]*, *Cause[FDCI]*, and *Cause[PCI]* bits provide a direct indication of pending interrupts from the on-CPU timer, fast debug channel, and performance counter subsystems (these interrupts are potentially shared with other interrupt inputs, and it previously required system-specific programming to discover the source of the interrupt and handle it appropriately).

The new interrupt options are enabled by the *IntCtl* register, whose fields are shown in [Figure 19.3](#).

**Figure 19.3 Fields in the IntCtl Register**



*IntCtl*[*IPTI*,*IPPCI*,*IPFDCI*]: *IPTI*, *IPPCI*, and *IPFDCI* are read-only 3-bit fields. These fields indicate how the internal timer, performance counter, and fast debug channel interrupts are wired up. They are relevant in non-vectored and simple-vectored ("VI") interrupt modes, but have no meaning when using an EIC interrupt controller.

Read this field to get the number of the *Cause*[*IPnn*] where the corresponding interrupt is seen. Because *Cause*[*IP1-0*] are software interrupt bits, unconnected to any input, legal values for *IntCtl*[*IPTI*], *IntCtl*[*IPPCI*], and *IntCtl*[*IPFDCI*] are between 2 and 7.

The timer, performance counter, and fast debug channel interrupt signals are taken out to the CPU interface and the SoC designer connects them back to one of the CPU’s interrupt inputs. The SoC designer is supposed to hard-wire some CPU inputs which show up as the *IntCtl*[*IPTI*,*IPPCI*,*IPFDCI*] fields to match.

These interrupt outputs are per-VPE, so there are two of them from the interAptiv core. The *IntCtl* register is also per-VPE, reflecting the local setup.

*IntCtl[VS]*: is writable to provide software control of the interrupt vector spacing. The spacing is calculated as  $32 \times 2^{(VS-1)}$  bytes.

VS values of 1, 2, 4, 8 and 16 work provide spacings of 32, 64, 128, 256, and 512 bytes respectively. A value of zero gives a zero spacing, so all interrupts arrive at the same address. This would be the legacy behavior.

### 19.3.2.1 Traditional MIPS® Interrupt Signalling and Priority

In previous generation MIPS processors, the CPU takes an interrupt exception on any cycle where one of the eight possible interrupt sources visible in *Cause[IP]* is active, enabled by the corresponding enable bit in *Status[IM]*, and not otherwise inhibited. When that happens control is passed to the general exception handler and is recognized by the “interrupt” value in *Cause[ExcCode]*. All interrupts are equal in the hardware, and the hardware does nothing special if two or more interrupts are active and enabled simultaneously. All priority decisions are made by software.

Six of the interrupt sources are hardware signals brought into the CPU, while the other two are “software interrupts” taking whatever value is written to them in the *Cause* register.

The original MIPS32 specification adds an option to this. If you set the *Cause[IV]* bit, the same priority-blind interrupt handling happens but control is passed to an interrupt exception entry point which is separate from the general exception handler.

### 19.3.2.2 VI Mode - Multiple Entry Points, Interrupt Signalling and Priority

The traditional interrupt system commonly has a single piece of code which does the housekeeping associated with interrupts prior to calling an individual device-interrupt handler. However, a single entry point does not always fit well with embedded systems using very low-level interrupt handlers. These types of applications perform best when multiple entry points are provided. To accommodate this, the interAptiv core implements the “VI interrupt mode” where interrupts are despatched to one of eight possible entry points. To make this happen:

1. *Config3[VInt]* must be 1 to indicate that the core has the vectored-interrupts feature. This is a read-only bit that is always set in the interAptiv core to indicate support for vectored interrupts.
2. Set the *Cause[IV]* bit to request that interrupts use the special interrupt entry point.
3. Set the *IntCtl[VS]* to a non-zero value, setting the spacing between successive interrupt entry points. The interAptiv core allows spacing of 32, 64, 128, 256, and 512 bytes between entry points.

Interrupt exceptions vector to one of eight distinct entry points. The bit-number in *Cause[IP]* corresponding to the highest-numbered active interrupt becomes the “vector number” in the range 0-7. The vector number is multiplied by the “spacing” implied by the OS-written field *IntCtl[VS]* (see above) to generate an offset. This offset is then added to the special interrupt entry point (already an offset of 0x200 from the value defined in *EBase*) to produce the entry point to be used.

If multiple interrupts are active and enabled, the entry point will be the one associated with the higher-numbered interrupt: in VI mode interrupts are no longer all equal, and the hardware now has some role in interrupt “priority”.

## 19.3.3 External Interrupt Controller (EIC) Mode

Embedded systems have lots of interrupts, typically far exceeding the six input signals traditionally available. Most systems have an external interrupt controller to allow these interrupts to be masked and selected. In the interAptiv

core, EIC mode allows the six interrupt input signals to be encoded to allow up to 63 distinct interrupt entry points. In a 4-core system, this would allow for up to 256 distinct interrupts.

To do this the same six hardware signals used in traditional and VI modes are redefined as a bus with 64 possible values: 0 means “no interrupt” and 1 - 63 represent distinct interrupts. That’s “EIC interrupt mode”, and you’re in EIC mode if you would be in VI mode (see previous section) and additionally the *Config3[VEIC]* bit is set. EIC mode is a little deceptive: the programming interface hardly seems to change, but the meaning of fields change quite a bit.

Firstly, once the interrupt bits are grouped the interrupt mask bits in *Status[IM]* can’t just be bitwise enables any more. Instead this field (strictly, the 6 high order bits of this field, excluding the mask bits for the software interrupts) is recycled to become a 6-bit *Status[IPL]* (“interrupt priority level”) field. Most of the time (when running application code, or even normal kernel code) *Status[IPL]* will be zero; the CPU takes an interrupt exception when the interrupt controller presents a number higher than the current value of *Status[IPL]* on its “bus” and interrupts are not otherwise inhibited.

As before, the interrupt handler will see the interrupt request number in *Cause[IP]* bits; the six MS of those bits are now relabelled as *Cause[RIPL]* (“requested IPL”). In EIC mode the software interrupt bits are not used in interrupt selection or prioritization: see below. But there’s an important difference; *Cause[RIPL]* holds a snapshot of the value presented to the CPU when it decided to take the interrupt, whereas the old *Cause[IP]* bits simply reflected the real-time state of the input signals<sup>2</sup>.

When an exception is triggered the new IPL - as captured in *Cause[RIPL]* - is used directly as the interrupt number; it’s multiplied by the interrupt spacing implied by *IntCtl[RS]* and added to the special interrupt entry point, as described in the previous section. *Cause[RIPL]* retains its value until the CPU next takes any exception.

**Software interrupts:** the two bits in *Cause[IP1-0]* are still writable, but now become real signals which are fed out of the core, and in most cases will become inputs - presumably low-priority ones - to the EIC-compliant interrupt controller.

In EIC mode the usual association of the internal timer, performance-counter overflow, and fast debug channel interrupts with individual bits of *Cause[IP]* is lost. These interrupts are turned into output signals from the core, and will themselves become inputs to the interrupt controller. Ask your system integrator how they are wired.

## 19.4 Exception Entry Points

Early versions of the MIPS architecture had a rather simple exception system, with a small number of architecture-fixed entry points.

But there were already complications. When a CPU starts up main memory is typically random and the MIPS caches are unusable until initialized; so MIPS CPUs start up in uncached ROM memory space and the exception entry points are all there for a while (in fact, for so long as *Status[BEV]* is set); these “ROM entry points” are clustered near the top of *kseg1*, corresponding to 0x1FC0.0000 physical<sup>3</sup>, which must decode as ROM.

ROM is slow and rigid; handlers for some exceptions are performance-critical, and OS’ want to handle exceptions without relying on ROM code. So once the OS boots up it’s essential to be able to redirect OS-handled exceptions into cached locations mapped to main memory (what exceptions are not OS-handled? well, there are no alternate entry points for system reset, NMI, and EJTAG debug).

- 
2. Since the incoming IPL can change at any time - depending on the priority views of the interrupt controller - this is essential if the handler is going to know which interrupt it’s servicing.
  3. Even this address can be changed by a brave and determined SoC integrator, see the note on RBASE in [Section 19.4.1 “Summary of Exception Entry Points”](#).



So when *Status[BEV]* is flipped to zero, OS-relevant exception entry points are moved to the bottom of *kseg0*, starting from 0 in the physical map. The cache error exception is an exception... it would be silly to respond to a cache error by transferring control to a cached location, so the cache error entry point is physically close to all the others, but always mapped through the uncached “*kseg1*” region.

In MIPS CPUs prior to the MIPS32 architecture (with a few infrequent special cases) only common TLB miss exceptions got their own entry point; interrupts and all other OS-handled exceptions were all funneled through a single “general” exception entry point.

### **The CP0 EBase Register**

The *EBase* register provides the ability for software to identify the specific processor within a multi-processor system, and allows the exception vectors for each processor to be different. Bits 31:12 of the *EBase* register are concatenated with zeros to form the base of the exception vectors when *Status<sub>BEV</sub>* is 0. The exception vector base address comes from the fixed defaults when *Status<sub>BEV</sub>* is 1, or for any EJTAG Debug exception. The reset state of bits 31:12 of the *EBase* register initialize the exception base register to `0x8000.0000`, providing backward compatibility with Release 1 implementations.

The size of the *ExcBase* field depends on the state of the *WG* bit. At reset, the *WG* bit is cleared by default. In this case, the *ExcBase* field is comprised of bits 29:12. Bits 31:30 of the *EBase* Register are not writeable and are forced to a value of `2'b10` by hardware so that the exception handler will be executed from the *kseg0/kseg1* segments.

When the *WG* bit is set, bits 31:30 of the *ExcBase* field become writeable and are used to relocate the *ExcBase* field to other segments after they have been setup using the *SegCtl0* through *SegCtl2* registers. Note that if the *WG* bit is set by software (allowing bits 31:30 to become part of the *ExcBase* field) and then cleared, bits 31:30 can no longer be written by software and the state of these bits remains unchanged for any writes after *WG* was cleared. Therefore, it is the responsibility of software to write a value of `2'b10` to bits 31:30 of the *EBase* register prior to clearing the *WG* bit if it wants to ensure that future exceptions will be executed from the *kseg0* or *kseg1* segments.

Refer to [Section 2.2.1.9, "Exception Base Address — EBase \(CP0 Register 15, Select 1\)"](#) for more information.

## **19.4.1 Summary of Exception Entry Points**

The incremental growth of exception entry points has left no one place where all the entry points are summarized; so here's [Table 19.5](#). But first:

*BASE* is `0x8000.0000`, as it will be where the software, ignoring the *EBase* register, leaves it at its power-on value — that's also compatible with older MIPS CPUs. Otherwise *BASE* is the 4Kbyte-aligned address found in *EBase* after you ignore the low 12 bits...

*RBASE* is the ROM/reset entry point base, usually `0xBFC0.0000`. However, the interAptiv core can be configured to use a different base address by fixing some input signals to the CPU. Specifically, if the CPU is wired with *SI\_UseExceptionBase* asserted, then *RBASE* bits 29-12 will be set by the values of the inputs *SI\_ExceptionBase[29:12]* (the two high bits will be “10” to select the *kseg0/kseg1* regions, and the low 12 bits are always zero). Relocating *RBASE* is strictly not compliant with the MIPS32 specification and may break all sorts of useful pieces of software, so it's not to be done lightly.

*DebugVectorAddr* is an alternative entry point for debug exceptions. It is specified via a *drseg* memory mapped register of the same name and enabled through the Debug Control Register. The probe handler still takes precedence, but this is higher priority than the regular ROM entry points.

**Table 19.5 All Exception entry points**

<i>Memory region</i>	<i>Entry point</i>	<i>Exceptions handled here</i>
EJTAG probe-mapped	0xFF20.0200	EJTAG debug, when mapped to “probe” memory.
Alternate Debug Vector	DebugVectorAddr	EJTAG debug, not probe, relocated, $DCR[RDVec]=1$
ROM-only entry points	RBASE+0x0480	EJTAG debug, when using normal ROM memory. $DCR[RDVec]=1$
	RBASE+0x0000	Post-reset and NMI entry point.
ROM entry points (when $Status[BEV]=1$ )	RBASE+0x0200	Simple TLB Refill ( $Status[EXL]=0$ ).
	RBASE+0x0300	Cache Error. Note that regardless of any relocation of RBASE (see above) the cache error entry point is always forced into kseg1.
	RBASE+0x0400	Interrupt special ( $Cause[IV]=1$ ).
	RBASE+0x0380	All others
“RAM” entry points ( $Status[BEV]=0$ )	BASE+0x100	Cache error - in RAM. but always through uncached kseg1 window.
	BASE+0x000	Simple TLB Refill ( $Status[EXL]=0$ ).
	BASE+0x200	Interrupt special ( $Cause[IV]=1$ ).
	BASE+0x200+ . . .	multiple interrupt entry points - seven more in “VI” mode, 63 in “EIC” mode; see <a href="#">Section 19.3.2, “Enhanced Interrupt System”</a> .
	BASE+0x180	All others

## 19.5 Saving Power

There are just a couple of facilities:

- The **wait** instruction: this puts the thread running to sleep. When this happens when all other threads are sleeping, halted or suspended, the core goes into a low-power mode with many clocks stopped, from which it will only emerge when it senses an interrupt. The interrupt will be delivered to any sleeping thread, but *all* sleeping threads will wake and return from their **wait**. That will usually be OK; it’s normal practice to loop over **wait**. The **wait** instruction causes the core to enter a low-power sleep mode until woken by an interrupt. Most of the core logic is stopped, but the *Count* register, in particular, continues to run.
- The *Status[RP]* bit: this doesn’t do anything inside the core, but its state is made available at the core interface as *SI\_RP*. Logic outside the core is encouraged to use this to control any logic which trades off power for speed - most often, that will be slowing the master clock input to the core.
- Via the Cluster Power Controller, it is possible to gate off the clocks or even the power going to an idle core. This functionality is described in the *MIPS32® interAptiv™ Multiprocessing System Hardware User’s Manual*.

## C and C++ Efficient Programming Principles

This chapter describes methods to improve performance for programs written in high-level-languages (HLL) such as C and C++. This chapter is targeted at application-level programming and as such does not discuss instruction scheduling rules and requirements at the assembler language level.

Throughout this chapter, the following conventions are used.

- *Courier font*. New is used for source code including instruction names.
- *Italicized Courier font*. New is used for command lines including compiler switches.

### 20.1 Basic Compiler Switches for Optimization

This section describes the basic GCC compiler switches that are always be used for optimization.

- The *-march* compiler switch is used to denote which instruction set to use. MIPS strongly recommends you use the newest instruction set that is supported by your MIPS CPU.
- The *-mtune* compiler switch is used to denote which CPU pipeline that the compiler will schedule the instructions.
- The third compiler switch denotes which optimization level to use: *-O0* or *-O1* or *-O2* or *-O3*.

Note: *-O0* is usually only used when debugging, so that code is not optimized away and it is easier to debug. Be aware that *-O3* enables many optimization options which can cause the code-size to increase greatly, such as loop unrolling and function in-lining.

To learn more about the optimization controls, refer to the GCC documentation

<https://gcc.gnu.org/onlinedocs/gcc-5.1.0/gcc/Optimize-Options.html#Optimize-Options>

This webpage describes what compiler optimizations are enabled by these levels: *-O1*, *-O2* and *-O3*.

Commonly used values for these switches when using the interAptiv CPU are:

```
gcc -march=mips32r2 -mtune=1004Kf -O2 helloworld.c
```

The interAptiv pipeline is inherited from the previous generation 1004K pipeline, hence the usage of the 1004K pipeline description. The "f" denotes hardware floating-point. Use *-mtune=1004Kc* for soft-float.

The MIPS specific compiler switches are listed here:

<https://gcc.gnu.org/onlinedocs/gcc-5.1.0/gcc/MIPS-Options.html#MIPS-Options>

Other more specialized compiler switches are mentioned later in this document.

## 20.2 Data Alignment

The most common issue that is encountered when using MIPS processors for the first time is the alignment of data. Most MIPS CPUs did not support unaligned access of data in hardware. Instead such accesses would cause a hardware exception. For this reason, it is important to align data as much as possible for good performance.

Starting in 2014, some newer MIPS CPUs included hardware support for unaligned data access. Earlier MIPS CPUs from Cavium, NetLogic, and Loongson also had unaligned access support in hardware.

Even with hardware support, unaligned access will always give lower performance. This lower performance is true for any CPU, not just a MIPS CPU. The reason for this lower performance is that an unaligned access will cause 2 memory transactions to happen as opposed to 1 memory transaction for an aligned access. On the memory interface, there is no such thing as a true unaligned access. Instead, the unaligned data must be constructed from two aligned memory locations.

For application porting, the location of unaligned data is easy to find because the hardware will take an exception at the first unaligned access.

If you want the application to work without any porting effort, add exception handlers which emulate the unaligned access support. At an unaligned access, the exception handler would access the memory appropriately, extract the unaligned data and then return to the application. If the frequency of unaligned access is low, this type of trap-and-emulate strategy is very usable.

The MIPS instruction set has special instructions to mimic unaligned data access, the *lwl / lwr* and *swl / swr* instruction pairs. These instructions are shown in the following examples.

### 20.2.1 Aligned Access Example

The following C code shows a structure in which the compiler aligns naturally (this is the default behavior of the compiler):

```
struct t1 {
    char a;           // placed at offset 0
    int b;           // placed at offset 4 due to regular alignment rules
    int c;           // placed at offset 8 due to regular alignment rules
};
int main()
{
    struct t1 x;
    x.b = x.c;
}
```

The compiler converts the C statement `x.b = x.c` into the following assembler code with aligned access:

```
lw      v0,16(s8)
sw      v0,12(s8)
```

### 20.2.2 Unaligned Access Example

Here's a data structure where the packed keyword is used to tell the compiler to ignore the normal alignment rules, creating unaligned data.

```

struct t2 {
    char a;           // placed at offset 0
    int b;           // placed at offset 1 due to packed keyword
    int c;           // placed at offset 5 due to packed keyword
} __attribute__((packed)); // pack the data members as close as possible
                               // ignore the normal data alignment rules

int main()
{
    struct t2 y;
    y.b = y.c;
}

```

The compiler converts the C statement `y.b = y.c` into the following assembler code with unaligned access:

```

lwl    v0,25(s8)    // get the left part of the data
lwr    v0,28(s8)    // get the right part of the data
swl    v0,21(s8)    // assign the left part of the data
swr    v0,24(s8)    // assign the right part of the data

```

## 20.3 Dealing with Caches

Another important performance aspect of CPU power management is working with and optimizing the use of caches.

### 20.3.1 Coherent Systems and Non-Coherent I/O

For multiprocessor systems, the coherency of caches of the multiple CPUs must be taken into consideration. If the hardware provides coherency support for all masters touching data in memory, no software maintenance of coherency is needed. For example, if all of the CPUs are in the same coherency domain with MIPS Coherency Manager and all non-CPU IO writes are using MIPS IOCU block.

On the other hand, some IO writes are not coherent (such as not using the IOCU or equivalent) making software maintenance of the caches necessary. For example, if a non-coherent IO device updates data that might be held in one of the CPU caches. In that case, the cache instruction is needed to invalidate the appropriate locations within CPU caches as those cache locations would be holding stale versions of the data.

Below is an example of how the cache instruction is used as an inlined assembly macro call.

```
#define Hit_Invalidate_D          0x11

#define cache_op(op,addr)

    __asm__ __volatile__(
        ".set    push          \n    \"\n
        ".set    noreorder     \n    \"\n
        "cache  %0, %1        \n    \"\n
        ".set    pop           \n    \"\n
        :
        : "i" (op), "R" (*(unsigned char*)(addr)))
```

1. The volatile keyword tells the smart inline assembler to not optimize this code away.
2. The push and pop pseudo-instructions tell the inline assembler to save its internal state.
3. The noreorder pseudo-instruction tells the inline assembler to not change the instruction sequencing (the MIPS assembler is smart and can optimize assembler instruction sequences by itself).
4. The "i" identifier tells the inline assembler that the first operand - %0 is an immediate value.
5. The "R" identifier tells the inline assembler that the second operand - %1 is an address register.

Documentation on how to use inline assembly with GCC can be found at the following website:

<https://gcc.gnu.org/onlinedocs/gcc-5.1.0/gcc/Using-Assembly-Language-with-C.html#Using-Assembly-Language-with-C>

Below is a C routine which calls the inline assembly routine. This routine invalidates one cache-line. This puts this C routine in a loop to invalidate a multiple cache-line sized region of the dcache.

```

static inline void flush_dcache_line(unsigned long addr)
{
    cache_op(Hit_Invalidate_D, addr);
}

```

Note: The Hit type of cache instructions are more appropriate for MIPS SMP products as the Hit types broadcast their operation across the entire CPU cluster to all of the coherent CPUs.

## 20.3.2 Data Prefetch

Data prefetching is one method of reducing performance loss due to memory latency. If a loop is accessing new locations from memory in each iteration, data prefetch instructions can be used to fetch ahead for the next loop iteration. In this way, when the next loop iteration starts, the requested data may already be present in the CPU caches, without having to wait for data to arrive from system memory.

The MIPS instructions for software prefetch are `pref` for integer data and `prefx` for indexing through arrays.

To avoid the latency of the memory transactions, the software prefetch instructions must be executed early and in time so that the data arrives into the caches before the data is used by the next loop iteration.

The following is an example of how the `pref` instruction is defined as an inline assembly macro routine.

```

#define Pref_Load      0      // cache-fill as shared
#define Pref_Store    1      // cache-fill as exclusive

__asm__ __volatile__(
    ".set    push \n      \"\n
    ".set    noreorder \n  \"\n
    "pref  %0, %1 \n    \"\n
    ".set    pop \n      \"\n
    :
    : "i" (op), "R" (*(unsigned char*)(addr)))

```

Note: On MIPS CPUs (including InterAptiv), each prefetch instruction needs an available FSB (Fill-Store-Buffer) to hold the incoming cache-line data. For that reason, the number of prefetch instructions that can be effectively used is limited by the number of available FSB entries.

## 20.3.3 Mixing Cached and non-Cached Writes

On most CPUs, cached writes normally first go to a write-buffer holding multiple cache-lines of data before being sent to the memory hierarchy after the L1 cache. This write-buffer allows the execution pipeline to not stall and also allows the memory transaction to happen opportunistically later.

On most CPUs, uncached writes normally go to a different write-buffer or go directly into the memory hierarchy external to the CPU.

Because cached writes and uncached writes don't use the same physical path to the memory hierarchy, it is possible that the actual order of when the transactions occur at external memory is different from the program order.

If the actual order of memory writes must match the program order, the sync instruction can be used to enforce ordering.

Below is an inline assembly macro of the sync instruction. It is used in the C example code to force the order of a cached write-back and an uncached write.

```
/* SMP write memory barrier */
/* Completion barrier, can be used for dev communication */
# define mb() __asm__ __volatile__("sync" ::: "memory")
    // cached write
    FrameBuffer[max_x][max_y] = 0xAAAA.AAAA;

    cache_op(Hit_Writeback_Inv_D, &(FrameBuffer[max_x][max_y]));

    // push-out frame buffer before sending new command
    mb()

    //uncached write
    FrameBuffer_ctrl_reg = new_command;
```

### 20.3.4 Self-modifying Code

For self-modifying code or code generated at run-time, care must be taken that the generated code is placed into a location within the memory hierarchy where the icache can access the new code.

For example, if the L2 cache is the first level of the memory hierarchy which is shared between the icache and dcache, the newly generated code must be written back to the L2 cache, so the icache can fetch the new code.

The self-modifying code procedure is shown below:

1. Place the new instructions into the level of the memory hierarchy which is shared between the icache and dcache. Here in our example, that is the L2 cache. You would do that by using the Hit\_Writeback\_D flavor of the cache instruction, using an inline assembly macro similar to the previous example.
2. Invalidate the stale copy of the instructions that might be held in the icache. You would do that using the Index\_Invalidate\_I flavor of the cache instruction. Again, using an inline assembly macro similar to the previous example.
3. Wait till the new instructions actually reach the L2 cache and then are written there. The sync instruction is used for this purpose.
4. Wait till the icache invalidation completes before fetching the new instructions. One of the hazard-barrier avoiding instructions is used for this purpose.

Below is an example of how steps 3 and 4 are defined as an inlined assembly routine. The sync instruction stalls the instruction fetch until there are no more instruction hazards. The sync instruction stalls the pipeline until the writebuffers are emptied. The jr hb instruction stalls the pipeline until any other type of hazard has completed. (This step isn't needed for our example, but this is a routine copied from real code and it is just being careful.)

Here the inline assembly macro is inside the C routine r4k\_flush\_cache\_sigtramp.



```

static void r4k_flush_cache_sigtramp(void * arg) {
    register unsigned long addr = (unsigned long) arg;

    __asm__ __volatile__(
        "synci 0(%0) \n" \
        "sync \n" \
        "jr hb $1 \n" \
        "l: \n" \
        ::"r"(addr):"memory"); \
}

```

### 20.3.5 Minimizing Memory Transactions

Because the CPU runs at much higher frequency than system memory, reducing cache misses is one way to improve performance. One way to reduce cache misses is to place data into locations that require the minimal number of cache lines to be filled from system memory.

There are several techniques that can be used to avoid dcache misses.

For example, in a CPU with 4-way set associative caches, the calculation might need data from 5 arrays. The 5 arrays might be arranged in memory, so the necessary values from all 5 arrays would use the same set of locations in the dcache. The 4-way cache might only be able to hold values from 4 of the arrays at the same time, but not all 5 arrays. In this case, each calculation would experience at least one dcache miss and thus have some performance loss. To avoid the cache miss, one of the data arrays could be moved in memory to use different addresses and thus use a different location within the dcache.

Reducing stack usage can minimize cache misses:

- Avoiding recursive calls
- Minimizing the scope of local variables, so they are brought in only when really needed.

Another strategy is called Tiling. Tiling is also known as loop blocking or strip-mining. The idea is to break up very large arrays into multiple smaller arrays. The data would be split up so that the calculations would only be accessing smaller arrays at the same time, which has better chance of all of the required data to be present within the dcache at the same time.

## 20.4 General Porting Considerations

- As previously discussed, unaligned memory access causes a hardware exception on most MIPS processors.
- While porting code from another CPU architecture, the data alignment rules might be different and thus any assumptions about data layout in memory might cause problems.
- MIPS processors can be configured for either Big-Endian or Little-Endian execution. If you are porting a program written for the opposite endian-ness of what you plan to use, sub-word memory accesses within the program must be fixed.
- Null pointers (with value of 0x0) might cause a TLB miss exception as the virtual address 0x0000.0000 is a mapped address.
- Conversions from ints to pointers (addresses) must be explicit. Pointers must be handled as unsigned values. This is to avoid overflow when doing address math on MIPS32 kernel-space addresses (which have bit 31 set).

- MIPS processors are either 32-bit or 64-bit in nature. Both the register files and arithmetic operations of these widths. These are the natural data-widths of the hardware. There are no 8-bit (char) nor 16-bit (short) arithmetic operations in hardware. Instead these 8-bit or 16-bit operations are synthesized from multiple 32-bit operations. For these reason, it is often more efficient to declare variables as full 32-bit values (ints).
- In most cases, it's more efficient to use 32-bit variables as opposed to shorter values. There are some exceptions though. For example: some old x86 programs might depend on overflowing 16-bit values.
- MIPS GCC treats chars as unsigned by default. Other architectures might treat chars as signed by default. It is recommended to be explicit and define variables as either signed char or unsigned char. You can also use the compiler flag `-fsigned-char` to change the default choice.
- Some old programs without function prototypes might have some function input parameters automatically cast as 32-bit ints.

## 20.5 Migrating from ARM to MIPS

Software is the key element in any project because you must choose an instruction set architecture (ISA) that is a scalable solution for your future development. MIPS architecture is that scalable solution. This section describes considerations you must explore for a smooth transition from ARM architecture to the MIPS scalable solution.

- As previously described, most MIPS processors do not have unaligned data support in hardware. Porting code originally created for any other non-MIPS processor usually means re-aligning some data.
- MIPS has hardware divide and modulo instructions. MIPS does not have or need to have helper library functions for these instructions.
- By default, ARM-based systems don't perform divide-by-zero checks while MIPS-based systems do this check. To remove this check, use `-mno-check-divide-by-zero` (to mimic ARM behavior).
- MIPS doesn't care about For loop direction for the index variable. It makes no difference whether the loop variable counts away from zero or towards zero.

## 20.6 Helping the Compiler

Casting a pointer to a different type (excluding casts to or from `'void *'`) is not recommended because you can generate misaligned memory references due to differing default alignments for the types. Additionally, the optimizer may make assumptions about the types that is incorrect due to the cast - incorrect code could be generated, or more conservative code could be generated that performs poorly.

Use the `__restrict` C99 keyword to tell the compiler that no other pointer is accessing the data referenced by this pointer. This helps the compiler know that there is no pointer aliasing for this data. Avoid using the `'restrict'` keyword on parameters of inline functions.

Inlining can cause improper code to be generated.

```
void mem_op (int *src, int *dest, int num)
{
    for (int i = 0; i < 16; i++)
    {
        src[i] = dest[i] + num;
    }
}
```

In the above example, compiler assumes that `src` and `dest` might point to overlapping memory blocks, thereby following strict order of loads and stores pertaining to loop. `__restrict__` is basically a promise to the compiler that for the scope of the pointer, the target of the pointer will only be accessed through that pointer. Because of this, compiler can have more freedom in moving load/store instructions without changing semantics of the program, resulting in better schedule.

Below is the same code with the `__restrict__` keywords added:

```
void mem_op (int * __restrict__ src, int * __restrict__ dest, int num)
{
    for (int i = 0; i < 16; i++)
    {
        src[i] = dest[i] + num;
    }
}
```

Referencing data through pointers prevents the compiler from using many optimizations as the compiler can't tell if other references are accessing the same data or not. For that reason, data variables created at compile time is preferred over data that is created at run-time and referenced through pointers.

The program might want to read data that is written by an I/O peripheral. If the I/O peripheral is always updating the value of its data register, your program ought to always load from the memory-mapped data register instead of reusing any stale value that was previously loaded into a CPU register. You can use the `volatile` keyword to denote that the data variable needs to be read from memory each time before being used by another instruction.

## 20.7 O32 ABI Issues

The most popular 32-bit ABI used for MIPS is called O32. The O32 ABI is the default for 32-bit compilation. If you want to specify it explicitly, use `-mabi=32`.

Here are some O32 specifics that might affect high-level-language (HLL) program tuning.

### 20.7.1 Data sizes and alignment

Table 20.1 Data Sizes and Alignment

Data Type	Int	Long Int	Long Long Int	Pointer	Char	Short Int	Float FP	Double FP	Quad FP
Width in bytes	4	4	8	4	1	2	4	8	8
Default alignment in bytes	4	4	8	4	1	2	4	8	8

### 20.7.2 Integer Calling Convention

The first 4 words (16 bytes) of the input argument list is passed in hardware integer registers. The remaining input arguments are passed on the stack. Functions which only use 4 words of input arguments would be the fastest because they avoid having to access the stack.

The return value is passed in a hardware integer register.

```
// all 4 inputs are passed in HW registers
int i_function1( int input1, int input2, int input3, int input4);

// the last 2 inputs are loaded from the stack
int i_function2( int input1, int input2, int input3, int input4, int input5, int input6);
```

### 20.7.3 Floating-Point Calling Convention

If the first input argument is a floating-point value, then up to the first 4 words (16 bytes) are passed in hardware floating-point registers. The remaining input arguments are passed on the stack. Functions which only use 4 words of input arguments would be the fastest as they avoid having to access the stack.

If the return value is a floating-point value, it is returned in a hardware floating-point register. If the return value is an integer value, it is returned in an hardware integer register.

```
// all 4 inputs are passed in HW registers
float fp_function1(float input1, float input2, float input3, float input4);

// the last 2 inputs are loaded from the stack
float fp_function2( float input1, float input2, float input3, float input4, float input5, float input6);
```

### 20.7.4 Large Structure as Return Value

If the returned object is a structure too large to fit in the 2 return registers, the structure is placed on the stack. A pointer to the structure is prepended to the input argument list.

For the function return, the return register holds the pointer to the structure on the stack.

### 20.7.5 Stack and Heap Overview

In most MIPS 32-bit systems:

1. The stack starts at a high address near the largest user-space address (0x7FFF.FFFF for MIPS32) and grows downwards towards smaller addresses (grows towards 0x0000.0000).
2. The program code resides within the lowest 1 GB of virtual address space.
3. The heap starts after the program code region and grows upwards toward larger addresses (grows towards 0x7FFF.FFFF).
4. When using the legacy address map, the top 2GB of virtual address space is only accessible while running in kernel mode.

## 20.8 Customer Code Optimization Overview

### 20.8.1 Profile your Application

On Linux, popular profiling tools include the older OProfile or newer perf tools. Using these profiling tools, you would identify which routines take up most of your execution time. This is done by doing statistical sampling of the program counter values. Those routines would be the ones most worthwhile to optimize first.

Setting the performance counter to other types of events - such as cache misses, branch mis-predicts, TLB misses - one can get hot-spot locations for those events as well.

## 20.8.2 Optimize Routines

There are several methods of optimizing the identified routines:

- Use additional compiler switches if the code has heavy usage of specific C language constructs (for example multiple levels of nested loops).
- Another method is to use data prefetching.
- Another possibility is to convert FP operations to lower-latency integer/fixed-point operations
- Another possibility is to inline functions to avoid the subroutine call cost. Use the inline keyword.
- Other methods include using intrinsics or inlined assembly to use instructions whose semantics aren't understood by the compiler.

## 20.8.3 Intrinsics

Intrinsics are another way to use assembly-language level instructions. Intrinsics are also known as built-in functions. They are similar to assembler macros except that the compiler handles the register allocation and scheduling.

Some examples for the MIPS DSP Module:

```
typedef int i32;

// a routine to do bit reversal of a 32-bit integer
i32 __builtin_mips_bitrev (i32);

typedef short v2q15 __attribute__((vector_size(4)));

// a routine to do element-wise addition for
// 2-element vectors of Q15 fixed-point values
v2q15 __builtin_mips_addq_s_ph (v2q15, v2q15);
```

## 20.9 Overview of Compiler Optimizations

### 20.9.1 List of Automatic Compiler Optimizations

The GCC compiler uses the default settings of the following optimizations. That is, you do not have to add any compiler switches for the following described optimizations.

#### 20.9.1.1 Data flow Analysis Optimizations

The following common optimizations are part of data-flow analysis and are done by the GCC compiler when the user program is represented in an intermediate-language that is not CPU specific. The code within the compiler that does this is shared among all CPU architectures.

- Common Sub-Expression Elimination - re-use any calculation as much as possible

- Constant folding - replace a calculation with constant if possible
- Induction variable elimination - recognize a variable is a simple function of the loop index
- Pointer Alias analysis - determine which pointers may reference same locations

### 20.9.1.2 Loop Optimizations

The following common optimizations are part of loop analysis and are done by the GCC compiler when the user program is represented in an intermediate-language that is not CPU specific. The code within the GCC compiler that does this is shared among all CPU architectures.

- Loop invariant code motion - execute code only once not every iteration
- Loop Unrolling - to have more instructions executing in parallel. MIPS CPUs have more registers, making it easier to do this
- Loop Tiling - explained earlier in this document
- Loop Interchange - to get better data locality
- Loop Fusion - combine multiple loops to avoid unnecessary loop index updates & branches
- Strength reduction - convert long latency operations to shorter latency operations
- Function Inlining - remove the overhead of subroutine calls

### 20.9.1.3 Code Generation Optimizations

The following common optimizations are part of the compiler back-end which is instruction set specific.

- Register allocation - the most heavily used variables are kept in hardware registers
- Instruction scheduling - avoiding pipeline stalls due to data-dependencies
- Removing redundant branches and jumps

## 20.9.2 List of Manually Chosen Compiler Optimizations

Here are some compiler switches which can be useful in optimizing benchmark performance.

**Note: Remember that your application might be different and the results you get by using these switches might be different from the historical results.**

For Linux-based benchmarks, these compiler switches have been used on multiple benchmarks:

```
-O3 -march=mips32r2 -mtune=1004Kf -funroll-loops -fprefetch-loop-arrays
```

Then, for specific tests, you can increase the strength of the unrolling, by using:

```
--param max-unrolled-insns=100
--param max-inline-insns-auto=5
```

For the bare-iron environment, there is a small-data section which allows more efficient access to small constants. You can control what size of values that can be placed in the small-data-section with this switch:

```
-G<number>Put global and static data smaller than <number> bytes into a special section (on some targets)
```

For example, *-G8* or *-G4* - the latter was tested to give higher perf for Coremark than the former. So this is a tunable factor.

### 20.9.3 Compiler Optimizations for Code-Size

The first compiler switch to use for code-size is *-Os*.

The next set of switches to try are:

*-fshort-enums, -fno-inline-small-functions, -Wl --relax, --gc-sections*

For each file that is compiled, you can choose between the "regular" instruction sets (MIPS32 or MIPS64) and one of the "compressed" instruction sets (MIPS16e or microMIPS). Use the *-minterlink-mips16* compiler switch to enable interlinking between MIPS16e and MIPS32.

The choice between "regular" instruction set and "compressed" instruction set can even be done at the function level within one file. There are *micromips, nomicromips, mips16* and *nomips16* function attribute keywords.

Refer to MIPS document number MD00842 "*Using GCC Toolchain Options to Optimize Code Size*" for a more complete listing of compiler switches which affect code-size.

### 20.9.4 Optimized Libraries

#### 20.9.4.1 Optimized C Library Routines

There are specific C library routines that are known to be heavily used and thus can affect application performance. Such routines include *memcpy, memset, strcmp, strcpy*, and so on.

MIPS supplies free open-source versions of these routines which have been optimized for the MIPS CPU products.

#### 20.9.4.2 The MIPS DSP Library

An optimized library of useful routines for signal processing, vector math, FIR/IIR filters, FFTs, video processing. There are versions for MIPS base architecture, DSP Module instructions, MSA Module instructions. Available from MIPS.

For more information, refer to MIPS document number MD00472 "*MIPS® DSP Library Reference Manual*".

#### 20.9.4.3 GoFast Floating-Point Library

This is a performance optimized soft-float library. Available at <http://www.smxrtos.com/ussw/gofast.html>

## 20.10 Linking to Libraries

The compiler usually comes with a C library that supplies routines to make the system more usable the ability to create I/O, the ability to access a file-system, the ability to interact with the operating system, the facilities to handle strings, and so on.

There are a couple of ways of using the C library.

Refer to the GCC documentation on linking:

<https://gcc.gnu.org/onlinedocs/gcc-5.1.0/gcc/Link-Options.html#Link-Options> .

### 20.10.1 Static Linking

With static linking, the C library routines which are used by the application are included within the application binary file. This gives the best performance but can make the application very large. You specify static linking with the *-static* switch.

For large systems like Linux, where there are numerous applications, static linking is normally not used as it is very space inefficient.

Static linking is often used for creating benchmark scores. But in real life, many larger-scale operating systems don't use static linking.

### 20.10.2 Dynamic Linking

In large systems, like Linux, with numerous applications, dynamic linking is normally used. Here, the libraries are stand-alone objects placed in their own location in the address map.

During application load-time, the addresses of the used libraries are loaded into system memory. The libraries are shared among all applications which are using the services of the library. This sharing saves a lot of memory space.

The libraries are loaded into memory in an "as-used" sequence and the memory location of the dynamically linked library might change from one time to another, depending on which other libraries were already loaded.

You specify dynamic linking with the *-shared* switch.

### 20.10.3 Position Independent Code

Because these libraries are loaded into memory only at usage, they might change locations due the previous loading of other libraries.

The ability for the code to reside at different locations requires that all jumps and branches to be PC-relative. This type of code is called "Position Independent Code".

There are GCC compiler switches which are used to emit PIC code, such as *-fpic* .

## 20.11 Thread Synchronization and Memory Ordering

Atomic operations are needed to get exclusive access to resources which are shared among multiple threads.

### 20.11.1 Atomic Operations

In the MIPS architecture, the *ll* (load-linked) and *sc* (store conditional) instructions are the primitives used to create the atomic operations.

These two instructions are mimicking an atomic sequence - first, it reads a memory location and then attempts to write that memory location. The write only succeeds if there is there no other access to that memory location and no other thread is executed between the load and the store.



This sequence does not cause any real locking of any memory interface nor cause any type of stall within the CPU. It just checks to see if the entire read-write code sequence was run without any interference nor any interruption.

Operations like spin-locks, counting semaphores, compare-and-swap, fetch-and-add operations all can be built from these two instructions.

Below is an example of a spin-lock implementation when the semaphore is cacheable.

Acquire:

```
ll t0, 0(a0)           // Read the semaphore.
bne t0, 0, Retry       // If non-zero, someone else already
li t0, 0x1             // set semaphore, so try again.
sc t0, 0(a0)          // Try writing semaphore.
bne t0, 0, Release     // Check if LL-SC sequence was atomic
nop                   // if not, try again.
```

Retry: // Need memory barrier here

```
pause                 // Wait for LLBit to clear
beq 0,0,Acquire // Restart the hole process
nop
```

// Critical section

Release:

```
sync 0
sw 0, 0(a0)
```

Note: The pause instruction causes the Virtual Processor to sleep until the lock is available. This decreases both system traffic and system power and allows the lock to be returned sooner since there is less contention for the lock. The PAUSE instruction will not wake up when the semaphore is uncached and therefore PAUSE should not be used in an uncached LL/SC sequence.

## 20.11.2 Memory Barriers

The MIPS architecture provides two types of memory barriers:

- A heavier type which ensures that data leaves the CPU boundary and enters the rest of the system before the memory barrier completes. This heavier type of barrier is known as "Completion barrier".
- A lighter type which only ensures the order of the data which leaves the CPU boundary, but does not stall until the data leaves the CPU boundary. This lighter type of barrier is known as "Ordering barrier".

## 20.11.3 Weak Memory Model

The MIPS architecture follows the weak memory model. Here is an example of what is allowed in a weakly ordered memory model.

Initially variables X=1 and Y=1. X and Y are located in addresses which would be held in different cache-lines. X and Y are accessed with a cached coherency. Each physical core has its own separate dcache.

**Table 20.2 Weak Memory Model**

Thread 1 running on Physical Core1	Thread 2 running on Physical Core2
Instr1 - Store a value of 2 to variable X	Instr3 - Load variable Y and get a value of 2
Instr2 - Store a value of 2 to variable Y	Instr4 - Load variable X and get a value of 1

Analysis:

Implementations are allowed to re-order the memory transaction of Instr2 ahead of the transaction for Instr1, because the stores are different cache-lines. It is possible Instr4 can receive the updated store value while Instr3 does not.

Implementations are allowed to reorder the memory transaction of Instr4 ahead of the memory transaction of Instr3, because the loads are to different cache-lines. It is possible Instr4 can be executed before the store value of Instr1 is globally visible.

To enforce program ordering of the visibility of the store values, you can place a `sync` instruction between Instr1 and Instr2. To enforce program ordering of the load memory transactions, a `sync` instruction between Instr3 and Instr4.

Note: Uncached writes (using Coherence Attribute = 2) are always strongly ordered. If you use Coherency Attribute = 2, there is less need for `sync` instructions. Coherency Attribute = 2 is normally used for access to IO registers.

The main features of the weak memory model are:

- Stores are allowed to be buffered within the CPU and for this reason the visible order of stores within the system can be different from the order seen by the local CPU which generated the store.
- If a specific ordering is required for stores to different addresses in different cache-lines, memory barriers are needed to be added.

Here is an example of local visibility of memory updates versus global visibility of memory updates. Initially variable X=1. X is accessed with a cached coherency. Each physical core has its own separate dcache.

**Table 20.3 Local and Global Visibility of Memory Updates**

Thread 1 Running on Physical Core 1	Thread 2 Running on Physical Core2
Instr1 - Store a value of 2 to variable X	Instr3 - Store a value of 3 to variable X
Instr2 - Load variable X and get a value of 2	Instr4 - Load variable X and get a value of 3

Analysis:

Processors are allowed to buffer writes. There can be a time delay between when the write is first buffered and when the write is made visible to non-local processors. There is neither requirement that either store has been made globally visible when Instr2 executes nor when Instr4 executes. Processors are allowed to forward data from the write buffer to subsequent loads within the local processor before the store is globally visible.

Because there is no pre-ordained ordering between the times when the store values of Instr1 and Instr3 becoming globally visible, Instr2 gets its value as if Instr1 was the last globally visible store to X at that time. Similarly Instr4 gets its value as if Instr3 was the last globally visible store at that time. After both stores are globally ordered, the

instruction which received the stale value can be considered to have executed before the latter store in the global order.

## 20.12 Programming Considerations for MIPS Virtual Processors

### 20.12.1 Use all Physical Cores

For systems with multiple physical processors, spread out the threads to use all of the physical processors as appropriate to maximize the performance.

Similarly, when trying to save power, move the threads to the appropriate number of physical processors and keep those physical processors powered up. Then power down the other physical processors which are idle due to the thread migration.

### 20.12.2 Sharing of Pipelines - Potentially Bad for Performance

Remember that Virtual Processors share the execution pipelines of the physical processor. User's should avoid situations where multiple threads are meant to be running simultaneously on the same physical processor (but different Virtual processor).

For better performance, it would be better if these simultaneous threads are spread out among multiple physical processors.

### 20.12.3 Sharing of Caches - Good for Temporal and Spatial Locality

Remember that Virtual Processors also share the caches of the physical processor. For threads that share a lot of data, it might make sense for those threads to be running on the same physical processor (but different Virtual Processor). The sharing of data would happen naturally as the caches are shared between the Virtual processors within the same physical processor.

For such cases, consider processor affinity for these data-sharing threads.

### 20.12.4 Producer-consumer Examples

Here are some examples to show the effects of the pipeline sharing and cache sharing issues.

Case 1: Let's say there is a producer thread that creates data and a consumer thread that uses that data. In this first case, the two threads are interlocked - the consumer thread can only run after the producer thread creates new data. That is, there is no overlap in time when the two threads are running at the same time. For this case, it is advantageous to run the two threads on two Virtual Processors on the same physical processor as there is no danger of resource contention of the execution pipelines and you get the benefit of the data locality of sharing the caches.

Case 2: There is still the producer thread and the consumer thread. But now, the two threads are always running. That is, both are running at the same time. For this case, you might encounter a situation where both threads are stalling each other by trying to use the execution pipelines at the same time. For this situation, it might be better to move one of the threads to another physical core to avoid such resource contention.

## 20.13 Multi-threading Software - An Overview

This section is a quick overview of multi-threaded software environments. A thread is created for execution that can be processed in parallel with other parts of the application. There are two types of such parallelism - one is for data that can be processed in parallel and the other is for tasks that can be processed in parallel.

### 20.13.1 Threading for Tasks

Historically, creating separate threads that can run in their own context from a single-threaded application has been a process done by manual inspection. Such manual inspection is prone to errors, which can cause bugs such as race-conditions (described later in this paper).

There are companies which sell automated tools which perform such analysis. Critical Blue and Vector Fabrics are two companies which sell such tools.

### 20.13.2 Fork Join Model

The most popular execution model of threading is the Fork-Join model. Here each individual thread is created and ended on an as-needed basis.

### 20.13.3 Thread Pool Model

Another execution model of threading is the Thread pool. Here a fixed number of threads are created upon the application start-up. These threads will be used over and over again. The threads will be moved among the "running", "ready to run" and "not-runnable/blocked" lists as execution goes on over time.

### 20.13.4 PThreads

The most popular API to create threads is the Posix Threads or Pthreads library. This is the library that is used for most UNIX-like operating systems such as Linux.

Some books covering the Pthreads API library include:

- Programming with POSIX Threads by David R. Butenhof
- PThreads Programming: A POSIX Standard for Better Multiprocessing by David Buttlar, et. al.

There are other libraries that deal with threads at higher abstraction levels such as OpenMP, Intel TBB, Cilk+, Apple Grand Central Dispatch. Some of these use Pthreads at the thread creation level.

### 20.13.5 Thread-Safety and Reentrancy

Routines can be made thread-safe (that is, runnable by multiple threads without problems) by ensuring that:

- State is stored using thread-local storage whenever possible
- Shared (global or static) state is always protected by semaphores, to ensure only one thread has access at any one time

Routines which are re-entrant (that is, are interrupt-able and re-startable without problems) have additional requirements:

- All data must be supplied by the caller
- The routine cannot hold shared (global or static) data over successive invocations
- The routine cannot return data in shared (global or static) state
- The routine cannot call other routines which are not re-entrant.

Only routines which are both thread-safe and re-entrant can be used for multi-threaded programs.

## 20.13.6 SMP/MT Bugs and Issues

Below is a list of performance and functional issues that commonly affect MP/MT software.

### 20.13.6.1 Race Condition

Correctly written code would protect shared data variables with locks/semaphores. If shared data is not protected in this way, a data-race is possible. Two threads would attempt to access the shared variable. Both threads believe it has exclusive access to the data variable. Now it is possible that the later thread can over-write the value from the earlier thread, thus some data is lost.

### 20.13.6.2 Priority Inversion

Priority inversion is where a lower-priority thread/task takes up execution time while a higher priority task is unexpectedly blocked from executing.

### 20.13.6.3 Deadlock

This is a system hang where two threads (or more) are waiting for events which are controlled by the other. In this case, there is no movement of the program counters as both threads are stalled.

### 20.13.6.4 LiveLock

This is another type of system hang where two threads (or more) are perpetually trying to complete some transaction, but cannot due to interference from the other thread. Each of the threads keeps re-trying their transaction, but their transaction keeps failing to complete. In this case, there is movement of the program counters, but the execution trace look like an endless loop is being executed.

### 20.13.6.5 Cache thrashing

This is a case of poor performance (or even live-lock) caused by data that is perpetually migrating from one cache to another due to multiple threads trying to access the same cache line. The data migrates before any of the threads have completed their tasks which require the contents of the shared cache-line.

### 20.13.6.6 False Sharing

This is a special case of cache thrashing. Each thread is trying to access its own data variable which is held within the same cache-line but not the same set of bytes as the variables owned by the other threads. It is called false sharing as none of the threads are actually sharing any data. Instead it is the spatial locality of the multiple data variables which is causing the cache thrashing.

### **20.13.7 MIPS PDTrace and Performance Counters for SMP**

MIPS Coherent Processing System products have real-time trace facilities for all of the CPU cores as well as the Coherency Manager. Similarly, there are Performance Counters within each CPU core as well as the Coherency Manager.

## Multithreading in the interAptiv Core

Multithreading is the ability of a single core in the multi-core interAptiv Multiprocessing System to execute multiple processes or threads concurrently in an attempt to increase the overall efficiency of the core. The distribution of threads is managed by the operating system, which is responsible for ensuring that all the threads can run at the same time without interfering with each other.

The interAptiv Multiprocessing System implements the MIPS® Multi-Threading (MT) Application Specific Extension (ASE) that provides hardware support for multithreading software applications through the implementation of virtual processing elements (VPE's) and thread contexts (TC's). Through the use of multi-threading, software applications can be executed by an interAptiv core in fewer cycles than on a typical single-threaded core.

The interAptiv Multiprocessing System can be configured with between 1 to 4 cores. Each core to be configured with between 1 to 9 threads.

### 21.1 MT ASE Definitions

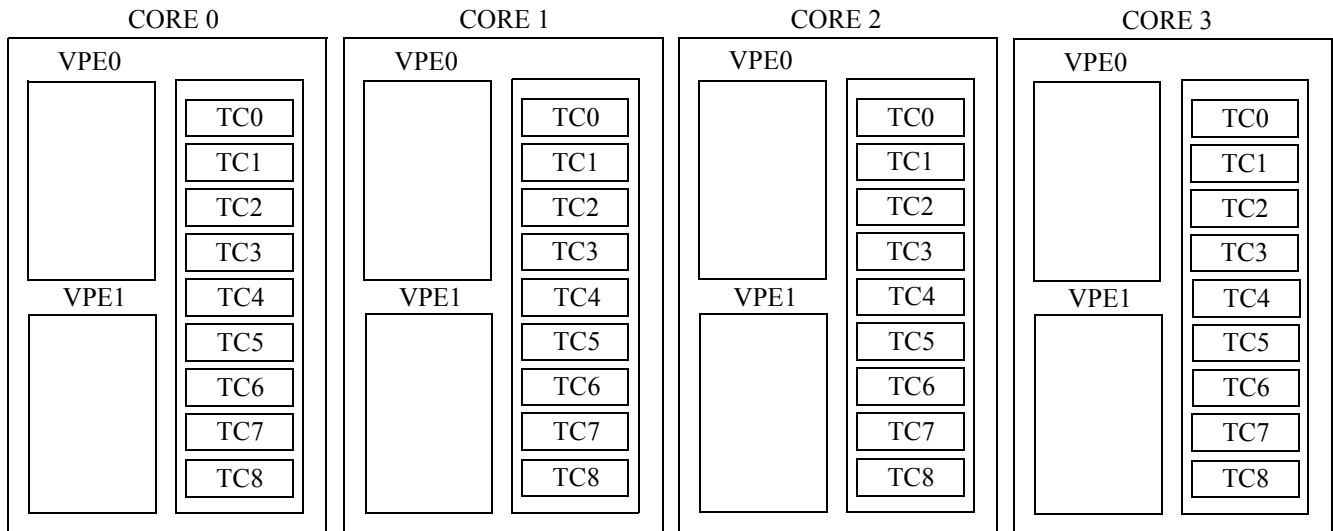
The MIPS MT ASE is an application-specific extension of the MIPS32 instruction set and privileged resource architecture. The MT ASE defines a thread context and a virtual processor as follows:

A **thread context**, or **TC** is a sequential MIPS32 instruction stream that contains the hardware state necessary to support a thread of execution. Each TC includes a set of general purpose registers (GPRs), a program counter (PC), and some multiplier and coprocessor state information. Each interAptiv core can support up to 9 threads.

A **virtual processing element**, or **VPE**, is an instantiation of the full MIPS32 ISA and privileged resource architecture (PRA), sufficient to run a per-processor operating system image. A VPE can be thought of as an “exception domain”, as exception state and priority apply globally within a VPE, and only one exception can be dispatched at a time on a VPE. Each interAptiv core support up to two VPE's.

**Figure 21.1** shows the maximum number of cores, virtual processors, and thread contexts supported by the interAptiv Multiprocessing System.

**Figure 21.1 Maximum Number of VPE's and TC's in the interAptiv Multiprocessing System**



## 21.2 Thread Context Resource Allocation in the interAptiv Core

Resources are allocated at IP configuration time depending on the number of cores, VPE's, and TC's in the system.

The following resources common to the Core level are:

- Pipeline
- Instruction Fetch Unit
- L1 Caches
- Load Store unit
- Multiply-Divide unit
- Arithmetic logic unit
- Memory
- CP0 registers that are shared by all VPE's
- Up to two VPE's per core

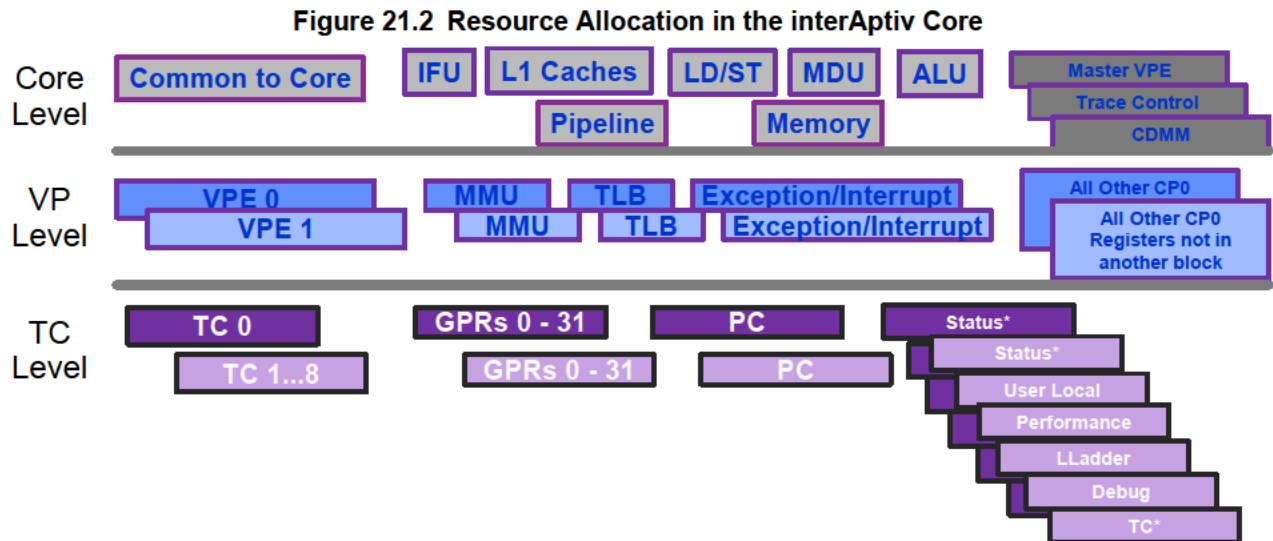
In addition to the core, each VPE has its own:

- MMU
- TLB
- Exception and interrupt logic.
- All CP0 registers that are not Common to the Core are duplicated for each VPE. In addition to the Standard MIPS CP0 registers there are additional registers are defined to be per-VPE, common for all TCs within that VPE.



Each Thread Context (TC) has its own General Purpose Registers and internal Program counter. There are also several CP0 registers in each Thread Context. There can be up to 9 Thread Contexts in the core. Each Thread context is associated with a specific VPE.

Figure 21.2 shows how resources are distributed within the interAptiv Multiprocessing System at the core, VPE, and TC levels.



## 21.3 MT ASE Instructions

The MIPS MT ASE extends both the instruction set and the privileged resources of MIPS32 architecture. These extensions allow an operating system to manipulate the hardware resources associated within a core with multiple VPE's and threads. The ISA extensions allow for efficient user-mode creation and destruction of threads so as to not require OS intervention in the typical cases.

### 21.3.1 Thread Creation

Threads can be directly created in user mode by using the new FORK instruction. The FORK instruction takes three operands. The first is the instruction address at which the new thread will begin execution, the second is an arbitrary register value to be passed to the new thread, typically a pointer to a block of thread-specific storage, and the third, an

output parameter, is the register in the new thread's TC which will receive that second input operand value.

### 21.3.2 Multithreading Software Design Considerations and the FORK Instruction

Software should consider the following design considerations.

#### 21.3.2.1 No Implicit Context Copy

Some high-level multithreading software paradigms require that each new thread receive a copy of the full register set. Others do not. A MIPS32 user mode thread context consists of 31 GPRs (register 0 being always 0), the Hi/Lo or ACX/Hi/Lo accumulator, and some coprocessor state. Copying the full state in a single cycle would require an unrealistically wide interconnect between TCs. Copying sequentially, at a rate of 1-2 registers per cycle, would be complex and time consuming.

The FORK instruction passes a single GPR value parameter to the newly spawned thread. If more than a single value is required for the computation, this value can be a pointer to a context block in memory that will contain register and other values needed by the thread computation.

A FORK also implicitly propagates some privileged state, such as the contents of the ASID register, but the objective is to minimize the “payload” of a FORK operation, in part to minimize the hardware implementation cost, but also in anticipation of multicore “remote” FORKS, where the information would need to be transmitted between processing elements.

### **21.3.2.2 No Value Provided to Forking Thread**

Some high-level multithreading paradigms require that thread creation return a value to the “parent” thread performing the FORK operation. This value is used as a handle or tag for future operations which may reference the thread. The proposed FORK instruction does not do this, for two reasons.

First, it would require a register-file write beyond the write of the GPR value parameter to the register file of the new thread, which creates an undesirable constraint on the design of multithreaded register files.

Second, it creates a name-space problem. In the course of its lifetime, the newly created software thread may end up executing on different register sets of the same CPU, due to context switching, and it may even be migrated to some other processing element. Having hardware provide, as an output of the FORK instruction, a system-unique identifier that would follow each new thread, would be possible, but rather complex to impose on small, embedded cores. Where traceability is required, it can be accomplished using software-based memory interactions.

### **21.3.2.3 Absolute Virtual Thread Starting Address**

The FORK instruction takes as an input operand a register value which is taken to be the starting instruction fetch address for the new thread. Two alternative semantics were considered and rejected. It would have been possible to define FORK such that the new thread simply begins executing at the address following that of the FORK instruction. The two threads could be distinguished by data addressable through the parameter register value passed by the FORK. That would imply that the parent, as well as the child thread, perform the load, evaluation, and control transfer, which was rejected as wasteful.

It would also have been possible to have the FORK instruction contain a “branch” offset in its encoding, rather than specify a register containing a full jump address. This would have made forks to thread code located quite close to the FORK instruction more efficient, but would have penalized all other FORKS by forcing them to perform a double control transfer.

## **21.3.3 Thread Overflow Exception**

Issuing a FORK instruction when there are no free, dynamically allocatable TCs available on the VPE causes a Thread “overflow” exception that can be used by system software to virtualize threads. A program can thus create and use a larger number of software threads than the available complement of TCs, so long as the OS provides higher level scheduling to swap them in and out.

Issuing a FORK instruction when multi-TC scheduling is inhibited on the issuing VPE does not necessarily result in a failure or exception. So long as a TC can be successfully allocated, it is set up to run by the FORK operation, and will begin execution once TC scheduling is enabled.

### 21.3.4 Thread Suspension Using the YIELD Instruction

Thread execution is suspended by the YIELD instruction, which takes as an input operand a descriptor of the circumstances under which the issuing thread should be resumed. If the operand has a value of zero, there are no circumstances under which the thread will resume, and it is de-allocated so that the associated TC may be re-used.

Negative descriptor values are reserved by MIPS for architecturally defined rescheduling conditions. A value of -1 requests that the YIELDing thread be rescheduled without waiting for any specific condition, but allowing other threads to “play through”, according to the implemented thread scheduling scheme. A value of -2 samples the YIELD qualifier inputs to the core without any rescheduling of the thread. Positive descriptor values represent a vector of up to 31 independent “YIELD Qualifier” bits which are hardware inputs to the processor.

The MIPS MT ASE provides mechanisms for an operating system to intercept and emulate YIELD operations. If a per-VPE enable is set, rescheduling YIELDS trap to the operating system with a designated exception code and sub-code identifying a YIELD scheduler intercept. The operating system can evaluate the current YIELD qualifier input state, check it against the contents of the register specified by the trapping YIELD instruction, and make its own determination whether the YQ input values (or some synthetic value) should be placed in the YIELD’s destination register, and the TC restarted at the instruction following the YIELD, or whether the TC contents should be swapped to memory and replaced with the context of another thread of execution.

Each TC has a status bit that is set whenever an instruction is completed for that TC, outside of low-level exception handlers. This bit has multiple software uses, and it is further used by hardware to enable the YIELD scheduler intercept exception. An operating system which wishes to allow a TC to resume and remain blocked on a YIELD after handling a YIELD scheduler intercept exception can clear this “DT” (Dirty Thread) bit before restarting the TC on the YIELD, and that particular TC will remain blocked until the YIELD qualifiers are satisfied, or until some other OS intervention takes place. If the YIELD completes due to the qualifiers being satisfied, the DT bit will be set, and the next blocking YIELD issued by that thread will trap if the YIELD scheduler intercept exceptions are still enabled.

### 21.3.5 Additional MT ASE Instructions

While the previously described FORK and YIELD are the two primitives on which user-mode multithreading is based, the interAptiv MT ASE also includes some privileged instructions to help manage thread and VPE resources.

- **MTTR** is a privileged, “COP0” instruction which moves information from a register of the issuing thread to a register of another thread context on the same processor.
- **MFTR** is a privileged, “COP0” instruction which moves information from a register of another thread context on the same processor to a register of the issuing thread.
- **EMT** is a privileged, “COP0” instruction which atomically enables multithreaded issue on a VPE.
- **EVPE** is a privileged, “COP0” instruction which atomically enables multi-VPE issue on a core with multiple VPE’s enabled.
- **DMT** is a privileged, “COP0” instruction which atomically disables multithreaded issue on a VPE.
- **DVPE** is a privileged, “COP0” instruction which atomically disables multi-VPE issue on a core with multiple VPE’s enabled.

Note that EMT, DMT, EVPE, and DVPE are all instances of the “MFMC0” instruction.

## 21.4 Multithreading CP0 Registers

Some new privileged resources are required to manage the multithreading capabilities of a VPE. Multithreading CP0 registers are instantiated at the core level, the VPE level, and the TC level as shown below.

### 21.4.1 Per-Core Multithreading Registers

- The *MVPControl* register which contains control bits for managing multi-VPE processors.
- The *MVPConf0* and optional *MVPConf1* registers contain information about global multithreaded processor resources which can be configured at boot time and bound to different VPE's.

### 21.4.2 Per-VPE Multithreading Registers

The following registers are instantiated at the VPE level.

- The *VPEControl* register which contains information about the configuration of threads within a VPE.
  - The *VPEConf0* and *VPEConf1* registers contain per-VPE information about the multithreading resources available to the VPE.
  - The *YQMask* register which allows certain YIELD qualifier bits to be masked, so that an attempt to suspend execution pending that state will result in an exception.
  - The *VPESchedule* register allows for the hardware scheduling algorithms of a processor to be manipulated to guarantee some “quality of service” to VPE's with hard real-time requirements.
  - The *VPEScheFBack* register is the counterpart to the *VPESchedule* register, providing per-VPE feedback from the core scheduler logic to system software.
  - The *VPEOpt* register which provides control/status information for optional features, such as run-time cache partitioning.
- The *SRSCConf0* - *SRSCConf4* registers allow for run-time binding of TC's to Shadow Register Sets.

### 21.4.3 Per-TC Multithreading Registers

Seven registers are defined to be per-TC.

- The *TCStatus* register which contains privileged resource information per-thread, such as the Kernel/User state of the thread, or whether it has access to a coprocessor.
- The *TCBind* register which defines a TC's binding to a VPE.
- The *TCRestart* register which contains the restart fetch and execution address of a TC.
- The *TCHalt* register which allows a TC to be put into or taken out of a halted state with a single register write.
- The *TCContext* register which is simply a storage register implemented per-thread, which allows the OS to have instant access to a value, typically a memory pointer such as a kernel stack pointer, that is unique per-thread.
- The *TCSchedule* register allows for the hardware scheduling algorithms of a processor to be manipulated to guarantee some “quality of service” to threads with hard real-time requirements.

- The *TCScheFBack* registers is the counterpart to the *TCSchedule* register, providing per-TC feedback from the core scheduler logic to system software.

## 21.5 Thread Level Exception Processing

By definition, parallelism at the VPE level introduces nothing new in the handling of exceptions for single-threaded VPE's within a multi-VPE core. In the explicit, fine-grained model, however, multiple threads of execution with multiple hardware thread contexts share common system coprocessor resources. This has a number of implications for hardware and software.

Since there is only one Cause register to contain the reason for an exception, a single VPE cannot manage concurrent exceptions. When a synchronous exception is provoked by a thread, as in the case of a TLB miss or a floating-point exception, the MIPS32 architecture stipulates that the EXL or ERL bits of the Status register be set, which block interrupts and further general exceptions from being taken.

In the MIPS MT ASE, the setting of EXL/ERL also prevents the scheduling of other threads until it is cleared by the exception handler. Short exception handling sequences like TLB miss handlers can reasonably be coded, and re-enable multithreading implicitly with the clearing of EXL by the ERET instruction.

More complex exception handling sequences, such as OS system calls, may explicitly re-enable the concurrent execution of non-privileged application threads by clearing EXL once the Cause information has been acquired and saved by the OS.

On a synchronous exception, the TC associated with the instruction stream causing the exception is the one which is associated with the exception: If the exception is not bound to a shadow register set, the associated TC is used to execute the exception handler, and if a shadow register set is used, the associated TC is used as the “previous shadow set”. Asynchronous exceptions, such as interrupts, can be associated with any available activated TC, with the restriction that TCs used by real-time service threads may be designated as exempt from use by interrupt service routines by setting a the IXMT per-thread control bit.

If all activated TCs are explicitly blocked via YIELD instructions or uncompleted loads/stores of gating storage locations, asynchronous exceptions, including Debug exceptions, must be associated with such a blocked TC. The associated handlers will be executed using the previously blocked context, aborting the YIELD or load/store, and the VPE resumes execution on an ERET by re-fetching and re-executing the YIELD or load/store. An aborted gating storage load or store must leave the state of the storage location as it would have been had the instruction never been issued.

EJTAG Debug exceptions are special in several regards with respect to MIPS MT. Like other exceptions, they execute within the context of a specific VPE, but whereas the setting of EXL or ERL by a “normal” exception inhibits thread scheduling only within the affected VPE, Debug mode execution inhibits thread scheduling across all VPE's of a core. And whereas other asynchronous exceptions, such as interrupts, require a TC that is activated and not halted (though it may have been blocked) to process the exception, an asynchronous Debug exception, such as that caused by the assertion of a DINT signal by an EJTAG probe, can be serviced by any TC bound to the targeted VPE, regardless of its halt or activation state. This makes it possible for EJTAG-based debuggers to recover from otherwise completely fatal OS errors, such as halting all TCs.

## 21.6 Fine-Grain Multithreading in the interAptiv Core

Finer-grained multithreading can exploit parallelism at levels that cannot be efficiently addressed by OS-level multithreading. The MT ASE implemented by the interAptiv core allows threads of execution to be created and destroyed very inexpensively by user-mode code. This requires that the applications or underlying libraries be explicitly built or coded to use the new instructions, and also requires the appropriate OS support.

Fine-grained multithreading in the interAptiv core is implemented as an execution model that allows multiple threads to exist within the context of one CPU. Threads share some CPU resources but are able to execute independently. Each thread is scheduled by the policy manager this is a way of controlling the priority of each thread.

### 21.6.1 Dedicated Register Set

Each thread has its own set of general purpose registers. This enables each stage of a multithreaded pipeline to contain instructions from different threads so that execution of those instructions affects only the registers of the thread the instruction is from.

The CPU scheduling through the policy manager takes advantage of stalls in the CPU pipeline such as a cache miss.

### 21.6.2 Automatic Fine-Grained Multithreading

Automatic parallelization algorithms can be employed in compilers to generate multithreaded code. This technique is the ultimate means by which a single C or Fortran program can be accelerated in terms of execution “clock time”.

The necessary compiler techniques exist in the research and high-performance computing communities. They would need to be adapted to MIPS, and used in conjunction with the OS support described above for explicitly fine-grained multithreading.

Operating system support for the fine-grained, FORK/YIELD parallelism of the MT ASE should include:

- Context switch code which dynamically checks the number of threads to be saved and restored each time a user task is switched.
- Fault handling code for Thread exceptions, which occur when there is an underflow or overflow of the number of available physical TCs.
- Allocation and memory management code for ITC storage, if present, as “special” memory.

If threads at runtime without OS intervention are to be able to take nested exceptions, it is anticipated that the Thread-Context register value of each TC is unique. The OS start-up code would assign context storage for each TC on a processor, and insert a pointer to it into that thread’s ThreadContext register prior to that thread’s being made available for FORK allocation.

## 21.7 Operating System Support

To provide the most optimum implementation of multithreading, system software should contain the features listed in the following subsections.

### 21.7.1 “Virtualization” of Threads and Hybrid Scheduling

If more software threads are active in a system than there are TCs available in a MIPS MT VPE, it is necessary to impose a layer of software scheduling on top of the hardware thread scheduling policy of the processor. MIPS MT contains architectural hooks to support this “virtualization” of threads.

Executing a FORK instruction when no dynamically allocatable TCs are free to accept the new instruction stream causes a thread overflow exception, which allows an operating system to detect the case of more software than hardware threads in systems where user-mode thread creation is allowed. If software threads are created only by the OS, the OS will be able to track the available resources without need for an exception.

So long as the number of software threads does not exceed the TC resources available, it is of no consequence from the standpoint of system performance whether a TC remains blocked on a qualified YIELD or a gating storage access, but when TCs are saturated, it becomes necessary to multiplex the software threads across the available TCs.

This can be achieved using simple scheduling algorithms that time-slice threads, regardless of whether or not they are making forward progress, but for high efficiency, it is highly desirable to use blockages as an opportunity to schedule other software threads. The MIPS MT ASE provides the option for a VPE to take an exception whenever a YIELD could cause a rescheduling or whenever a gating storage access blocks.

If blockages will generally be of a short duration, generating exceptions on each blockage may not be desirable, and it may be better to allow TCs to be blocked for some period of time before swapping out their contents. An operating system can do this by periodically sampling and clearing the “dirty” bit associated with each TC, which is set whenever the state of the TC is modified by instruction execution. If the dirty bit remains clear after a sample interval, it may be deduced that the TC has been blocked for the full interval.

## 21.7.2 Software Security

If dynamic FORK/YIELD thread creation and resource allocation is in use simultaneously in different security domains, i.e. by multiple applications or by both an OS and an application, there can be a risk of information “leakage” in the form of register values inherited by an application. It is the responsibility of a secure operating system to manage this risk.

The MIPS MT ASE provides one simple mechanism to facilitate this task; a “dirty” bit associated with each TC, which can be cleared by software and which is set whenever the context is modified. An OS can initialize all TCs to a known “clean” state, and clear all associated dirty bits, prior to scheduling a task. On a task switch, those TCs which are dirty must be “scrubbed” to the clean state before another task can be allowed to allocate and use them.

If a secure operating system wishes to make use of dynamic thread creation and allocation for privileged service threads, the associated TCs must be scrubbed before they are freed for potential use by applications.

Note that the MIPS MT ASE provides no mechanisms to guarantee that two independent, untrusted tasks running concurrently on the same VPE and executing FORK and YIELD instructions, will not exchange TC storage, and thus register values. As such, programs which cannot “trust” one another should be run on distinct VPE’s.

## 21.7.3 Manipulation of Dynamic Allocation Properties of TCs

Each TC has an associated DA bit which makes it available for dynamic allocation by FORK instructions. The interactions of FORK and YIELD with the set of DA bits makes possible several TC management algorithms.

Interrupt-exempt real-time threads may have the DA bit of their associated TC cleared so that a YIELD 0 of the last dynamically allocated thread will cause an underflow Thread exception on the YIELDing thread without interfering with the realtime thread execution, and without leaving the processor in a state where no interrupt-capable TCs are active.

In response to an overflow Thread exception on a FORK, where no more DA TCs are available, the OS can, after having saved a copy of the previous values, clear the DA bits of all TCs, so that the next YIELD 0 will cause an underflow Thread exception which can be used by the OS to restore DA bits and schedule a replay of the failed FORK.

## 21.7.4 Virtual Multiprocessor

Most mainstream operating systems implement some form of symmetric multiprocessing (SMP). Several Microsoft operating systems support SMP platforms, as does Linux. “Multithreaded” applications exist which exploit the parallelism of such platforms, using “heavyweight” threads provided by the operating system. The MIPS MT ASE is designed to provide maximum leverage to this technology.

A multithreaded processor, configured as two single-threaded VPE’s, is indistinguishable to applications software from a 2-way SMP multiprocessor. The operating system would have no need to use any of the new instructions or privileged resources defined by the ASE.

Each MIPS MT TC has its own interrupt “exemption” bit and its own MMU address space identifier (ASID), which allows operating systems to be modified or written to use a “symmetric multi-TC” (SMTC) model, wherein each TC is treated as an independent “processor”. Because multiple TCs may share the privileged resources of a single VPE, an SMTC operating system requires additional logic and complexity to coordinate the use of the shared resources, relative to a standard MIPS32 OS, but the SMTC model allows SMP-like concurrency up to the limit of available TCs.

## 21.7.5 Master / Slave VPE’s

One or more VPE’s on a processor may power-up as a “master” VPE, indicated by the MVP field of the VPConf0 register. A master VPE can access the registers of other VPE’s by using MTTR/MFTR instructions, and can, via the DVPE instruction, suspend all other VPE’s in a processor.

This Master/Slave model allows a multi-tasking master “application processor” VPE running an operating system such as Linux to dispatch real-time processing tasks on another VPE on behalf of various applications. While this could be done using an SMP paradigm, handing work off from one OS to another, MIPS MT also allows this to be done more directly.

A master VPE can take control of another VPE of the same processor at any time. Once a DVPE instruction has been issued by the master VPE, the slave VPE’s CP0 privileged resource state can be set up as needed using MTTR instructions targeting TCs that are bound to the slave VPE, the necessary instructions and data can be set up in memory visible to the slave VPE, one or more TCs of the slave VPE can be initialized using MTTR instructions to set up their TCRstart addresses (and indeed their GPR register values, if appropriate), and the slave VPE can be dispatched to begin execution using the configured TCs by the master VPE executing an EVPE instruction.



## SMP Linux Programming

This chapter describes the Symmetric Multiprocessing (SMP) Linux implementation for the interAptiv Multiprocessing System. The SMP Linux implementation is selected by `CONFIG_MIPS_CPS` and added to the kernel. A kernel with `CONFIG_MIPS_CPS` enabled detects the number of cores and VP's present in the system during boot.

The boot exception vector (BEV) of each core is moved to point to some kernel code rather than the boot loader. Each VP within each core, up to the maximum supported by the kernel (`CONFIG_NR_CPUS` or the *maxcpus* parameter), will then be started by the kernel with no involvement from the boot loader.

MIPS recommends opening the following source files in order to follow along:

- `arch/mips/mti-malta/malta-init.c`
- `arch/mips/mti-malta/malta-int.c`
- `arch/mips/kernel/mips-cpc.c`
- `arch/mips/include/asm/mips-cpc.h`
- `arch/mips/kernel/smp-cps.c`
- `arch/mips/kernel/cps-vec.S`
- `arch/mips/kernel/head.S`

In the above files, VP's are referred to by the notation `cXvY`, where `X` is the zero-based index of the core which the VP is a part of, and `Y` is the zero-based index of the VP within that core. For example, `c0v0` refers to the first VPE of the first core, `c0v1` refers to the second VPE of the first core, `c1v0` refers to the first VPE of the second core, etc.

The interAptiv SMP Linux implementation treats each VP as a logical CPU. For example, in a system with two VP's in each of two cores, Linux will see four CPUs. Linux CPU numbers begin with `0 0`, representing `c0v0` and progressing through the VP's of core 0, then through the VP's of core 1 etc. For example in a system with two VP's in each of two cores, the various CPU's would be referred to as follows:

- `cpu0` refers to `c0v0` (core 0, VP 0)
- `cpu1` refers to `c0v1` (core 0, VP 1)
- `cpu2` refers to `c1v0` (core 1, VP 0)
- `cpu3` refers to `c1v1` (core 1, VP 1)

## 22.1 Overview of Boot Sequence

The following steps denote the boot sequence.

Early during boot the platform code will probe for the CM. In kernel versions  $\geq 3.15$  this is done by calling `mips_cm_probe`. In older kernels this is done by calling `gcmp_probe`.

1. After finding the CM2, the platform code probes for the CPC. This is done by calling `mips_cpc_probe`. As the CPC base address is programmable the platform code must implement a `mips_cpc_default_phys_base` function which returns the address to map the CPC registers to if the CPC is not already enabled.
2. The platform probes for the GIC, likely after finding the CM2. This is currently a platform-specific step.
3. Platform code calls `register_cps_smp_ops` to initialise the SMP implementation. This will return zero on success or non-zero if an error occurred (eg. CONFIG\_MIPS\_CPS is disabled, or the CM or CPC probing was unsuccessful).
4. The `cps_smp_setup` function is called (via `start_kernel -> setup_arch -> plat_smp_setup`). This detects the cores and VP's present within the system by inspecting the CM2 registers. It calls `init_core` for core 0, enters core 0 into the coherent domain and performs some initialization functions.
5. The `cps_prepare_cpus` function is called (via `kernel_init -> kernel_init_freeable -> smp_prepare_cpus`). This currently only performs some initialization relevant to debugging MT cores and the ITC feature of the interAptiv core (see `mips_mt_set_cpuoptions`).
6. The `cps_boot_secondary` function is called, running on c0v0, for each secondary CPU (VP) to be started. If the VP to be started is in a powered down core then the `boot_core` function is called. Otherwise `boot_vpe` is called on a VP that has already been started within the same core. If the VP to be started is in a core which is not core 0 and is already powered up, then this means that an Inter-Processor Interrupt (IPI) will be sent to a VP within that core during this step.
7. If `boot_core` was called due to the core being powered down, then the BEV of the core is moved to point to `mips_cps_core_entry` within uncached (kseg1) space. The `mips_cps_bootcfg` structure is filled in with the appropriate pc, sp, and gp registers for the newly started core to proceed with execution once it is running.

The program counter points to the `smp_bootstrap` function. The core is then reset via the CPC or the CM (if using CM1) and begins to execute from `mips_cps_core_entry`. This code checks whether it is running due to an NMI, which is not the case. The code then initializes some basic coprocessor 0 state and the L1 caches of the core, enters the coherent domain, transitions to the run state from cached (kseg0) space, and loads the pc, sp, and gp registers from the `mips_cps_bootcfg` structure. Finally, the code branches to that loaded program counter and continue executing from `smp_bootstrap`.

8. If the `boot_vpe` function is called on a powered up core, then the registers of the TC associated with the VP to be started are set appropriately to execute generic kernel code. They take the same values as the `mips_cps_bootcfg` structure would in the `boot_core` case. The VP is then activated and allowed to execute, and proceeds into `smp_bootstrap`.
9. The `smp_bootstrap` function is a generic MIPS kernel function which is run on each secondary VP once it has started executing. It simply initializes some coprocessor 0 state such as the `Status` register, then branches to the `start_secondary` function.

10. The *start\_secondary* function performs setup generic to the MIPS architecture such as probing for the type of CPU and the features it implements. It calls *cps\_init\_secondary* fairly early. This function calls *init\_core* if it is the first VP executing upon that core, and set the CP0 *Status* register appropriately to later allow interrupts to be processed.

## 22.2 Inter-Processor Interrupts

Inter-Processor Interrupts (IPIs) are sent via the General Interrupt Controller (GIC) using interrupts which are not connected to external hardware. There must be two such interrupts for each VP within the system. One interrupt per VP is used to signal that the VP should reschedule to run another task. The other is used to signal that the VP should execute a function, which is used in many places throughout the kernel but notably to execute cache and TLB maintenance functions on the appropriate VPEs.

The assignment of GIC interrupts to be used for IPIs is currently entirely platform specific. On the Malta board, the interrupts  $(GIC\_NUM\_INTRS - (NR\_CPUS * 2))..(GIC\_NUM\_INTRS - 1)$  are used. See the *gic\_call\_int\_base* and *gic\_resched\_int\_base* variables, plus the *fill\_ipi\_map* function in *arch/mips/mti-malta/maltaint.c*.

The SMP code in *arch/mips/kernel/smp-gic.c* calls the platform-specific functions *plat\_ipi\_call\_int\_xlate* and *plat\_ipi\_resched\_int\_xlate* in order to map from a CPU number to the GIC interrupt to be used for an IPI, and then triggers that interrupt using the write edge register of the GIC.

## 22.3 Clock Source and Events

The SMP Linux implementation provides no restrictions upon what is used as a clock source or for clock events. As always the clock source should ideally be a global counter, which the GIC clock source (*CONFIG\_CSRC\_GIC*) can provide. Clock events would typically also be provided by the GIC (*CONFIG\_CEVT\_GIC*).

## 22.4 Platform Setup

The following steps should be undertaken to make use of the SMP Linux implementation for your platform:

1. Ensure that the bootloader sets the cache coherency attribute in the *Config.K0* register for c0v0 to something coherent during boot. Currently the secondary cores have this set to CWB (cacheable, coherent, write-back, write-allocate, read misses request Shared) or 0x5. It is recommended that the same be set for c0v0 by the boot loader.
2. Probe the CM2 early during platform initialization by calling *mips\_cm\_probe* on kernel version revision 3.15 and above, or *gcmp\_probe* on older kernels.
3. Probe the Cluster Power Controller (CPC) after the CM2 by calling *mips\_cpc\_probe*. Provide a *mips\_cpc\_default\_phys\_base* function which returns the physical address at which the CPC registers should be mapped.
4. Initialise the GIC. This is currently largely platform-specific but at a minimum involves generating an array of struct *gic\_intr\_map* to configure GIC interrupts and map them to VP's, calling *gic\_init* and providing a *gic\_platform\_init* function. The CPU interrupts which are connected to the GIC also need to be configured appropriately to a function which calls *gic\_get\_int* to demux the pending GIC interrupt and process it accordingly. See *malta\_ipi\_irqdispatch* for an example.

5. Provide *plat\_ipi\_call\_int\_xlate* and *plat\_ipi\_resched\_int\_xlate* functions to map from CPU numbers to the GIC interrupts to be used for IPIs. These interrupts should have been routed to the appropriate VP's through the array of struct *gic\_intr\_map* earlier passed to *gic\_init*.
6. Call *register\_cps\_smp\_ops* to register SMP as the SMP implementation to be used for your platform. \
7. Select the `SYS_SUPPORTS_MIPS_CPS` Kconfig option for your platform. This will allow selection of the `MIPS_CPS` Kconfig option.
8. Enable the `MIPS_CPS` Kconfig option (Kernel type -> MIPS Coherent Processing System support).

## Instruction Set Overview

This chapter provides an overview of the interAptiv™ core instruction set, including the instruction formats and the basic instruction types.

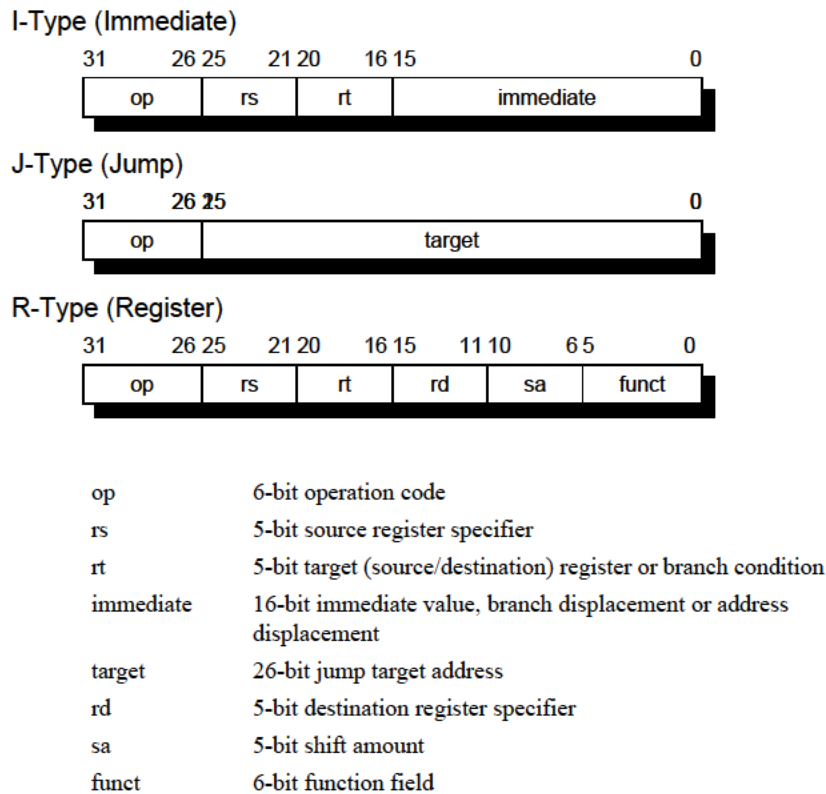
This chapter discusses the following topics:

- [Section 23.1 “CPU Instruction Formats”](#)
- [Section 23.2 “Load and Store Instructions”](#)
- [Section 23.3 “Computational Instructions”](#)
- [Section 23.4 “Jump and Branch Instructions”](#)
- [Section 23.5 “Control Instructions”](#)
- [Section 23.6 “Coprocessor Instructions”](#)
- [Section 23.7 “New Instructions for the interAptiv™ Core”](#)
- [Section 23.8 “Base Instruction Set for the interAptiv™ Core”](#)

### 23.1 CPU Instruction Formats

A CPU instruction consists of a single 32-bit word, aligned on a word boundary. There are three instruction formats: *immediate* (*I*-type), *jump* (*J*-type), and *register* (*R*-type). The use of a small number of instruction formats simplifies instruction decoding, allowing the compiler to synthesize more complicated (and less frequently used) operations and addressing modes from these three formats as needed. The instruction formats are shown in [Figure 23.1](#).

**Figure 23.1 Instruction Formats**



## 23.2 Load and Store Instructions

Load and store instructions are immediate (I-type) instructions that move data between memory and the general registers. The only addressing mode that integer load and store instructions directly support is *base register plus 16-bit signed immediate offset*. Floating point load and store instructions can use either that addressing mode or *register plus register* indexed addressing.

### 23.2.1 Scheduling a Load Delay Slot

A load instruction that does not allow its result to be used by the instruction immediately following is called a *delayed load instruction*. The instruction slot immediately following this delayed load instruction is referred to as the *load delay slot*.

In an interAptiv core, the instruction immediately following a load instruction can use the contents of the loaded register; however in such cases hardware interlocks insert additional real cycles. Although not required, the scheduling of load delay slots can be desirable, both for performance and R-Series processor compatibility.

### 23.2.2 Defining Access Types

*Access type* indicates the size of a core data item to be loaded or stored, set by the load or store instruction opcode.

Regardless of access type or byte ordering (endianness), the address given specifies the low-order byte in memory. For a big-endian configuration, the low-order byte is the most-significant byte; for a little-endian configuration, the low-order byte is the least-significant byte.

The access type, together with the three low-order bits of the address, defines the bytes accessed within the addressed word, as shown in Table 23.1. Only the combinations shown in Table 23.1 are permissible; other combinations cause address-error exceptions.

Instruction fetches are either halfword accesses (MIPS16e™ code) or word accesses (32b code). These references will be impacted by endianness in the same way as load/store references of those sizes.

**Table 23.1 Byte Access Within a Doubleword**

Access Type	Low-Order Address Bits			Bytes Accessed															
				Big Endian (63-----31-----0)								Little Endian (63-----31-----0)							
	2	1	0	Byte								Byte							
Doubleword	0	0	0	0	1	2	3	4	5	6	7	7	6	5	4	3	2	1	0
Word	0	0	0	0	1	2	3									3	2	1	0
	1	0	0					4	5	6	7	7	6	5	4				
Triplebyte	0	0	0	0	1	2											2	1	0
	0	0	1		1	2	3									3	2	1	
	1	0	0					4	5	6			6	5	4				
	1	0	1						5	6	7	7	6	5					
Halfword	0	0	0	0	1													1	0
	0	1	0			2	3									3	2		
	1	0	0					4	5					5	4				
	1	1	0							6	7	7	6						
Byte	0	0	0	0															0
	0	0	1		1														1
	0	1	0			2											2		
	0	1	1				3									3			
	1	0	0					4							4				
	1	0	1						5					5					
	1	1	0							6			6						
	1	1	1								7	7							

### 23.3 Computational Instructions

Computational instructions can be either in register (R-type) format, in which both operands are registers, or in immediate (I-type) format, in which one operand is a 16-bit immediate.

Computational instructions perform the following operations on register values:

- Arithmetic
- Logical
- Shift
- Count Leading Zeros/Ones
- Multiply
- Divide

### 23.3.1 Cycle Timing for Multiply and Divide Instructions

Multiply instruction in the integer pipeline are transferred to the multiplier while remaining instructions continue through the pipeline; the product of the multiply instruction is saved in the HI and LO registers. If the multiply instruction is followed by an MFHI or MFLO before the product is available, the pipeline interlocks until this product does become available.

## 23.4 Jump and Branch Instructions

Jump and branch instructions change the control flow of a program. All MIPS32 R3 jump and branch instructions occur with a delay of one instruction: that is, the instruction immediately following the jump or branch (the instruction in the so-called *delay slot*) always executes while the target instruction is being fetched from storage.

### 23.4.1 Overview of Jump Instructions

Subroutine calls in high-level languages are usually implemented with Jump (J) or Jump and Link (JAL) instructions, both of which have the J-type format. In J-type format, the 26-bit target address shifts left 2 bits and combines with the high-order 4 bits of the current program counter to form an absolute address.

Returns and large cross-page jumps are usually implemented with the Jump Register (JR) or Jump and Link Register (JALR) instructions. Both are R-type instructions that use the 32-bit byte address contained in one of the general purpose registers.

For more information about jump instructions, refer to the individual instructions in *MIPS32® Architecture Reference Manual, Volume II: The MIPS32® Instruction Set*.

### 23.4.2 Overview of Branch Instructions

All branch instruction target addresses are computed by adding the address of the instruction in the delay slot to the 16-bit *offset* (shifted left 2 bits and sign-extended to 32 bits). All MIPS32 R3 branches occur with a delay of one instruction.

If a conditional branch likely is not taken, the instruction in the delay slot is nullified.

Branches, jumps, ERET, and DERET instructions should not be placed in the delay slot of a branch or jump.



## 23.5 Control Instructions

Control instructions allow the software to initiate traps; they are always R-type.

## 23.6 Coprocessor Instructions

CP0 instructions perform operations on the System Control Coprocessor registers to manipulate the memory management and exception handling facilities of the processor.

CP1 instructions perform operations on the floating point unit (FPU).

CP2 instructions are not implemented in interAptiv core.

## 23.7 New Instructions for the interAptiv™ Core

This section describes the new instructions added to the MIPS32 Release 3 architecture. [Table 23.2](#) lists the instructions in alphabetical order. Additional instruction information in the *MIPS32® Architecture Reference Manual* is not duplicated here.

**Table 23.2 interAptiv™ CPU Instruction Set**

Instruction	Description	Function
CACHEE	Cache Instruction EVA	See the description of the CACHE/CACHEE instruction in Chapter 19.
LBE	Load Byte EVA	Same as the LB instruction, except that it allows kernel address to user addresses. Refer to the LB instruction in <a href="#">Table 23.3</a> for functional information.
LBUE	Load Byte Unsigned EVA	Same as the LBU instruction, except that it allows kernel address to user addresses. Refer to the LBU instruction <a href="#">Table 23.3</a> for functional information.
LHE	Load Halfword EVA	Same as the LH instruction, except that it allows kernel address to user addresses. Refer to the LH instruction <a href="#">Table 23.3</a> for functional information.
LHUE	Load Halfword Unsigned EVA	Same as the LB instruction, except that it allows kernel address to user addresses. Refer to the LH instruction <a href="#">Table 23.3</a> for functional information.
LLE	Load Linked EVA	Same as the LL instruction, except that it allows kernel address to user addresses. Refer to the LL instruction <a href="#">Table 23.3</a> for functional information.
LWE	Load Word EVA	Same as the LW instruction, except that it allows kernel address to user addresses. Refer to the LW instruction in <a href="#">Table 23.3</a> for functional information.
LWLE	Load Word Left EVA	Same as the LWL instruction, except that it allows kernel address to user addresses. Refer to the LWL instruction in <a href="#">Table 23.3</a> for functional information.
LWRE	Load Word Right EVA	Same as the LWR instruction, except that it allows kernel address to user addresses. Refer to the LWR instruction in <a href="#">Table 23.3</a> for functional information.

**Table 23.2 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
PREFE	Prefetch EVA	Load Specified Line into Cache. See also the description of the PREF / PREFE instruction in Chapter 19.
SBE	Store Byte EVA	Same as the SB instruction, except that it allows kernel address to user addresses. Refer to the SB instruction in Table 23.3 for functional information.
SCE	Store Conditional EVA	Same as the SC instruction, except that it allows kernel address to user addresses. Refer to the SC instruction in Table 23.3 for functional information.
SHE	Store Halfword EVA	Same as the SH instruction, except that it allows kernel address to user addresses. Refer to the SH instruction in Table 23.3 for functional information.
SWE	Store Word EVA	Same as the SW instruction, except that it allows kernel address to user addresses. Refer to the SW instruction in Table 23.3 for functional information.
SWLE	Store Word Left EVA	Same as the SWL instruction, except that it allows kernel address to user addresses. Refer to the SWL instruction in Table 23.3 for functional information.
SWRE	Store Word Right EVA	Same as the SWR instruction, except that it allows kernel address to user addresses. Refer to the SWR instruction in Table 23.3 for functional information.
TLBINV	TLB Invalidate	TLBINV invalidates a set of TLB entries based on ASID and <i>Index</i> match. On execution of the TLBINV instruction, the set of TLB entries with matching ASID are marked invalid, excluding those TLB entries which have their G bit set to 1. For more information, refer to the TLBINV instruction in Chapter 19.
TLBINVF	TLB Invalidate Flush	TLBINV invalidates a set of TLB entries based on ASID and <i>Index</i> match. On execution of the TLBINVF instruction, all entries within range of <i>Index</i> are invalidated. For more information, refer to the TLBINVF instruction in Chapter 19.

## 23.8 Base Instruction Set for the interAptiv™ Core

This section describes the base instructions for the MIPS32 Release 3 architecture. [Table 23.3](#) lists the instructions in alphabetical order. Following the table, the instructions that have implementation-dependent behavior in the interAptiv core are described individually. The descriptions of other instructions that exist in the *MIPS32® Architecture Reference Manual* are not duplicated here.

Refer to *Volume II* of the *MIPS32® Architecture Reference Manual* for more information about the instruction descriptions. That document contains a description of the instruction fields, a definition of terms, and a description function notation.

**Table 23.3 interAptiv™ CPU Instruction Set**

Instruction	Description	Function
ADD	Integer Add	$Rd = Rs + Rt$
ADDI	Integer Add Immediate	$Rt = Rs + Immed$
ADDIU	Unsigned Integer Add Immediate	$Rt = Rs +_U Immed$
ADDIUPC	Unsigned Integer Add Immediate to PC (MIPS16 only)	$Rt = PC +_U Immed$
ADDU	Unsigned Integer Add	$Rd = Rs +_U Rt$
AND	Logical AND	$Rd = Rs \& Rt$
ANDI	Logical AND Immediate	$Rt = Rs \& (0_{16}    Immed)$
B	Unconditional Branch (Assembler idiom for: BEQ r0, r0, offset)	$PC += (int)offset$
BAL	Branch and Link (Assembler idiom for: BGEZAL r0, offset)	$GPR[31] = PC + 8$ $PC += (int)offset$
BEQ	Branch On Equal	if $Rs == Rt$ $PC += (int)offset$
BEQL	Branch On Equal Likely	if $Rs == Rt$ $PC += (int)offset$ else Ignore Next Instruction
BGEZ	Branch on Greater Than or Equal To Zero	if $!Rs[31]$ $PC += (int)offset$
BGEZAL	Branch on Greater Than or Equal To Zero And Link	$GPR[31] = PC + 8$ if $!Rs[31]$ $PC += (int)offset$
BGEZALL	Branch on Greater Than or Equal To Zero And Link Likely	$GPR[31] = PC + 8$ if $!Rs[31]$ $PC += (int)offset$ else Ignore Next Instruction
BGEZL	Branch on Greater Than or Equal To Zero Likely	if $!Rs[31]$ $PC += (int)offset$ else Ignore Next Instruction

**Table 23.3 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
BGTZ	Branch on Greater Than Zero	if !Rs[31] && Rs != 0 PC += (int)offset
BGTZL	Branch on Greater Than Zero Likely	if !Rs[31] && Rs != 0 PC += (int)offset else Ignore Next Instruction
BLEZ	Branch on Less Than or Equal to Zero	if Rs[31]    Rs == 0 PC += (int)offset
BLEZL	Branch on Less Than or Equal to Zero Likely	if Rs[31]    Rs == 0 PC += (int)offset else Ignore Next Instruction
BLTZ	Branch on Less Than Zero	if Rs[31] PC += (int)offset
BLTZAL	Branch on Less Than Zero And Link	GPR[31] = PC + 8 if Rs[31] PC += (int)offset
BLTZALL	Branch on Less Than Zero And Link Likely	GPR[31] = PC + 8 if Rs[31] PC += (int)offset else Ignore Next Instruction
BLTZL	Branch on Less Than Zero Likely	if Rs[31] PC += (int)offset else Ignore Next Instruction
BNE	Branch on Not Equal	if Rs != Rt PC += (int)offset
BNEL	Branch on Not Equal Likely	if Rs != Rt PC += (int)offset else Ignore Next Instruction
BREAK	Breakpoint	Break Exception
CACHE	Cache Operation	See the description of the CACHE instruction. Refer to the Caches chapter for more information.
CLO	Count Leading Ones	Rd = NumLeadingOnes(Rs)
CLZ	Count Leading Zeroes	Rd = NumLeadingZeroes(Rs)
COPYW	Copy words between memory locations (MIPS16 only)	See Section 24.4.
COPO	Coprocessor 0 Operation	See Software User's Manual
DERET	Return from Debug Exception	PC = DEPC Exit Debug Mode
DI	Atomically Disable Interrupts	Rt = Status; Status <sub>IE</sub> = 0

**Table 23.3 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
DIV	Divide	LO = (int)Rs / (int)Rt HI = (int)Rs % (int)Rt
DIVU	Unsigned Divide	LO = (uns)Rs / (uns)Rt HI = (uns)Rs % (uns)Rt
EHB	Execution Hazard Barrier	Stop instruction execution until execution hazards are cleared
EI	Atomically Enable Interrupts	Rt = Status; Status <sub>IE</sub> = 1
ERET	Return from Exception	if SR[2] PC = ErrorEPC else PC = EPC SR[1] = 0 SR[2] = 0 LL = 0
EXT	Extract Bit Field	Rt = ExtractField(Rs, pos, size)
INS	Insert Bit Field	Rt = InsertField(Rs, Rt, pos, size)
J	Unconditional Jump	PC = PC[31:28]    offset<<2
JAL	Jump and Link	GPR[31] = PC + 8 PC = PC[31:28]    offset<<2
JALR	Jump and Link Register	Rd = PC + 8 PC = Rs
JALR.HB	Jump and Link Register with Hazard Barrier	Like JALR, but also clears execution and instruction hazards
JALRC	Jump and Link Register Compact - do not execute instruction in jump delay slot (MIPS16 only)	Rd = PC + 2 PC = Rs
JR	Jump Register	PC = Rs
JR.HB	Jump Register with Hazard Barrier	Like JR, but also clears execution and instruction hazards
JRC	Jump Register Compact - do not execute instruction in jump delay slot (MIPS16 only)	PC = Rs
LB	Load Byte	Rt = (byte)Mem[base+offset]
LBU	Unsigned Load Byte	Rt = (ubyte)Mem[base+offset]
LH	Load Halfword	Rt = (half)Mem[base+offset]
LHU	Unsigned Load Halfword	Rt = (uhalf)Mem[base+offset]
LL	Load Linked Word	Rt = Mem[base+offset] LL = 1 See also the description of the LL instruction on <a href="#">page 893</a> .
LUI	Load Upper Immediate	Rt = immediate << 16
LW	Load Word	Rt = Mem[Rs+offset]

**Table 23.3 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
LWPC	Load Word, PC relative	$Rt = Mem[PC+offset]$
LWL	Load Word Left	See Architecture Reference Manual
LWR	Load Word Right	See Architecture Reference Manual
MADD	Multiply-Add	$HI   LO += (int)Rs * (int)Rt$
MADDU	Multiply-Add Unsigned	$HI   LO += (uns)Rs * (uns)Rt$
MFC0	Move From Coprocessor 0	$Rt = CPR[0, Rd, sel]$
MFHI	Move From HI	$Rd = HI$
MFLO	Move From LO	$Rd = LO$
MOVN	GPR Conditional Move on Not Zero	if $Rt \neq 0$ then $Rd = Rs$
MOVZ	GPR Conditional Move on Zero	if $Rt = 0$ then $Rd = Rs$
MSUB	Multiply-Subtract	$HI   LO -= (int)Rs * (int)Rt$
MSUBU	Multiply-Subtract Unsigned	$HI   LO -= (uns)Rs * (uns)Rt$
MTC0	Move To Coprocessor 0	$CPR[0, n, Sel] = Rt$
MTHI	Move To HI	$HI = Rs$
MTLO	Move To LO	$LO = Rs$
MUL	Multiply with register write	$HI   LO = Unpredictable$ $Rd = ((int)Rs * (int)Rt)_{31..0}$
MULT	Integer Multiply	$HI   LO = (int)Rs * (int)Rd$
MULTU	Unsigned Multiply	$HI   LO = (uns)Rs * (uns)Rd$
NOP	No Operation (Assembler idiom for: SLL r0, r0, r0)	
NOR	Logical NOR	$Rd = \sim(Rs   Rt)$
OR	Logical OR	$Rd = Rs   Rt$
ORI	Logical OR Immediate	$Rt = Rs   Immed$
PREF	Prefetch	Load Specified Line into Cache. See also the description of the PREF instruction on <a href="#">page 895</a> .
RDHWR	Read Hardware Register	Allows unprivileged access to registers enabled by HWREna register
RDPGPR	Read GPR from Previous Shadow Set	$Rt = SGPR[SRSctl_{pss}, Rd]$
RESTORE	Restore registers and deallocate stack frame	See RESTORE instruction in Chapter 25.
ROTR	Rotate Word Right	$Rd = Rt_{sa-1..0}    Rt_{31..sa}$

**Table 23.3 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
ROTRV	Rotate Word Right Variable	$Rd = Rt_{Rs-1..0} \parallel Rt_{31..Rs}$
SAVE	Save registers and allocate stack frame	See SAVE instruction in Chapter 25.
SB	Store Byte	$(byte)Mem[base+offset] = Rt$
SC	Store Conditional Word	<pre>if LL = 1     mem[base+offset] = Rt Rt = LL</pre> <p>See also the description of the SC instruction on <a href="#">page 898</a>.</p>
SDBBP	Software Debug Break Point	Trap to SW Debug Handler
SEB	Sign Extend Byte	$Rd = (byte)Rs$
SEH	Sign Extend Half	$Rd = (half)Rs$
SH	Store Half	$(half)Mem[base+offset] = Rt$
SLL	Shift Left Logical	$Rd = Rt \ll sa$
SLLV	Shift Left Logical Variable	$Rd = Rt \ll Rs[4:0]$
SLT	Set on Less Than	<pre>if (int)Rs &lt; (int)Rt     Rd = 1 else     Rd = 0</pre>
SLTI	Set on Less Than Immediate	<pre>if (int)Rs &lt; (int)Immed     Rt = 1 else     Rt = 0</pre>
SLTIU	Set on Less Than Immediate Unsigned	<pre>if (uns)Rs &lt; (uns)Immed     Rt = 1 else     Rt = 0</pre>
SLTU	Set on Less Than Unsigned	<pre>if (uns)Rs &lt; (uns)Immed     Rd = 1 else     Rd = 0</pre>
SRA	Shift Right Arithmetic	$Rd = (int)Rt \gg sa$
SRAV	Shift Right Arithmetic Variable	$Rd = (int)Rt \gg Rs[4:0]$
SRL	Shift Right Logical	$Rd = (uns)Rt \gg sa$
SRLV	Shift Right Logical Variable	$Rd = (uns)Rt \gg Rs[4:0]$
SSNOP	Superscalar Inhibit No Operation	NOP
SUB	Integer Subtract	$Rt = (int)Rs - (int)Rd$
SUBU	Unsigned Subtract	$Rt = (uns)Rs - (uns)Rd$
SW	Store Word	$Mem[base+offset] = Rt$
SWL	Store Word Left	See Architecture Reference Manual

**Table 23.3 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
SWR	Store Word Right	See Architecture Reference Manual
SYNC	Synchronize	See the description of the SYNC instruction on <a href="#">page 887</a> .
SYNCI	Synchronize Caches to Make Instruction Writes Effective	For D-cache writeback and I-cache invalidate on specified address
SYSCALL	System Call	SystemCallException
TEQ	Trap if Equal	if Rs == Rt TrapException
TEQI	Trap if Equal Immediate	if Rs == (int)Immed TrapException
TGE	Trap if Greater Than or Equal	if (int)Rs >= (int)Rt TrapException
TGEI	Trap if Greater Than or Equal Immediate	if (int)Rs >= (int)Immed TrapException
TGEIU	Trap if Greater Than or Equal Immediate Unsigned	if (uns)Rs >= (uns)Immed TrapException
TGEU	Trap if Greater Than or Equal Unsigned	if (uns)Rs >= (uns)Rt TrapException
TLBWI	Write Indexed TLB Entry	See the description of the TLBWI instruction on <a href="#">page 906</a> .
TLBWR	Write Random TLB Entry	See the description of the TLBWR instruction on <a href="#">page 904</a> .
TLBP	Probe TLB for Matching Entry	See Software Users Manual
TLBR	Read Index for TLB Entry	See the description of the TLBR instruction on <a href="#">page 905</a> .
TLT	Trap if Less Than	if (int)Rs < (int)Rt TrapException
TLTI	Trap if Less Than Immediate	if (int)Rs < (int)Immed TrapException
TLTIU	Trap if Less Than Immediate Unsigned	if (uns)Rs < (uns)Immed TrapException
TLTU	Trap if Less Than Unsigned	if (uns)Rs < (uns)Rt TrapException
TNE	Trap if Not Equal	if Rs != Rt TrapException
TNEI	Trap if Not Equal Immediate	if Rs != (int)Immed TrapException
UCOPYW	Copy unaligned words between memory locations	See Section 24.4.
WAIT	Wait for Interrupts	Stall until interrupt occurs. See the description of the WAIT instruction on <a href="#">page 910</a> .



**Table 23.3 interAptiv™ CPU Instruction Set (continued)**

Instruction	Description	Function
WRPGPR	Write to GPR in Previous Shadow Set	$SGPR[SRSCtl_{pss}, Rd] = Rt$
WSBH	Word Swap Bytes Within HalfWords	$Rd = Rt_{23..16}    Rt_{31..24}    Rt_{7..0}    Rt_{15..8}$
XOR	Exclusive OR	$Rd = Rs \wedge Rt$
XORI	Exclusive OR Immediate	$Rt = Rs \wedge (uns) Immed$
ZEB	Zero extend byte (MIPS16 only)	$Rt = (ubyte) Rs$
ZEH	Zero extend half (MIPS16 only)	$Rt = (uhalf) Rs$



# MIPS16e Application-Specific Extension to the MIPS32® Instruction Set

This chapter describes the MIPS16e™ ASE as implemented in the interAptiv core. Refer to *Volume IV-a* of the *MIPS32® Architecture Reference Manual* for a general description of the MIPS16e ASE and detailed descriptions of the instructions.

The MIPS16e2 ASE includes additional instructions not accessible in the previous generation to enhance code compression and performance. Refer to the document entitled *MIPS16e2 Application Specific Extension Technical Reference Manual* for a listing and definition of the new MIPS16e2 instructions.

This chapter covers the following topics:

- [Section 24.1 “Instruction Bit Encoding”](#)
- [Section 24.2 “MIPS16e and MIPS16e2 Instruction Listing”](#)

## 24.1 Instruction Bit Encoding

[Table 24.2](#) through [Table 24.9](#) describe the encoding used for the MIPS16e ASE. [Table 24.1](#) describes the meaning of the symbols used in the tables.

**Table 24.1 Symbols Used in the Instruction Encoding Tables**

*	Operation or field codes marked with this symbol are reserved for future use. Executing such an instruction causes a Reserved Instruction Exception.
$\delta$	(Also <i>italic</i> field name.) Operation or field codes marked with this symbol denote a field class. The instruction word must be further decoded by examining additional tables that show values for another instruction field.
$\beta$	Operation or field codes marked with this symbol represent a valid encoding for a higher-order MIPS ISA level. Executing such an instruction causes a Reserved Instruction Exception.
$\theta$	Operation or field codes marked with this symbol are available to licensed MIPS partners. To avoid multiple conflicting instruction definitions, the partner must notify MIPS Technologies, Inc. when one of these encodings is used. If no instruction is encoded with this value, executing such an instruction must cause a Reserved Instruction Exception ( <i>SPECIAL2</i> encodings or coprocessor instruction encodings for a coprocessor to which access is allowed) or a Coprocessor Unusable Exception (coprocessor instruction encodings for a coprocessor to which access is not allowed).
$\sigma$	Field codes marked with this symbol represent an EJTAG support instruction and implementation of this encoding is optional for each implementation. If the encoding is not implemented, executing such an instruction must cause a Reserved Instruction Exception. If the encoding is implemented, it must match the instruction encoding as shown in the table.
$\varepsilon$	Operation or field codes marked with this symbol are reserved for MIPS Application Specific Extensions. If the ASE is not implemented, executing such an instruction must cause a Reserved Instruction Exception.

**Table 24.1 Symbols Used in the Instruction Encoding Tables(continued)**

*	Operation or field codes marked with this symbol are reserved for future use. Executing such an instruction causes a Reserved Instruction Exception.
$\phi$	Operation or field codes marked with this symbol are obsolete and will be removed from a future revision of the MIPS64 ISA. Software should avoid using these operation or field codes.

**Table 24.2 MIPS16e Encoding of the Opcode Field**

<b>opcode</b>		bits 13..11							
		0	1	2	3	4	5	6	7
bits 15..14		000	001	010	011	100	101	110	111
0	00	ADDIUSP <sup>1</sup>	ADDIUPC <sup>2</sup>	B	<i>JAL(X)</i> $\delta$	BEQZ	BNEZ	<i>SHIFT</i> $\delta$	$\beta$
1	01	<i>RRI-A</i> $\delta$	ADDIU <sup>3</sup>	SLTI	SLTIU	<i>I8</i> $\delta$	LI	CMPI	$\beta$
2	10	LB	LH	LWSP <sup>4</sup>	LW	LBU	LHU	LWPC <sup>5</sup>	$\beta$
3	11	SB	SH	SWSP <sup>6</sup>	SW	<i>RRR</i> $\delta$	<i>RR</i> $\delta$	<i>EXTEND</i> $\delta$	$\beta$

1. The ADDIUSP opcode is used by the ADDIU rx, sp, immediate instruction
2. The ADDIUPC opcode is used by the ADDIU rx, pc, immediate instruction
3. The ADDIU8 opcode is used by the ADDIU rx, immediate instruction
4. The LWSP opcode is used by the LW rx, offset(sp) instruction
5. The LWPC opcode is used by the LW rx, offset(pc) instruction
6. The SWSP opcode is used by the SW rx, offset(sp) instruction

**Table 24.3 MIPS16e JAL(X) Encoding of the x Field**

<b>x</b>	bit 26	
	0	1
	JAL	JALX

**Table 24.4 MIPS16e SHIFT Encoding of the f Field**

<b>f</b>	bits 1..0			
	0	1	2	3
	00	01	10	11
	SLL	$\beta$	SRL	SRA

**Table 24.5 MIPS16e RRI-A Encoding of the f Field**

<b>f</b>	bit 4	
	0	1
	ADDIU <sup>1</sup>	$\beta$

1. The ADDIU function is used by the ADDIU ry, rx, immediate instruction

**Table 24.6 MIPS16e I8 Encoding of the funct Field**

<b>funct</b>	bits 10..8							
	0	1	2	3	4	5	6	7
	000	001	010	011	100	101	110	111
	BTEQZ	BTNEZ	SWRASP <sup>1</sup>	ADJSP <sup>2</sup>	SVRS $\delta$	MOV32R <sup>3</sup>	*	MOVR32 <sup>4</sup>

1. The SWRASP function is used by the SW ra, offset(sp) instruction
2. The ADJSP function is used by the ADDIU sp, immediate instruction
3. The MOV32R function is used by the MOVE r32, rz instruction
4. The MOVR32 function is used by the MOVE ry, r32 instruction

**Table 24.7 MIPS16e RRR Encoding of the f Field**

<b>f</b>	bits 1..0			
	0	1	2	3
	00	01	10	11
	$\beta$	ADDU	$\beta$	SUBU

**Table 24.8 MIPS16e RR Encoding of the Funct Field**

<b>funct</b>	bits 2..0								
	0	1	2	3	4	5	6	7	
bits 4..3	000	001	010	011	100	101	110	111	
0	00	$J(AL)R(C) \delta$	SDBBP	SLT	SLTU	SLLV	BREAK	SRLV	SRAV
1	01	$\beta$	*	CMP	NEG	AND	OR	XOR	NOT
2	10	MFHI	$CNVT \delta$	MFLO	$\beta$	$\beta$	*	$\beta$	$\beta$
3	11	MULT	MULTU	DIV	DIVU	$\beta$	$\beta$	$\beta$	$\beta$

**Table 24.9 MIPS16e I8 Encoding of the s Field when funct=SVRS**

<b>s</b>	bit 7	
	0	1
	RESTORE	SAVE

**Table 24.10 MIPS16e RR Encoding of the ry Field when funct=J(AL)R(C)**

<b>ry</b>	bits 7..5							
	0	1	2	3	4	5	6	7
	000	001	010	011	100	101	110	111
	JR rx	JR ra	JALR	*	JRC rx	JRC ra	JALRC	*

**Table 24.11 MIPS16e RR Encoding of the ry Field when funct=CNVT**

<b>ry</b>	bits 7..5							
	0	1	2	3	4	5	6	7
	000	001	010	011	100	101	110	111
	ZEB	ZEH	$\beta$	*	SEB	SEH	$\beta$	*

## 24.2 MIPS16e and MIPS16e2 Instruction Listing

For the MIPS16e-2 instruction set, many new instructions have been added as shown in [Table 24.12](#). The legacy MIPS16e instructions are listed by instruction type in [Table 24.13](#) through [Table 24.20](#).

For more information on the original MIPS16e instructions, refer to the document MD00076, *MIPS32 Architecture for Programmers Volume IV-a: The MIPS16e Application Instructions to the MIPS32 Architecture*.

For more information on the MIPS16e2 instructions, refer to the document MD01172, *MIPS32 Architecture for Programmers: MIPS16e2 Applications Specific Extension Technical Reference Manual*.

**Table 24.12 New Instructions in the MIPS16e2 ASE**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Type</b>
ADDIUGP	Add immediate unsigned global pointer	MIPS16e2
ANDI	Logical AND immediate	MIPS16e2
CACHE	Perform Cache operation	MIPS16e2
DI	Disable interrupts	MIPS16e2
DMT	Disable multithreading	MIPS16e2
DVPE	Disable virtual processor element.	MIPS16e2
EHB	Execution hazard barrier	MIPS16e2
EI	Enable interrupts	MIPS16e2
EMT	Enable multithreading	MIPS16e2
EVPE	Enable virtual processor element.	MIPS16e2
EXT	Extract bit field.	MIPS16e2
INS	Insert bit field.	MIPS16e2
LBGP	Load byte global pointer	MIPS16e2
LBUGP	Load byte unsigned global pointer	MIPS16e2
LHGP	Load half-word global pointer	MIPS16e2
LHUGP	Load half-word unsigned global pointer	MIPS16e2
LL	Load linked word.	MIPS16e2
LUI	Load upper immediate	MIPS16e2
LWGP	Load word global pointer	MIPS16e2
LWL	Load word left.	MIPS16e2
LWR	Load word right	MIPS16e2

**Table 24.12 New Instructions in the MIPS16e2 ASE (continued)**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Type</b>
MOVN	Move conditional on NOT zero.	MIPS16e2
MOVZ	Move conditional on zero.	MIPS16e2
MTC0	Move to Coprocessor 0.	MIPS16e2
MFC0	Move from Coprocessor 0.	MIPS16e2
ORI	Logical OR immediate	MIPS16e2
PAUSE	Wait for LL bit to clear.	MIPS16e2
PREF	Prefetch.	MIPS16e2
RDHWR	Read hardware register.	MIPS16e2
SBGP	Store byte global pointer	MIPS16e2
SC	Store conditional	MIPS16e2
SHGP	Store half-word global pointer	MIPS16e2
SWG	Store word global pointer	MIPS16e2
SWL	Store word left.	MIPS16e2
SWR	Store word right	MIPS16e2
SYNC	Synchronize shared memory.	MIPS16e2
XORI	Exclusive OR immediate.	MIPS16e2
COPYW	Copy word	MIPS16e2 ASMACRO
UCOPYW	Unaligned copy word	MIPS16e2 ASMACRO

**Table 24.13 MIPS16e Load and Store Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
LB	Load Byte	Yes
LBU	Load Byte Unsigned	Yes
LH	Load Halfword	Yes
LHU	Load Halfword Unsigned	Yes
LW	Load Word	Yes
SB	Store Byte	Yes
SH	Store Halfword	Yes
SW	Store Word	Yes

**Table 24.14 MIPS16e Save and Restore Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
RESTORE	Restore Registers and Deallocate Stack Frame	Yes
SAVE	Save Registers and Setup Stack Frame	Yes

**Table 24.15 MIPS16e ALU Immediate Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
ADDIU	Add Immediate Unsigned	Yes
CMPI	Compare Immediate	Yes
LI	Load Immediate	Yes
SLTI	Set on Less Than Immediate	Yes
SLTIU	Set on Less Than Immediate Unsigned	Yes

**Table 24.16 MIPS16e Arithmetic Two or Three Operand Register Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
ADDU	Add Unsigned	No
AND	AND	No
CMP	Compare	No
MOVE	Move	No
NEG	Negate	No
NOT	Not	No
OR	OR	No
SEB	Sign-Extend Byte	No
SEH	Sign-Extend Halfword	No
SLT	Set on Less Than	No
SLTU	Set on Less Than Unsigned	No
SUBU	Subtract Unsigned	No
XOR	Exclusive OR	No
ZEB	Zero-Extend Byte	No
ZEH	Zero-Extend Halfword	No



**Table 24.17 MIPS16e Special Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
BREAK	Breakpoint	No
SDBBP	Software Debug Breakpoint	No
EXTEND	Extend	No

**Table 24.18 MIPS16e Multiply and Divide Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
DIV	Divide	No
DIVU	Divide Unsigned	No
MFHI	Move From HI	No
MFLO	Move From LO	No
MULT	Multiply	No
MULTU	Multiply Unsigned	No

**Table 24.19 MIPS16e Jump and Branch Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
B	Branch Unconditional	Yes
BEQZ	Branch on Equal to Zero	Yes
BNEZ	Branch on Not Equal to Zero	Yes
BTEQZ	Branch on T Equal to Zero	Yes
BTNEZ	Branch on T Not Equal to Zero	Yes
JAL	Jump and Link	No
JALR	Jump and Link Register	No
JALRC	Jump and Link Register Compact	No
JALX	Jump and Link Exchange	No
JR	Jump Register	No
JRC	Jump Register Compact	No

**Table 24.20 MIPS16e Shift Instructions**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Extensible Instruction</b>
SRA	Shift Right Arithmetic	Yes
SRAV	Shift Right Arithmetic Variable	No
SLL	Shift Left Logical	Yes
SLLV	Shift Left Logical Variable	No
SRL	Shift Right Logical	Yes
SRLV	Shift Right Logical Variable	No

### **24.3 MIPS16e2 Implementation Specific Instructions**

In addition to the architecturally defined MIPS16e2 instructions in the document entitled *MIPS16e2 Application Specific Extension Technical Reference Manual*, there are two new MIPS16e2 implementation specific instructions, COPYW and UCOPY. The encoding and description for these instructions are indicated below.

31	27 26	24 23	21 20	16 15	11 10	8 7	5 4	2 1	0
EXTEND 11110	sel 000	001	Offset/16	RRR 11100	dst	src	000	Count - 1	
5	3	3	5	5	3	3	5		

**Format:** COPYW dst, src, offset, count

**MIPS16e2 ASMACRO**

**Purpose:** Copy Data to Aligned Memory Location Extended

Copy a specified number of words to a memory location when the source and destination addresses are aligned.

**Description:**  $\text{memory}[\text{GPR}[\text{dst}]] \leftarrow \text{memory}[\text{GPR}[\text{src}]]$

The COPYW instruction is implemented as an ASMACRO instruction. This instruction copies 1-4 words of data to an aligned memory location. The unsigned offset field allows multiple COPYW instructions to be placed in line without requiring the source and destination registers to be updated. It uses registers \$12 - \$15 as temporary registers. If an exception is detected, the entire macro sequence can be restarted.

Support for the instruction is indicated to software by Config7.CP (bit 22) = 1.

**Restrictions:**

Source and destination must be word-aligned addresses.

**Operation:**

This instruction does a copy of 1-4 words when source and destination are known to be aligned. This is converted into a sequence of loads and stores. To minimize pipeline hazards the instructions are interleaved as shown.

```

LW r12, offset(Xlat[src])
if (count >= 1) then
    LW r13, (offset+4)(Xlat[src])
endif
if (count >= 2) then
    LW r14, (offset+8)(Xlat[src])
endif
if (count == 3) then
    LW r15, (offset+12)(Xlat[src])
endif
SW r12, offset(Xlat[dst])
if (count >= 1) then
    SW r13, (offset+4)(Xlat[dst])
endif
if (count >= 2) then
    SW r14, (offset+8)(Xlat[dst])
endif
if (count == 3) then
    SW r15, (offset+12)(Xlat[dst])
endif

```

**Exceptions:**

Address Error

TLB Refill

TLB Invalid

TLB Modified

MPU Protection

Watch

Reserved Instruction

**Programming Notes:**

This instruction executes for a variable number of cycles and performs a variable number of loads from memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.

31	27 26	24 23	21 20	16 15	11 10	8 7	5 4	2 1	0
EXTEND 11110	sel 000	000	Offset/16	RRR 11100	dst	src	000	Count - 1	
5	3	3	5	5	3	3	5		

**Format:** UCOPYW dst, src, offset, count

**MIPS16e2 ASMACRO**

**Purpose:** Copy Data to Unaligned Memory Location Extended

Copy a specified number of words to a memory location when the source and destination addresses are unaligned.

**Description:** `memory[GPR[dst]] ← memory[GPR[src]]`

The UCOPYW macro instruction is implemented as an ASMACRO instruction. This instruction copies 1-4 words of data to an unaligned memory location. The unsigned offset field allows multiple UCOPYW instructions to be placed in line without requiring the source and destination registers to be updated. It uses registers \$12 - \$15 as temporary registers. If an exception is detected, the entire macro sequence can be restarted.

Support for this instruction is indicated to software by Config7.CP (bit 22) = 1.

**Restrictions:**

None

**Operation:**

This instruction does an unaligned copy of 1-4 words. It can overwrite \$12, \$13, \$14, \$15 as temporary registers. This is converted into a sequence of unaligned loads and stores. To minimize pipeline hazards the instructions are interleaved as shown.

if (BigEndian) then

```

LWL r12, offset(Xlat[src])
if (count >= 1) then
    LWL r13, (offset+4)(Xlat[src])
endif
if (count >= 2) then
    LWL r14, (offset+8)(Xlat[src])
endif
if (count == 3) then
    LWL r15, (offset+12)(Xlat[src])
endif
LWR r12, (offset+3)(Xlat[src])
if (count >= 1) then
    LWR r13, (offset+7)(Xlat[src])
endif
if (count >= 2) then
    LWR r14, (offset+11)(Xlat[src])
endif
if (count == 3) then
    LWR r15, (offset+15)(Xlat[src])
endif
SWL r12, offset(Xlat[dst])

```

```

    SWR r12, (offset+3)(Xlat[dst])
    if (count >= 1) then
        SWL r13, (offset+4)(Xlat[dst])
        SWR r13, (offset+7)(Xlat[dst])
    endif
    if (count >= 2) then
        SWL r14, (offset+8)(Xlat[dst])
        SWR r14, (offset+11)(Xlat[dst])
    endif
    if (count == 3) then
        SWL r15, (offset+12)(Xlat[dst])
        SWR r15, (offset+15)(Xlat[dst])
    endif
else
    LWR r12, offset(Xlat[src])
    if (count >= 1) then
        LWR r13, (offset+4)(Xlat[src])
    endif
    if (count >= 2) then
        LWR r14, (offset+8)(Xlat[src])
    endif
    if (count == 3) then
        LWR r15, (offset+12)(Xlat[src])
    endif
    LWL r12, (offset+3)(Xlat[src])
    if (count >= 1) then
        LWL r13, (offset+7)(Xlat[src])
    endif
    if (count >= 2) then
        LWL r14, (offset+11)(Xlat[src])
    endif
    if (count == 3) then
        LWL r15, (offset+15)(Xlat[src])
    endif
    SWR r12, offset(Xlat[dst])
    SWL r12, (offset+3)(Xlat[dst])
    if (count >= 1) then
        SWR r13, (offset+4)(Xlat[dst])
        SWL r13, (offset+7)(Xlat[dst])
    endif
    if (count >= 2) then
        SWR r14, (offset+8)(Xlat[dst])
        SWL r14, (offset+11)(Xlat[dst])
    endif
    if (count == 3) then

```

```
        SWR r15, (offset+12)(Xlat[dst])
        SWL r15, (offset+15)(Xlat[dst])
    endif
endif
```

**Exceptions:**

- Address Error
- TLB Refill
- TLB Invalid
- TLB Modified
- MPU Protection
- Watch
- Reserved Instruction

**Programming Notes:**

This instruction executes for a variable number of cycles and performs a variable number of loads from memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.

## Implementation-specific Instructions

This chapter describes the architectural definition for the following implementation-specific instructions in the inter-Aptiv Multiprocessing System.

- CACHE: Cache Operation
- LL: Load Linked Word
- PREF: Prefetch
- SC: Store Conditional
- SYNC: Synchronize Shared Memory
- TLBR: Read Indexed TLB Entry
- TLBWI: Write Indexed TLB Entry
- TLBWR: Write Random TLB Entry
- WAIT: Enter Standby Mode
- CACHEE: Cache Operation EVA
- LLE: Load Link EVA
- PREFE: Prefetch EVA
- SCE: Store Conditional EVA
- SAVE
- RESTORE

In addition to the above MIPS32 instructions, the MIPS16e2 version also implements the CACHE, LL, PREF, SC, and SYNC instructions. The implementation specific behavior will be the same as the MIPS32 descriptions below. For more information on the encoding of these instructions, refer to the document entitled MD01172, *MIPS16e2 Application Specific Extension Technical Reference Manual*.





Error Exception due to an non-aligned address.

A Cache Error exception may occur as a by-product of some operations performed by this instruction. For example, if a Writeback operation detects a cache or bus error during the processing of the operation, that error is reported via a Cache Error exception. Similarly, a Bus Error Exception may occur if a bus operation invoked by this instruction is terminated in an error. However, cache error exceptions must not be triggered by an Index Load Tag or Index Store tag operation, as these operations are used for initialization and diagnostic purposes.

An address Error Exception (with cause code equal AdEL) occurs if the effective address references a portion of the kernel address space which would normally result in such an exception. The preferred implementation is not to match on the cache instruction.

Bits [17:16] of the instruction specify the cache on which to perform the operation, as follows

**Table 25.2 Encoding of Bits[17:16] of CACHE Instruction**

Code	Name	Cache	Cop0 Registers Used
2'b00	I	Primary Instruction	<i>I</i> TagLo, <i>I</i> DataLo, <i>I</i> DataHi, <i>ErrCtl</i>
2'b01	D	Primary Data	<i>D</i> TagLo, <i>D</i> TagHi, <i>D</i> DataLo, <i>ErrCtl</i>
2'b10	T	Tertiary - Not supported	
2'b11	S	Secondary	<i>L2</i> TagLo

Some of the operations use coprocessor0 registers as either sources or destinations. Each of the caches has a separate set of Tag and Data registers. The last column in [Table 25.2](#) lists which registers are used by operations to each cache.

Bits [20:18] of the instruction specify the operation to perform. On Index Load Tag and Index Store Data operations, the specific word (primary D) or double-word (primary I, secondary) that is addressed is loaded into or read from the *D*DataLo (primary D), *L23*DataLo, and *L23*DataHi (secondary), or *I*DataLo and *I*DataHi (primary I) registers. All other cache instructions are line-based, and the word and byte indexes will not affect their operation

[Table 25.3](#) shows the normal mode condition where the *ErrCtl*<sub>WST</sub>, *ErrCtl*<sub>SPR</sub>, and *ErrCtl*<sub>ITC</sub> bits of the CP0 *ErrCtl* register are all cleared. Refer to the *ErrCtl* register of the CP0 chapter for more information.

**Table 25.3 Encoding of Bits [20:18] of the CACHE Instruction, ErrCtl[WST, SPR, ITC] Cleared**

Code	Caches	Name	Effective Address Operand Type	Operation
3'b000	I	Index Invalidate	Index	Set the state of the cache line at the specified index to invalid. This encoding may be used by software to invalidate the entire instruction cache by stepping through all valid indices.
	D, S	Index Writeback Invalidate	Index	If the state of the cache line at the specified index is valid and dirty, write the line back to the memory address specified by the cache tag. After that operation is completed, set the state of the cache line to invalid. If the line is valid but not dirty, set the state of the line to invalid.  This encoding may be used by software to invalidate the entire data cache by stepping through all valid indices. Note that Index Store Tag should be used to initialize the cache at powerup.

Table 25.3 Encoding of Bits [20:18] of the CACHE Instruction, ErrCtl[WST, SPR, ITC] Cleared (*continued*)

Code	Caches	Name	Effective Address Operand Type	Operation
3'b001	I	Index Load Tag	Index	<ul style="list-style-type: none"> <li>Read the tag for the cache line at the specified index into the <i>ITagLo</i> register.</li> <li>Read the data corresponding to the dword index into the <i>IDataLo</i> and <i>IDataHi</i> registers.</li> <li>If parity is implemented, read the parity bits corresponding to the data into <i>ErrCtl<sub>PI</sub></i> field.</li> </ul>
3'b001	D	Index Load Tag	Index	<ul style="list-style-type: none"> <li>Read the tag for the cache line at the specified index into the CP0 <i>DTagLo</i> and <i>DTagHi</i> registers.</li> <li>Read the data corresponding to the word index into the <i>DDataLo</i> register.</li> <li>If parity is enabled, data array parity bits are read into the <i>ErrCtl</i> register.</li> <li>If ECC is enabled, data and tag array ECC bits are read into the <i>DTagHi</i> register.</li> </ul>
3'b001	S	Index Load Tag	Index	<ul style="list-style-type: none"> <li>Read the tag for the cache line at the specified index into the CP0 <i>L23TagLo</i> register.</li> <li>Read the data corresponding to the dword index into the <i>L23DataLo</i> and <i>L23DataHi</i> registers.</li> </ul>
3'b010	I	Index Store Tag	Index	<ul style="list-style-type: none"> <li>Write the tag for the cache block at the specified index from the <i>ITagLo</i> register.</li> <li>If parity is implemented, the parity written into the cache is generated by the hardware if <i>ErrCtl<sub>PO</sub></i> = 0, or it is obtained from <i>ITagLo</i> if <i>ErrCtl<sub>PO</sub></i> = 1.</li> </ul>
3'b010	D	Index Store Tag	Index	<p>Write the tag for the cache line at the specified index from the CP0 <i>DTagLo</i> (data cache) register.</p> <p>By default, the tag parity value will be automatically calculated. For test purposes, the parity/ECC bits from the <i>DTagLo</i> register will be used if <i>ErrCtl<sub>PO</sub></i> is set. The L1 data cache ECC bits come from the <i>DTagHi</i> register.</p> <p>This encoding may be used by software to initialize the entire instruction or data caches by stepping through all valid indices. Doing so requires that the <i>TagLo</i> register associated with the cache be initialized first.</p>
3'b010	S	Index Store Tag	Index	<p>Write the tag for the L2 cache line at the specified index from the CP0 <i>L23TagLo</i> (L2 cache) register.</p> <p>By default, the tag parity value will be automatically calculated. For test purposes, the parity/ECC bits from the <i>L23TagLo</i> register will be used if <i>ErrCtl<sub>PO</sub></i> is set. The L2 cache ECC bits come from the <i>L23TagLo</i> register.</p>
3'b011	I, D	Reserved	Unspecified	Executed as a no-op
3'b011	S	Index Store Data	Index	<p>Write the <i>L23DataHi</i> and <i>L23DataLo</i> Coprocessor 0 register contents at the way and dword index specified.</p> <p>The ECC bits are always generated by the hardware (if present).</p>

Table 25.3 Encoding of Bits [20:18] of the CACHE Instruction, ErrCtl[WST, SPR, ITC] Cleared (*continued*)

Code	Caches	Name	Effective Address Operand Type	Operation
3'b100	All	Hit Invalidate	Address	If the cache line contains the specified address, set the state of the cache line to invalid. This encoding may be used by software to invalidate a range of addresses from the instruction cache by stepping through the address range by the line size of the cache.
3'b101	I	Fill	Address	Fill the cache from the specified address.  The cache line is refetched even if it is already in the cache. In that case, the existing copy in the cache is invalidated
	D, S	Hit WriteBack Invalidate	Address	If the cache line contains the specified address and it is valid and dirty, write the contents back to memory. After that operation is completed, set the state of the cache line to invalid. If the line is valid but not dirty, set the state of the line to invalid.  This encoding may be used by software to invalidate a range of addresses from the data cache by stepping through the address range by the line size of the cache.
3'b110	D, S	Hit WriteBack	Address	If the cache line contains the specified address and it is valid and dirty, write the contents back to memory. After the operation is completed, leave the state of the line valid, but clear the dirty state.
3'b111	All	Fetch and Lock	Address	If the cache does not contain the specified address, fill it from memory, performing a writeback if required, and set the state to valid and locked. If the cache already contains the specified address, set the state to locked. The way selected on fill from memory is the least recently used.  The lock state is cleared by executing an Index Invalidate, Index Writeback Invalidate, Hit Invalidate, or Hit Writeback Invalidate operation to the locked line, or via an Index Store Tag operation with the lock bit reset in the associated <i>TagLo</i> register.  It is illegal to lock all ways at a given cache index. If all ways are locked, subsequent references to that index will displace one of the locked lines.

Table 25.4 shows the condition for the way select test where the  $ErrCtl_{WST}$  bit is set, and the  $ErrCtl_{SPR}$  and  $ErrCtl_{ITC}$  bits of the CP0  $ErrCtl$  register are cleared. Refer to the  $ErrCtl$  register of the CP0 chapter for more information.

**Table 25.4 Encoding of Bits [20:18] of the CACHE Instruction,  $ErrCtl_{WST}$  Set,  $ErrCtl_{SPR}$ ,  $ITC$  Cleared**

Code	Caches	Name	Effective Address Operand Type	Operation
3'b001	All	Index Load WS	Index	Read the WS RAM at the specified index into the associated $ITagLo$ , $DTagLo$ , or $L23TagLo$ CP0 register.
3'b010	I	Index Store WS	Index	Update the WS RAM at the specified index from the $ITagLo$ CP0 register.
3'b010	D	Index Store WS	Index	Update the WS RAM at the specified index from the $DTagLo$ CP0 register.  If $ErrCtl_{PO}$ is set, the dirty parity values in the $DTagLo$ register will be written to the WS RAM. Otherwise, the parity will be calculated for the write data.
3'b010	S	Index Store WS	Index	Update the WS RAM at the specified index from the $L23TagLo$ CP0 register.  If $ErrCtl_{PO}$ is set, the dirty parity values in the $L23TagLo$ register will be written to the WS RAM. Otherwise, the parity will be calculated for the write data.
3'b011	I	Index Store Data	Index	Write the $IDataLo$ and $IDataHi$ CP0 register contents at the way and dword index specified.  If $ErrCtl_{PO}$ is set, the dirty parity values in the $ITagLo$ register will be written to the WS RAM. Otherwise, the parity will be calculated for the write data.
3'b011	D	Index Store Data	Index	Write the $DDataLo$ CP0 register contents at the way and word index specified.  If $ErrCtl_{PO}$ is set, the ECC value comes from the $DTagHi$ register.
3'b011	S	Index Store ECC	Index	Write the $L23DataLo$ register contents to the $ECC$ bits at the way and dword index specified.
All Others	All			Other codes should not be used while $ErrCtl_{WST}$ is set.

Table 25.5 shows the condition for the SPRAM access test where the  $ErrCtl_{SPR}$  bit is set, and the  $ErrCtl_{WTS}$  and  $ErrCtl_{ITC}$  bits of the CP0  $ErrCtl$  register are cleared. Refer to the  $ErrCtl$  register of the CP0 chapter for more information.

**Table 25.5 Encoding of Bits [20:18] of the CACHE Instruction,  $ErrCtl[SPR]$  Set,  $ErrCtl[WST, ITC]$  Cleared**

Code	Caches	Name	Effective Address Operand Type	Operation
3'b001	I	Index Load Tag	Index	Read the SPRAM tag at the specified index into the $I_{TagLo}$ Coprocessor 0 register. Also read the instruction data and precode information corresponding to the byte index into the $IDataHi$ , $IDataLo$ , and $ErrCtl$ registers
3'b001	D	Index Load Tag	Index	Read the SPRAM tag at the specified index into the $TagLo1$ Coprocessor 0 register.
3'b010	D	Index Store Tag	Index	Update the SPRAM tag at the specified index from the $TagLo$ Coprocessor 0 register.
3'b010	I	Index Store Tag	Index	Update the SPRAM tag at the specified index from the $TagLo$ Coprocessor 0 register.
3'b011	I	Index Store Data	Index	Write the $IDataLo$ and $IDataHi$ Coprocessor 0 register contents into the SPRAM at the dword index specified.  If $ErrCtl_{P0}$ is set, the dirty parity values in the $I_{TagLo}$ register will be written to the WS RAM. Otherwise, the parity will be calculated for the write data.
3'b011	D	Index Store Data	Index	Write the $DDataLo$ Coprocessor 0 register contents into the SPRAM at the word index specified.  If $ErrCtl_{P0}$ is set, $ErrCtl_{P1}$ is used for the parity value. Otherwise, the parity value is calculated for the write data. If ECC is enabled, the ECC value comes from the $D_{TagHi}$ register.
All Others	D			Other codes should not be used while $ErrCtl_{SPR}$ is set.
All	S			Secondary and Tertiary operations should not be performed while $ErrCtl_{SPR}$ is set.

Table 25.6 shows the condition for the duplicate tag array access where the  $ErrCtl_{WST}$  and  $ErrCtl_{SPR}$  bits are set, and the  $ErrCtl_{ITC}$  bit of the CP0  $ErrCtl$  register is cleared. Refer to the  $ErrCtl$  register of the CP0 chapter for more information.

**Table 25.6 Encoding of Bits [20:18] of the CACHE Instruction,  $ErrCtl_{[WST, SPR]}$  Set,  $ErrCtl_{[ITC]}$  Cleared**

Code	Caches	Name	Effective Address Operand Type	Operation
3'b001	D	Index Load Tag	Index	Read the duplicate tag array into the CP0 $DTagLo$ register.
3'b010	D	Index Store Tag	Index	Writes the duplicate tag array from the CP0 $DTagLo$ register. By default, the tag parity value will be automatically calculated. For test purposes, the parity/ECC bits from the $DTagLo/DTagHi$ registers will be used if $ErrCtl_{PO}$ is set.
All Others	D			Other codes should not be used while $ErrCtl_{WST}$ and $ErrCtl_{SPR}$ are set.

Table 25.6 shows the condition for the duplicate tag array access where the  $ErrCtl_{ITC}$  bit is set, and the  $ErrCtl_{ITC}$  and  $ErrCtl_{SPR}$  bits of the CP0  $ErrCtl$  register are cleared. Refer to the  $ErrCtl$  register of the CP0 chapter for more information.

**Table 25.7 Encoding of Bits [20:18] of the CACHE Instruction,  $ErrCtl_{[ITC]}$  Set,  $ErrCtl_{[WST, SPR]}$  Cleared**

Code	Caches	Name	Effective Address Operand Type	Operation
3'b001	D	Index Load Tag	Index	Read the ITC tag into the CP0 $DTagLo$ register.
3'b010	D	Index Store Tag	Index	Writes the ITC tag from the CP0 $DTagLo$ register.
All Others	D			Other codes should not be used while $ErrCtl_{ITC}$ is set.

**Restrictions:**

The operation of this instruction is **UNDEFINED** for any operation/cache combination that is not implemented.

The operation of this instruction is **UNDEFINED** if the operation requires an address, and that address is uncacheable.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

**Operation:**

```

if IsCoprocesorEnabled(0) then
    vAddr ← GPR[base] + sign_extend(offset)
    (pAddr, uncached) ← AddressTranslation(vAddr, DataReadReference)
    CacheOp(op, vAddr, pAddr)
else
    SignalException(CoprocessorUnusable, 0)
endif

```

**Exceptions:**

TLB Refill Exception.

TLB Invalid Exception

Coprocessor Unusable Exception

Address Error Exception

Cache Error Exception

Bus Error Exception

**Programming Notes:**


For cache operations that require an index, it is implementation dependent whether the effective address or the translated physical address is used as the cache index. Therefore, the index value should always be converted to a kseg0 address by ORing the index with 0x80000000 before being used by the cache instruction. For example, the following code sequence performs a data cache Index Store Tag operation using the index passed in GPR a0:

```

li    a1, 0x80000000    /* Base of kseg0 segment */
or    a0, a0, a1       /* Convert index to kseg0 address */
cache DCIndexStTag, 0(a1) /* Perform the index store tag operation */

```





**Format:** LL *rt*, *offset*(*base*)  
 LLE *rt*, *offset*(*base*) – Extended Virtual Address (EVA)

MIPS32

**Purpose:** Load Linked Word

To load a word from memory for an atomic read-modify-write. The LL and LLE instructions perform identical operations with one exception — when the processor is configured in Enhanced Virtual Address (EVA) mode, the LLE instruction is used to perform the virtual address translation using the user mapping of the address rather than the kernel mapping. The LL instruction can be used with unmapped addresses, in non-EVA mode, or when the kernel mapping is required.

**Description:**  $GPR[rt] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The LL and SC instructions provide the primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and written into GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

This begins a RMW sequence on the current TC. There can be only one active RMW sequence per TC. When an LL is executed, it starts an active RMW sequence replacing any other sequence that was active. The read-modify-write (RMW) sequence is completed by a subsequent SC instruction that either completes the RMW sequence atomically and succeeds, or does not and fails.

Executing LL on one TC does not cause an action that, by itself, causes an SC for the same block to fail on another TC.

An execution of LL does not have to be followed by execution of SC; a program is free to abandon the RMW sequence without attempting a write.

**Restrictions:**

The addressed location must be synchronizable by all TC's and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the SC instruction for the formal definition.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the effective address is non-zero, an Address Error exception occurs.

**Operation:**


```
vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ... 02 then
  SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
memword ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
GPR[rt] ← memword
LLbit ← 1
```

**Exceptions:**

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Watch

**Programming Notes:**

There is no Load Linked Word Unsigned operation corresponding to Load Word Unsigned.



**Format:** PREF *hint*, *offset*(*base*)  
 PREFE *hint*, *offset*(*base*) – Extended Virtual Address (EVA)

MIPS32

**Purpose:** Prefetch

To move data between memory and cache. The PREF and PREFE instructions perform identical operations with one exception — when the processor is configured in Enhanced Virtual Address (EVA) mode, the PREFE instruction is used to perform the virtual address translation using the user mapping of the address rather than the kernel mapping. The PREF instruction can be used with unmapped addresses, in non-EVA mode, or when the kernel mapping is required.

**Description:** `prefetch_memory(GPR[base] + offset)`

PREF adds the 16-bit signed *offset* to the contents of GPR *base* to form an effective byte address. The *hint* field supplies information about the way that the data is expected to be used.

PREF does not cause addressing-related exceptions, including TLB exceptions. If the address specified would cause an addressing exception, the exception condition is ignored and no data movement occurs. However, even if no data is moved, some action that is not architecturally visible, such as writeback of a dirty cache line, can take place.

It is implementation-dependent whether a Bus Error or Cache Error exception is reported, if such an error is detected as a by-product of the action taken by the PREF instruction.

PREF neither generates a memory operation nor modifies the state of a cache line for a location with an *uncached* memory access type, whether this type is specified by the address segment (e.g., *kseg1*), the programmed coherency attribute of a segment (e.g., the use of the K0, KU, or K23 fields in the *Config* register), or the per-page coherency attribute provided by the TLB.

If PREF results in a memory operation, the memory access type and coherency attribute used for the operation are determined by the memory access type and coherency attribute of the effective address, just as it would be if the memory operation had been caused by a load or store to the effective address.

The PREF instruction supports the PrepareForStore hint.

**Table 25.8 Values of *hint* Field for PREF Instruction**

Value	Name	Data Use and Desired Prefetch Action
0	load	Use: Prefetched data is expected to be read (not modified). Action: Fetch data as if for a load.
1	store	Use: Prefetched data is expected to be stored or modified. Action: Fetch data as if for a store.
2-3	Reserved	
4	load_streamed	Use: Prefetched data is expected to be read (not modified) but not reused extensively; it “streams” through cache.

Table 25.8 Values of *hint* Field for PREF Instruction

5	store_streamed	Use: Prefetched data is expected to be stored or modified but not reused extensively; it “streams” through cache.
6	load_retained	Use: Prefetched data is expected to be read (not modified) and reused extensively; it should be “retained” in the cache.
7	store_retained	Use: Prefetched data is expected to be stored or modified and reused extensively; it should be “retained” in the cache.
8-24	Reserved	
25	writeback_invalidate (also known as “nudge”)	Use: Data is no longer expected to be used.  Action: For a writeback cache, schedule a writeback of any dirty data. At the completion of the writeback, mark the state of any cache lines written back as invalid. If the cache line is not dirty, it is implementation dependent whether the state of the cache line is marked invalid or left unchanged. If the cache line is locked, no action is taken.
26-29	Reserved	
30	PrepareForStore	Use: Prepare the cache for writing an entire line, without the overhead involved in filling the line from memory.  Action: If the reference hits in the cache, no action is taken. If the reference misses in the cache, a line is selected for replacement, any valid and dirty victim is written back to memory, the entire line is filled with zero data, and the state of the line is marked as valid and dirty.  Programming Note: Because the cache line is filled with zero data on a cache miss, software must not assume that this action, in and of itself, can be used as a fast bzero-type function.
31	Reserved	

**Restrictions:**

None

**Operation:**

```

vAddr ← GPR[base] + sign_extend(offset)
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
Prefetch(CCA, pAddr, vAddr, DATA, hint)

```

**Exceptions:**

Bus Error, Cache Error

Prefetch does not take any TLB-related or address-related exceptions under any circumstances.


**Programming Notes:**

Prefetch cannot move data to or from a mapped location unless the translation for that location is present in the TLB. Locations in memory pages that have not been accessed recently may not have translations in the TLB, so prefetch may not be effective for such locations.

Prefetch does not cause addressing exceptions. A prefetch may be used using an address pointer before the validity of the pointer is determined without worrying about an addressing exception.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a by-product of the action taken by the PREF instruction. Typically, this only occurs in systems which have high-reliability requirements.

Prefetch operations have no effect on cache lines that were previously locked with the CACHE instruction.



**Format:** SC *rt*, *offset*(*base*)  
 SCE *rt*, *offset*(*base*) – Extended Virtual Address (EVA)

MIPS32

**Purpose:** Store Conditional Word

To store a word to memory to complete an atomic read-modify-write. The SC and SCE instructions perform identical operations with one exception — when the processor is configured in Enhanced Virtual Address (EVA) mode, the SCE instruction is used to perform the virtual address translation using the user mapping of the address rather than the kernel mapping. The SC instruction can be used with unmapped addresses, in non-EVA mode, or when the kernel mapping is required.

**Description:** if `atomic_update` then `memory[GPR[base] + offset] ← GPR[rt]`, `GPR[rt] ← 1`  
 else `GPR[rt] ← 0`

The LL and SC instructions provide primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

The 32-bit word in GPR *rt* is conditionally stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

The SC completes the RMW sequence begun by the preceding LL instruction executed on the processor. To complete the RMW sequence atomically, the following occur:

- The 32-bit word of GPR *rt* is stored into memory at the location specified by the aligned effective address.
- A 1, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

If the following event occurs between the execution of LL and SC, the SC fails:

- An ERET instruction is executed.

If either of the following events occurs between the execution of LL and SC, the SC may succeed or it may fail; the success or failure is not predictable. Portable programs should not cause one of these events.

- A memory access instruction (load, store, or prefetch) is executed on the processor executing the LL/SC.
- The instructions executed starting with the LL and ending with the SC do not lie in a 2048-byte contiguous region of virtual memory. (The region does not have to be aligned, other than the alignment required for instruction words.)

The following conditions must be true or the result of the SC is **UNPREDICTABLE**:

- Execution of SC must have been preceded by execution of an LL instruction.
- An RMW sequence executed without intervening events that would cause the SC to fail must use the same address in the LL and SC. The address is the same if the virtual address, physical address, and cache-coherence algorithm are identical.

**Restrictions:**

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

**Operation:**

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ... 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
dataword ← GPR[rt]
if LLbit then
    StoreMemory (CCA, WORD, dataword, pAddr, vAddr, DATA)
endif
GPR[rt] ← 0 || LLbit

```

**Exceptions:**

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

**Programming Notes:**

LL and SC are used to atomically update memory locations, as shown below.

```

L1:
LL    T1, (T0) # load counter
ADDI  T2, T1, 1 # increment
SC    T2, (T0) # try to store, checking for atomicity
BEQ   T2, 0, L1 # if not atomic (0), try again
NOP                   # branch-delay slot


```

Exceptions between the LL and SC cause SC to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, and floating point operations that trap or require software emulation assistance.

LL and SC function on a single processor for *cached noncoherent* memory so that parallel programs can be run on uniprocessor systems that do not support *cached coherent* memory access types.

Some systems may support atomic operations with an uncached memory access type. Uncached atomic operations must use a CCA of UC. Uncached atomic operations with a CCA of UCA are not supported. Uncached LL will be marked appropriately and system can begin monitoring accesses to that address. An uncached SC will first check the LL bit and if still set, will send out a conditional write request. The system will drop the write if another core or device has written to the location and process the write if there has not been a conflicting write. The success/failure of the write will be indicated in the SC destination register. Systems without this support might cause a bus error upon seeing an uncached LL or SC or may unconditionally allow the write.

A PAUSE instruction cannot be used to wait for an uncached semaphore to be unlocked.



**Format:** SYNC (stype = 0 implied)  
 SYNCE (stype = 0 implied) – Extended Virtual Address (EVA)

MIPS32

**Purpose:** Synchronize Shared Memory

To order loads and stores. The SYNC and SYNCE instructions perform identical operations with one exception — when the processor is configured in Enhanced Virtual Address (EVA) mode, the SYNCE instruction is used to perform the virtual address translation using the user mapping of the address rather than the kernel mapping. The SYNC instruction can be used with unmapped addresses, in non-EVA mode, or when the kernel mapping is required.

**Description:**

These types of ordering guarantees are available through the SYNC instruction:

- Completion Barriers
- Ordering Barriers

*Simple Description of Completion Barrier:*

- SYNC affects only *uncached* and *cached coherent* loads and stores. The specified memory instruction (loads or stores or both) that occur before the SYNC / SYNCE instruction must be completed before the loads and stores after the SYNC / SYNCE are allowed to start.
- Loads are completed when the destination register is written. Stores are completed when the stored value is visible to every other processor in the system.

*Detailed Description of Completion Barrier:*

- Every synchronizable specified memory instruction (loads or stores or both) that occurs in the instruction stream before the SYNC instruction must be already globally performed before any synchronizable specified memory instruction that occurs after the SYNC are allowed to be performed, with respect to any other processor or coherent I/O module.
- The barrier does not guarantee the order in which instruction fetches are performed.
- A stype value of zero will always be defined such that it performs the most complete set of synchronization operations that are defined. This means stype zero always does a completion barrier that affects both loads and stores preceding the SYNC instruction and both loads and stores that are subsequent to the SYNC instruction. Non-zero values of stype may be defined by the architecture or specific implementations to perform synchronization behaviors that are less complete than that of stype zero. If an implementation does not use one of these non-zero values to define a different synchronization behavior, then that non-zero value of stype must act the same as stype zero completion barrier. This allows software written for an implementation with a lighter-weight barrier to work on another implementation which only implements the stype zero completion barrier.
- A completion barrier is required, potentially in conjunction with SSNOP (in Release 1 of the Architecture) or EHB (in Release 2 of the Architecture), to guarantee that memory reference results are visible across operating mode changes. For example, a completion barrier is required on entry to and exit from Debug Mode to guarantee that memory affects are handled correctly.



- For the purposes of this description, the CACHE (CACHEE), PREF (PREFE) and PREFX instructions are treated as loads and stores. That is, these instructions and the memory transactions sourced by these instructions obey the ordering and completion rules of the completion barrier SYNC instruction.

*Completion Barrier Types:*

All completion barrier types will flush any pending writes and generate an external SYNC request. The core will wait for all pending reads to complete as well as the SYNC response.

- 0x2 - Implementation specific stype. Intervention SYNC. When coherence is enabled, this SYNC will generate a CoherentSync request. The CoherenceManager will respond to the SYNC when the interventions for all older coherent requests have been completed. If coherence is not enabled, will default to stype 0x0.
- 0x3 - Implementation specific stype. Memory SYNC. When coherence is enabled, this SYNC will also generate a CoherentSync request. When interventions for all older coherent requests have completed, the sync will be sent to memory interface unit. All pending transactions will be sent out. If the next level device (L2 or system) supports legacy SYNC transactions, as indicated by SI\_CM\_SyncTxEn = 1, and CM\_SYNC\_TX\_DISABLE in the CM Control GCR is 0, an external SYNC request will also be generated. The CM will send a response to the CoreType-lowercase when all prior requests have completed and a SYNC response is received (if it was externalized). If coherence is not enabled, will default to stype 0x0.
- 0x0 - If coherence is enabled, this will be mapped to either a type 0x2 or 0x3 based on the value of the SYNCCTL bit in the CM Control GCR. If coherence is not enabled, a legacy SYNC request will be generated. This will bypass the intervention pipeline in the CM and go directly to the memory unit. If SyncTxEn = 1 and CM\_SYNC\_TX\_DISABLE in the CM Control GCR is 0, an external SYNC request will be generated.

*Simple Description of Ordering Barrier:*

- The specified memory instructions (loads or stores or both) that occur before the SYNC instruction must always be ordered and globally visible to all cores before the specified memory instructions after the SYNC.
- Memory instructions which are ordered before other memory instructions are processed by the load/store datapath first before the other memory instructions.

*Detailed Description of Ordering Barrier:*

- Every synchronizable specified memory instruction (loads or stores or both) that occurs in the instruction stream before the SYNC instruction must reach a stage in the load/store datapath after which no instruction re-ordering is possible before any synchronizable specified memory instruction which occurs after the SYNC instruction in the instruction stream reaches the same stage in the load/store datapath.

NOTE: cached and uncached operations proceed down different data paths within the Coherence Manager. Because of that, this type of barrier will not enforce the ordering between cached and uncached requests. A Completion Barrier should be used if that ordering is required.

- If any memory instruction before the SYNC instruction in program order, generates a memory request to the external memory and any memory instruction after the SYNC instruction in program order also generates a memory request to external memory of the same type (cached or uncached), the memory request belonging to the older instruction must be globally performed before the time the memory request belonging to the younger instruction is globally performed.
- The barrier does not guarantee the order in which instruction fetches are performed.
- This barrier does not enforce the ordering of CACHE instructions. To ensure ordering of a CACHE instruction with other operations, a completion barrier type of SYNC should be used.

- For the purposes of this description, PREF and PREFX instructions are treated as loads and obey the same ordering rules as loads.

As compared to the completion barrier, the ordering barrier is a lighter-weight operation as it does not require the specified instructions before the SYNC to be already completed. Instead it only requires that those specified instructions which are subsequent to the SYNC in the instruction stream are never re-ordered for processing ahead of the specified instructions which are before the SYNC in the instruction stream. This potentially reduces how many cycles the barrier instruction must stall before it completes.

The Acquire and Release barrier types are used to minimize the memory orderings that must be maintained and still have software synchronization work.

The interAptiv core handles all ordering barriers identically.

- No external SYNC request will be generated and the core does not wait for pending transactions to complete.
- The LSU completes any pending evictions and wait until self interventions have been received for all Fill Store Buffers before proceeding.

NOTE: globalized CACHE instructions do not use an FSB entry, thus an Ordering Barrier does not wait for those to be completed. A Completion Barrier should be used to ensure that all prior CACHE instructions have completed.

- The BIU will stop merging on all Write Back Buffer (WBB) entries and put them into the external request queue.

Table 25.9 lists the available completion barrier and ordering barriers behaviors that can be specified using the stype field.

#### *MMIO SYNC STypes*

The interAptiv core defines two new MMIO specific SYNC SType values for determining how the SYNC instruction is handled within the MMIO. These two STypes are defined as follows:

- SYNC with stype = 0x6 - MMIO Only SYNC. This SYNC acts like a SYNC(0), but after the normal SYNC(0) processing, an uncached SYNC transaction is issued to all IOcus, The uncached SYNC request is ordered behind all previous MMIO requests. An external SYNC request can also be driven onto the IOcu MMIO port(s) if the device on the MMIO port supports legacy SYNC transactions, as indicated by the *CM\_IOC\_TX\_DISABLE/CM\_IOC1\_TX\_DISABLE* in the *CM3 Control GCR* is 0. The IOcu sends a response to the CM3 when all previous IOcu MMIO requests have been issued and a SYNC response is received (if it was externalized). The CM3 returns a response to the core when it has received a response from IOcus. If the cluster is built with no IOcus (indicated by *GCR\_CONFIG.NUMIOCU = 0*), then this sync is treated like an uncached memory sync.
- SYNC with stype = 0x7 - Implementation specific stype. MMIO and Memory SYNC. When coherence is enabled, this SYNC acts like a combination of SYNC(3) (coherent memory SYNC) and a SYNC(6) (MMIO only SYNC). The CM3 issues SYNCs to the memory interface and all IOcus' MMIO interface. The CM3 issues a response to the Core once a response has been received from the core interface and all IOcus. If the cluster is built with no IOcus (indicated by *GCR\_CONFIG.NUMIOCU = 0*), then this sync is treated like an uncached memory sync.

Table 25.9 Encodings of the Bits[10:6] of the SYNC Instruction; the STYPE Field

Code	Name	Type	Older instructions which must reach the load/store ordering point before the SYNC instruction completes	Younger instructions which must reach the load/store ordering point only after the SYNC instruction completes	Older instructions which must be globally performed when the SYNC instruction completes
0x0	SYNC or SYNC(0)	Completion	Loads, Stores	Loads, Stores	Loads, Stores
0x2	SYNC(2)	Completion	Loads, Stores	Loads, Stores	Loads, Stores
0x3	SYNC(3)	Completion	Loads, Stores	Loads, Stores	Loads, Stores
0x4	SYNC_WMB or SYNC(4)	Ordering	Stores	Stores	
0x6	MMIO only SYNC	Completion	Loads, Stores	Loads, Stores	Loads, Stores
0x7	MMIO and Memory SYNC	Completion	Loads, Stores	Loads, Stores	Loads, Stores
0x10	SYNC_MB or SYNC(16)	Ordering	Loads, Stores	Loads, Stores	
0x11	SYNC_ACQUIRE or SYNC(17)	Ordering	Loads	Loads, Stores	
0x12	SYNC_RELEASE or SYNC(18)		Loads, Stores	Stores	
0x13	SYNC_RMB or SYNC(19)		Loads	Loads	
All Others	RESERVED				

**Restrictions:**

None

**Operation:**

SyncOperation(stype)

**Exceptions:**

None

Software written to use a SYNC instruction with a non-zero stype value, expecting one type of barrier behavior, should only be run on hardware that actually implements the expected barrier behavior for that non-zero stype value or on hardware which implements a superset of the behavior expected by the software for that stype value. If the hardware does not perform the barrier behavior expected by the software, the system may fail.

31	26	25	24	6	5	0	
COP0 010000	CO 1	0 000 0000 0000 0000 0000				TLBWR 000110	
6	1	19				6	

**Format:** TLBWR

**MIPS32**

**Purpose:** Write Random TLB Entry

To write a TLB entry indexed by the *Random* register.

**Description:**

The TLB entry pointed to by the *Random* register is written from the contents of the *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers. If multiple TLB matches are detected on a TLB WR, the entries are silently invalidated. The information written to the TLB entry may be different from that in the *EntryHi*, *EntryLo0*, and *EntryLo1* registers, in that:

- The single G bit in the TLB entry is set from the logical AND of the G bits in the *EntryLo0* and *EntryLo1* registers.

**Restrictions:**

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

**Operation:**

```

if IsCoprocessorEnabled(0) then
    i ← Random
    TLB[i]Mask ← PageMaskMask
    TLB[i]VPN2_invalid ← 0
    TLB[i]VPN2 ← EntryHiVPN2 and not PageMaskMask # Implementation dependent
    TLB[i]ASID ← EntryHiASID
    TLB[i]G ← EntryLo1G and EntryLo0G
    TLB[i]PFN1 ← EntryLo1PFN and not PageMaskMask # Implementation dependent
    TLB[i]C1 ← EntryLo1C
    TLB[i]D1 ← EntryLo1D
    TLB[i]V1 ← EntryLo1V
    TLB[i]PFN0 ← EntryLo0PFN and not PageMaskMask # Implementation dependent
    TLB[i]C0 ← EntryLo0C
    TLB[i]D0 ← EntryLo0D
    TLB[i]V0 ← EntryLo0V
else
    SignalException(CoprocessorUnusable, 0)
endif

```

**Exceptions:**

Coprocessor Unusable

31	26	25	24	6	5	0
COP0 010000	CO 1	0 000 0000 0000 0000 0000			TLBR 000001	
6	1	19			6	

**Format:** TLBR

**MIPS32**

**Purpose:** Read Indexed TLB Entry

To read an entry from the TLB.

**Description:**

The *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers are loaded with the contents of the TLB entry pointed to by the *Index* register. Note that the value written to the *EntryHi*, *EntryLo0*, and *EntryLo1* registers may be different from that originally written to the TLB via these registers in that:

- The value returned in the G bit in both the *EntryLo0* and *EntryLo1* registers comes from the single G bit in the TLB entry. Recall that this bit was set from the logical AND of the two G bits in *EntryLo0* and *EntryLo1* when the TLB was written.

**Restrictions:**

The operation is **UNDEFINED** if the contents of the Index register are greater than or equal to the number of TLB entries in the processor.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

**Operation:**

```

if IsCoprocessorEnabled(0) then
  i ← Index
  if i > (TLBEntries - 1) then
    UNDEFINED
  endif
  PageMaskMask ← TLB[i]Mask
  EntryHi ← TLB[i]VPN2 ||
             02 || TLB[i]VPN2_invalid || 02 || TLB[i]ASID
  EntryLo1 ← 02Fill ||
             TLB[i]PFN1 ||
             TLB[i]C1 || TLB[i]D1 || TLB[i]V1 || TLB[i]G
  EntryLo0 ← 02Fill ||
             TLB[i]PFN0 ||
             TLB[i]C0 || TLB[i]D0 || TLB[i]V0 || TLB[i]G

else
  SignalException(CoprocessorUnusable, 0)
endif

```

**Exceptions:**

Coprocessor Unusable

31	26	25	24	6	5	0	
COP0 010000	CO 1	0 000 0000 0000 0000 0000				TLBWI 000010	
6	1	19				6	

**Format:** TLBWI

**MIPS32**

**Purpose:** Write Indexed TLB Entry

To write a TLB entry indexed by the *Index* register.

**Description:**

The TLB entry pointed to by the *Index* register is written from the contents of the *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers. If multiple TLB matches are detected on a TLBWI, the existing entries are silently invalidated. The information written to the TLB entry may be different from that in the *EntryHi*, *EntryLo0*, and *EntryLo1* registers, in that:

- The single G bit in the TLB entry is set from the logical AND of the G bits in the *EntryLo0* and *EntryLo1* registers.

**Restrictions:**

The operation is **UNDEFINED** if the contents of the Index register are greater than or equal to the number of TLB entries in the processor.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

**Operation:**

```

if IsCoprocessorEnabled(0) then
  i ← Index
  TLB[i]Mask ← PageMaskMask
  TLB[i]VPN2_invalid ← EntryHiVPN2_invalid
  TLB[i]VPN2 ← EntryHiVPN2
  TLB[i]ASID ← EntryHiASID
  TLB[i]G ← EntryLo1G and EntryLo0G
  TLB[i]PFN1 ← EntryLo1PFN
  TLB[i]C1 ← EntryLo1C
  TLB[i]D1 ← EntryLo1D
  TLB[i]V1 ← EntryLo1V
  TLB[i]PFN0 ← EntryLo0PFN
  TLB[i]C0 ← EntryLo0C
  TLB[i]D0 ← EntryLo0D
  TLB[i]V0 ← EntryLo0V
else
  SignalException(CoprocessorUnusable, 0)
endif

```

**Exceptions:**

Coprocessor Unusable



**Exceptions:**

Coprocessor Unusable





31	26	25	24	6	5	0
COP0 010000	CO 1	Implementation-Dependent Code			WAIT 100000	
6	1	19			6	

**Format:** WAIT

**MIPS32**

**Purpose:** Enter Standby Mode

Wait for Event

**Description:**

If the pipeline restarts as the result of an interrupt, that interrupt is taken between the WAIT instruction and the following instruction (*EPC* for the interrupt points to the instruction following the WAIT instruction).

**Restrictions:**

The operation of the processor is **UNDEFINED** if a WAIT instruction is placed in the delay slot of a branch or a jump.


If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

**Operation:**

```
I: Enter lower power mode
I+1:/* Potential interrupt taken here */
```

**Exceptions:**

Coprocessor Unusable Exception



**Format:** SAVE {ra,}{xsregs,}{aregs,}{framesize} (All arguments optional) **MIPS32 CorExtend**

**Purpose:** Save Registers and Setup Stack Frame Extended

To set up a stack frame on entry to a subroutine, saving return address, static, and argument registers, and adjusting the stack

**Description:** **Stack** ← GPR[ra] and/or **Stack** ← GPR[18-23,30] and/or **Stack** ← GPR[17] and/or **Stack** ← GPR[16] and/or **Stack** ← GPR[4-7], sp ← sp - (framesize \* 8)

SAVE is a MIPS32 CorExtend instruction that saves registers GPR[4-7] specified to be treated as incoming arguments by the aregs field. Save the ra register on the stack if the ra bit of the instruction is set. Save the number of registers in the set GPR[18-23, 30] indicated by the value of the xsregs field, and/or GPR 16 and/or GPR 17 (s0 and s1 in the MIPS ABI calling convention) on the stack if the corresponding s0 and s1 bits of the instruction are set. Save the number of registers in the range GPR[4-7] that are to be treated as static registers as indicated by the aregs field, and adjust the stack pointer by 8 times the 8-bit concatenated framesize value. Registers are stored with higher numbered registers at higher stack addresses.

#### Interpretation of the aregs Field

In the standard MIPS ABIs, GPR[4-7] are designated as argument passing registers, a0-a3. When they are so used, they must be saved on the stack at locations allocated by the caller of the routine being entered. In other MIPS16e calling sequences, however, it is possible that some of the registers GPR[4-7] will need to be saved as static registers on the local stack instead of on the caller stack. The encoding of the aregs field allows for 0-4 arguments, 0-4 statics, and for mixtures of the two. Registers are bound to arguments in ascending order, a0, a1, a2, and a3, and thus assigned to static values in the reverse order, GPR[7], GPR[6], GPR[5], and GPR[4]. The following table shows the encoding of the aregs field.

aregs Encoding (binary)	Registers Saved as Arguments	Registers Saved as Static Registers
0 0 0 0	None	None
0 0 0 1	None	GPR[7]
0 0 1 0	None	GPR[6], GPR[7]
0 0 1 1	None	GPR[5], GPR[6], GPR[7]
0 1 0 0	a0	None
0 1 0 1	a0	GPR[7]
0 1 1 0	a0	GPR[6], GPR[7]
0 1 1 1	a0	GPR[5], GPR[6], GPR[7]
1 0 0 0	a0, a1	None
1 0 0 1	a0, a1	GPR[7]
1 0 1 0	a0, a1	GPR[6], GPR[7]

<b>args Encoding (binary)</b>	<b>Registers Saved as Arguments</b>	<b>Registers Saved as Static Registers</b>
1 0 1 1	None	GPR[4], GPR[5], GPR[6], GPR[7]
1 1 0 0	a0, a1, a2	None
1 1 0 1	a0, a1, a2	GPR[7]
1 1 1 0	a0, a1, a2, a3	None
1 1 1 1	Reserved	Reserved

**Restrictions:**

If either of the 2 least-significant bits of the stack pointer are not zero, and any of the *ra*, *s0*, *s1*, or *xsregs* fields are non-zero or the *args* field contains an value that implies a register store, then an Address Error exception will occur.

**Operation:**

```

temp ← GPR[29]
temp2 ← GPR[29]
case args of
  0b0000 0b0001 0b0010 0b0011 0b1011: args ← 0
  0b0100 0b0101 0b0110 0b0111: args ← 1
  0b1000 0b1001 0b1010: args ← 2
  0b1100 0b1101: args ← 3
  0b1110: args ← 4
  otherwise: UNPREDICTABLE
endcase
if args > 0 then
  StoreStackWord(temp, GPR[4])
  if args > 1 then
    StoreStackWord(temp + 4, GPR[5])
    if args > 2 then
      StoreStackWord(temp + 8, GPR[6])
      if args > 3 then
        StoreStackWord(temp + 12, GPR[7])
      endif
    endif
  endif
endif
if ra = 1 then
  temp ← temp - 4
  StoreStackWord(temp, GPR[31])
endif
if xsregs > 0 then
  if xsregs > 1 then
    if xsregs > 2 then
      if xsregs > 3 then
        if xsregs > 4 then
          if xsregs > 5 then
            if xsregs > 6 then
              temp ← temp - 4
              StoreStackWord(temp, GPR[30])
            endif
          endif
          temp ← temp - 4
          StoreStackWord(temp, GPR[23])
        endif
      endif
    endif
  endif
endif

```

```

        endif
        temp ← temp - 4
        StoreStackWord(temp, GPR[22])
    endif
    temp ← temp - 4
    StoreStackWord(temp, GPR[21])
endif
temp ← temp - 4
StoreStackWord(temp, GPR[20])
endif
temp ← temp - 4
StoreStackWord(temp, GPR[19])
endif
temp ← temp - 4
StoreStackWord(temp, GPR[18])
endif
if s1 = 1 then
    temp ← temp - 4
    StoreStackWord(temp, GPR[17])
endif
if s0 = 1 then
    temp ← temp - 4
    StoreStackWord(temp, GPR[16])
endif
case aregs of
    0b0000 0b0100 0b1000 0b1100 0b1110: astatic ← 0
    0b0001 0b0101 0b1001 0b1101: astatic ← 1
    0b0010 0b0110 0b1010: astatic ← 2
    0b0011 0b0111: astatic ← 3
    0b1011: astatic ← 4
    otherwise: UNPREDICTABLE
endcase
if astatic > 0 then
    temp ← temp - 4
    StoreStackWord(temp, GPR[7])
    if astatic > 1 then
        temp ← temp - 4
        StoreStackWord(temp, GPR[6])
        if astatic > 2 then
            temp ← temp - 4
            StoreStackWord(temp, GPR[5])
            if astatic > 3 then
                temp ← temp - 4
                StoreStackWord(temp, GPR[4])
            endif
        endif
    endif
endif
temp ← temp2 - (0 || (framesize << 3))
GPR[29] ← temp

StoreStackWord(vaddr, value)
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
dataword ← value

```

```
    StoreMemory (CCA, WORD, dataword, pAddr, vAddr, DATA)
endfunction StoreStackWord
```


**Exceptions:**

TLB refill, TLB invalid, TLB modified, Address error, Bus Error

**Programming Notes:**

This instruction executes for a variable number of cycles and performs a variable number of stores to memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.

Behavior of the processor is **UNPREDICTABLE** for Reserved values of *aregs*.



**Format:** RESTORE {ra,}{xsregs,}{asregs,}{framesize} (All args optional) **MIPS32 CorExtend**

**Purpose:** Restore Registers and Deallocate Stack Frame Extended

To deallocate a stack frame before exit from a subroutine, restoring return address and static registers from and extended static register set, and adjusting the stack.

**Description:** GPR[ra] ← Stack and/or GPR[18-23,30] ← Stack and/or GPR[17] ← Stack and/or GPR[16] ← Stack and/or GPR[4-7] ← Stack, sp ← sp + (framesize\*8)

RESTORE is a MIPS32 CorExtend instruction that allows restoring of a number of registers at the start or end of a subroutine. Making these available to MIPS32 code can reduce the code size for functions that are not in MIPS16e mode. These feature the same arguments as the extended MIPS16e version and are encoded in the CorExtend opcode space. Support for these instructions is indicated to software by Config7.MCRO (bit 23) = 1.

This instruction restores the *ra* register from the stack if the *ra* bit is set in the instruction. Restore from the stack the number of registers in the set GPR[18-23,30] indicated by the value of the *xsregs* field. Restore from the stack GPR 16 and/or GPR 17 (*s0* and *s1* in the MIPS ABI calling convention) from the stack if the corresponding *s0* and *s1* bits of the instruction are set, restore from the stack the number of registers in the range GPR[4-7] indicated by the *aregs* field, and adjust the stack pointer by 8 times the 8-bit concatenated framesize value. Registers are loaded from the stack assuming higher numbered registers are stored at higher stack addresses.

The opcode and function field describe a general restore operation, with the *s* fields as variables. The individual RESTORE and SAVE instructions have specific values for this variable.

#### Interpretation of the *aregs* Field

In the standard MIPS ABIs, GPR[4-7] are designated as argument passing registers, *a0-a3*. When they are so used, they must be saved on the stack at locations allocated by the caller of the routine being entered, but need not be restored on a subroutine exit. In other MIPS16e calling sequences, however, it is possible that some of the registers GPR[4-7] will need to be saved as static registers on the local stack instead of on the caller stack, and restored before return from the subroutine. The encoding used for the *aregs* field of an extended RESTORE instruction is the same as that used for the extended SAVE, but since argument registers can be ignored for the purposes of a RESTORE, only the registers treated as static need be handled. The following table shows the encoding of the *aregs* field.

<b><i>aregs</i> Encoding (binary)</b>	<b>Registers Restored as Static Registers</b>
0 0 0 0	None
0 0 0 1	GPR[7]
0 0 1 0	GPR[6], GPR[7]
0 0 1 1	GPR[5], GPR[6], GPR[7]
0 1 0 0	None
0 1 0 1	GPR[7]
0 1 1 0	GPR[6], GPR[7]
0 1 1 1	GPR[5], GPR[6], GPR[7]

<i>aregs</i> Encoding (binary)	Registers Restored as Static Registers
1 0 0 0	None
1 0 0 1	GPR[7]
1 0 1 0	GPR[6], GPR[7]
1 0 1 1	GPR[4], GPR[5], GPR[6], GPR[7]
1 1 0 0	None
1 1 0 1	GPR[7]
1 1 1 0	None
1 1 1 1	Reserved

**Restrictions:**

If either of the 2 least-significant bits of the stack pointer are not zero, and any of the *ra*, *s0*, *s1*, or *xsregs* fields are non-zero or the *aregs* field contains an encoding that implies a register load, then an Address Error exception will occur.

**Operation:**

```

temp ← GPR[29] + (0 || (framesize << 3))
temp2 ← temp
if ra = 1 then
    temp ← temp - 4
    GPR[31] ← LoadStackWord(temp)
endif
if xsregs > 0 then
    if xsregs > 1 then
        if xsregs > 2 then
            if xsregs > 3 then
                if xsregs > 4 then
                    if xsregs > 5 then
                        if xsregs > 6 then
                            temp ← temp - 4
                            GPR[30] ← LoadStackWord(temp)
                        endif
                    endif
                endif
            endif
        endif
    endif
    temp ← temp - 4
    GPR[23] ← LoadStackWord(temp)
endif
    temp ← temp - 4
    GPR[22] ← LoadStackWord(temp)
endif
    temp ← temp - 4
    GPR[21] ← LoadStackWord(temp)
endif
    temp ← temp - 4
    GPR[20] ← LoadStackWord(temp)
endif
    temp ← temp - 4
    GPR[19] ← LoadStackWord(temp)
endif
    temp ← temp - 4
    GPR[18] ← LoadStackWord(temp)

```



```

endif
if s1 = 1 then
    temp ← temp - 4
    GPR[17] ← LoadStackWord(temp)
endif
if s0 = 1 then
    temp ← temp - 4
    GPR[16] ← LoadStackWord(temp)
endif
case aregs of
    0b0000 0b0100 0b1000 0b1100 0b1110: astatic ← 0
    0b0001 0b0101 0b1001 0b1101: astatic ← 1
    0b0010 0b0110 0b1010: astatic ← 2
    0b0011 0b0111: astatic ← 3
    0b1011: astatic ← 4
    otherwise: UNPREDICTABLE
endcase
if astatic > 0 then
    temp ← temp - 4
    GPR[7] ← LoadStackWord(temp)
    if astatic > 1 then
        temp ← temp - 4
        GPR[6] ← LoadStackWord(temp)
        if astatic > 2 then
            temp ← temp - 4
            GPR[5] ← LoadStackWord(temp)
            if astatic > 3 then
                temp ← temp - 4
                GPR[4] ← LoadStackWord(temp)
            endif
        endif
    endif
endif
GPR[29] ← temp2
LoadStackWord(vaddr)
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
memword ← LoadMemory (CCA, WORD, pAddr, vAddr, DATA)
LoadStackWord ← memword
enfunction LoadStackWord

```

**Exceptions:**

TLB refill, TLB invalid, MPU Protection, Address Error, Bus Error

**Programming Notes:**

This instruction executes for a variable number of cycles and performs a variable number of loads from memory. A full restart of the sequence of operations will be performed on return from any exception taken during execution.

## References

This appendix lists other publications available from MIPS. that are referenced elsewhere in this document. These documents are included in the interAptiv core release, or in some cases may be available on MIPS web site, <https://www.mips.com>.

1. MIPS32® interAptiv™ Multiprocessing System Hardware User's Manual  
MIPS Document: MD00905
2. MIPS32® Architecture For Programmers, Volume I: Introduction to the MIPS32® Architecture  
MIPS Document: MD0082
3. MIPS32® Architecture For Programmers, Volume II: The MIPS32® Instruction Set  
MIPS Document: MD0086
4. MIPS32® Architecture For Programmers, Volume III: The MIPS32® Privileged Resource Architecture  
MIPS Document: MD0090
5. MIPS32® Architecture For Programmers, Volume IV-a: The MIPS16e™ Application-Specific Extension to the MIPS32® Architecture  
MIPS Document: MD00076
6. MIPS32® Architecture For Programmers, Volume IV-e: The MIPS® DSP Application-Specific Extension to the MIPS32® Architecture  
MIPS Document: MD00374
7. Open Core Protocol Specification  
Available from the OCP International Partnership at <http://www.ocpip.org>



## Revision History

Change bars (vertical lines) in the margins of this document indicate significant changes in the document since its last release. Change bars are removed for changes that are more than one revision old.

Revision	Date	Description
01.24	August 1, 2013	Updated Section 3.2.1, Instruction TLB. Updated Figure 3.6, Address Translation Flow. Added Section 3.5.2, EVA Initial Configuration Parameters. Updated Section 4.2, L1 Instruction Cache. Updated Table 4.3, L1 Data Cache Organization. Updated Figure 4.1, L1 Data Cache Organization. Updated Section 4.3.2, L1 Data Cache Parity. Updated and reorganized Section 4.4, L1 Instruction and Data Cache Software Testing, and related subsections. Update Table 13.47, TCBCONTROLE register. Added new chapter 11, Policy Manager.
01.25	December 3, 2013	Updated references to DSP revision throughout document.
01.26	May 8, 2014	Miscellaneous updates from internal review.
02.00	May 13, 2016	Initial version of interAptiv MR2 SUM. - Updated Config.MT description in Chapter 2. - Added Chapter 4, MPU - Updated Chapter 5, Caches, to include support for 32K and 64K L2 cache sizes, support for 128-bit OCP bus width, and updated L2 cache write allocate policy. - Updated Chapter 6 Exceptions, to include support for new Memory Protection exception. Updated Exception priorities table. - Updated Chapter 7, CPC, to include section on CPC-basic build time option. Added CPC_TYPE field to CPC Revision register at offset 0x020. - Updated Chapter 8, CM2. Added Arbiter Priority global register at offset 0x0160. Added MMIO SYNC disable bits to the CM2 CGR Control register at offset 0x0010. Added encoding to the CM2_ACCESS_EN field of the CM2 Access Privilege register at offset 0x0020. Added encoding to the COH_DOMAIN_EN field of the Core Local Coherence Control register at offset 0x0008. Added ARB_PRI_PRESENT bit (23) the CM Arbiter Priority register at offset 0x0160. - Updated Chapter 9, GIC, to include that the GIC is a build time option. - Chapter 17, CM2 Debug. Removed L2 Data Width Counter table. Updated selected entries of the L2 Hit Qualifier table. Added bits 21 and 20 to this table. - Chapter 20, C++ Efficient programming. Update contents to include semaphore behavior. - Updated Chapter 24, MIPS16e2 ASE, to include a table of new instructions in the MIPS16e2 release. Updated to include information on new instruction buffer width. Added encoding and description of the COPYW and UCOPYW implementation specific instructions. - Added SAVE and RESTORE instructions to Chapter 25, Implementation Specific Instructions.

Revision	Date	Description
02.01	June 15, 2016	<p>Updated Segment Control registers in Chapter 4, MPU.</p> <p>Updated VPESchedule register in Chapter 10, Policy Manager.</p> <p>Updated TCSchedule register in Chapter 10, Policy Manager.</p> <p>Updated RESTORE instruction with edits from internal review.</p> <p>Updated MIPS16e2 instruction list in Table 24.12 of Chapter 24, MIPS16e ASE.</p> <p>Updated COPYW instruction definition in Chapter 24.</p>