



**MIPS32® M5100 Processor Core Family
Software User's Manual**

Document Number: MD00964

Revision 1.04

March 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.

Contents

Chapter 1: Introduction to the MIPS32® M5100™ Processor Core	19
1.1: Features	20
1.2: M5100™ Core Block Diagram	24
1.2.1: Required Logic Blocks	25
1.2.2: Optional Logic Blocks.....	33
Chapter 2: Pipeline of the M5100™ Core.....	40
2.1: Pipeline Stages.....	40
2.1.1: I Stage: Instruction Fetch	42
2.1.2: E Stage: Execution.....	42
2.1.3: M Stage: Memory Fetch.....	43
2.1.4: A Stage: Align	43
2.1.5: W Stage: Writeback	43
2.2: Multiply/Divide Operations	43
2.3: MDU Pipeline with DSP Module Enabled	44
2.3.1: MDU	44
2.3.2: DSP Module Instruction Latencies	45
2.3.3: High-performance MDU Pipeline Stages	48
2.3.4: High-performance MDU Divide Operations.....	49
2.4: MDU Pipeline — High-performance MDU with DSP Module Disabled	50
2.4.1: 32x16 Multiply (High-Performance MDU)	53
2.4.2: 32x32 Multiply (High-Performance MDU)	53
2.4.3: Divide (High-Performance MDU)	54
2.5: MDU Pipeline — Area-Efficient MDU with DSP Module Disabled.....	55
2.5.1: Multiply (Area-Efficient MDU).....	56
2.5.2: Multiply Accumulate (Area-Efficient MDU)	56
2.5.3: Divide (Area-Efficient MDU)	57
2.6: Branch Delay	57
2.7: Data Bypassing	58
2.7.1: Load Delay	60
2.7.2: Move from HI/LO and CP0 Delay.....	60
2.8: Coprocessor 2 Instructions	61
2.9: Interlock Handling.....	62
2.10: Slip Conditions.....	63
2.11: Instruction Interlocks.....	63
2.12: Hazards	64
2.12.1: Types of Hazards	65
2.12.2: Instruction Listing	66
2.12.3: Eliminating Hazards	66
Chapter 3: Floating-Point Unit of the M5100™ Core	67
3.1: Features Overview	67
3.1.1: IEEE Standard 754	68
3.2: Enabling the Floating-Point Coprocessor	69
3.3: Data Formats.....	69
3.3.1: Floating-Point Formats.....	69
3.3.2: Fixed-Point Formats.....	72

3.4: Floating-Point General Registers	73
3.4.1: FPRs and Formatted Operand Layout	73
3.4.2: Formats of Values Used in FP Registers	73
3.4.3: Binary Data Transfers (32-Bit and 64-Bit)	75
3.5: Floating-Point Control Registers	76
3.5.1: Floating-Point Implementation Register (FIR, CP1 Control Register 0)	77
3.5.2: Floating-Point Condition Codes Register (FCCR, CP1 Control Register 25)	79
3.5.3: Floating-Point Exceptions Register (FEXR, CP1 Control Register 26)	80
3.5.4: Floating-Point Enables Register (FENR, CP1 Control Register 28)	80
3.5.5: Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)	81
3.5.6: Operation of the FS/FO/FN Bits	84
3.5.7: FCSR Cause Bit Update Flow	87
3.6: Instruction Overview	88
3.6.1: Data Transfer Instructions	88
3.6.2: Arithmetic Instructions	90
3.6.3: Conversion Instructions	91
3.6.4: Formatted Operand-Value Move Instructions	92
3.6.5: Conditional Branch Instructions	92
3.6.6: Miscellaneous Instructions	93
3.7: Exceptions	93
3.7.1: Precise Exception Mode	94
3.7.2: Exception Conditions	94
3.8: Pipeline and Performance	97
3.8.1: Pipeline Overview	97
3.8.2: Bypassing	98
3.8.3: Repeat Rate and Latency	99
3.9: 2008 FPU Support	99
Chapter 4: The MIPS® DSP Module	101
4.1: Additional Register State for the DSP Module	101
4.1.1: HI-LO Registers	101
4.1.2: DSPControl Register	101
4.2: Software Detection of the DSP Module Revision 2	103
Chapter 5: The MIPS® Virtualization Module	104
5.1: Additional Register State for the VZ Module	104
5.2: Additional Instructions for the Virtualization Module	105
5.3: Software Detection of the VZ Module	106
Chapter 6: Memory Management of the M5100™ Core	107
6.1: Introduction	107
6.1.1: Memory Management Unit (MMU) in Guest Modes	107
6.1.2: Memory Management Unit (MMU) in Root Modes	107
6.2: Modes of Operation	108
6.2.1: Virtual Memory Segments	109
6.2.2: User Mode	111
6.2.3: Kernel Mode	112
6.2.4: Debug Mode	114
6.3: Fixed Mapping MMU	116
6.4: System Control Coprocessor	118
Chapter 7: Exceptions and Interrupts in the M5100™ Core	119

7.1: Exception Conditions.....	119
7.2: Exception Priority.....	120
7.3: Interrupts	121
7.3.1: Interrupt Modes.....	121
7.3.2: Generation of Exception Vector Offsets for Vectored Interrupts	130
7.3.3: MCU ASE Enhancement for Interrupt Handling.....	131
7.4: GPR Shadow Registers.....	132
7.5: Exception Vector Locations	133
7.6: General Exception Processing	135
7.7: Debug Exception Processing	137
7.8: Exception Descriptions	138
7.8.1: Reset/SoftReset Exception	138
7.8.2: Debug Single Step Exception	139
7.8.3: Debug Interrupt Exception	140
7.8.4: Non-Maskable Interrupt (NMI) Exception.....	140
7.8.5: Interrupt Exception.....	141
7.8.6: Debug Instruction Break Exception.....	141
7.8.7: Address Error Exception — Instruction Fetch/Data Access.....	141
7.8.8: SRAM Parity Error Exception.....	142
7.8.9: Bus Error Exception — Instruction Fetch or Data Access.....	142
7.8.10: Protection Exception	143
7.8.11: Debug Software Breakpoint Exception	143
7.8.12: Execution Exception — System Call.....	143
7.8.13: Execution Exception — Breakpoint.....	144
7.8.14: Execution Exception — Reserved Instruction.....	144
7.8.15: Execution Exception — Coprocessor Unusable	144
7.8.16: Execution Exception — CorExtend Unusable.....	145
7.8.17: Execution Exception — DSP Module State Disabled	145
7.8.18: Execution Exception — Coprocessor 2 Exception.....	145
7.8.19: Execution Exception — Implementation-Specific 1 Exception.....	146
7.8.20: Execution Exception — Integer Overflow.....	146
7.8.21: Execution Exception — Trap.....	146
7.8.22: Debug Data Break Exception.....	147
7.8.23: Complex Break Exception.....	147
7.9: Exception Handling and Servicing Flowcharts	147

Chapter 8: CP0 Registers of the M5100™ Core 151

8.1: CP0 Register Summary.....	151
8.2: CP0 Register Descriptions	153
8.2.1: Index Register (CP0 Register 0, Select 0).....	154
8.2.2: Random Register (CP0 Register 1, Select 0)	155
8.2.3: EntryLo0 and EntryLo1 Registers (CP0 Registers 2 and 3, Select 0).....	155
8.2.4: Context Register (CP0 Register 4, Select 0).....	157
8.2.5: UserLocal Register (CP0 Register 4, Select 2).....	158
8.2.6: PageMask Register (CP0 Register 5, Select 0).....	158
8.2.7: PageGrain Register (CP0 Register 5, Select 1).....	160
8.2.8: Wired Register (CP0 Register 6, Select 0).....	161
8.2.9: HWREna Register (CP0 Register 7, Select 0).....	162
8.2.10: BadVAddr Register (CP0 Register 8, Select 0).....	163
8.2.11: BadInstr Register (CP0 Register 8, Select 1).....	163
8.2.12: BadInstrP Register (CP0 Register 8, Select 2)	164
8.2.13: Count Register (CP0 Register 9, Select 0)	165
8.2.14: EntryHi Register (CP0 Register 10, Select 0).....	166

8.2.15: GuestCtl1 Register (CP0 Register 10, Select 4)	167
8.2.16: GuestCtl2 Register (CP0 Register 10, Select 5)	168
8.2.17: GuestCtl3 Register (CP0 Register 10, Select 6)	171
8.2.18: Compare Register (CP0 Register 11, Select 0)	172
8.2.19: GuestCtl0Ext Register (CP0 Register 11, Select 4)	173
8.2.20: Status Register (CP0 Register 12, Select 0)	175
8.2.21: IntCtl Register (CP0 Register 12, Select 1)	179
8.2.22: SRSCtl Register (CP0 Register 12, Select 2)	184
8.2.23: SRSSMap Register (CP0 Register 12, Select 3)	187
8.2.24: View_IPL Register (CP0 Register 12, Select 4)	188
8.2.25: SRSSMap2 Register (CP0 Register 12, Select 5)	188
8.2.26: GuestCtl0 Register (CP0 Register 12, Select 6)	189
8.2.27: GTOffset Register (CP0 Register 12, Select 7)	196
8.2.28: Cause Register (CP0 Register 13, Select 0)	196
8.2.29: View_RIPL Register (CP0 Register 13, Select 4)	201
8.2.30: NestedExc (CP0 Register 13, Select 5)	202
8.2.31: Exception Program Counter (CP0 Register 14, Select 0)	203
8.2.32: NestedEPC (CP0 Register 14, Select 2)	204
8.2.33: Processor Identification (CP0 Register 15, Select 0)	205
8.2.34: EBase Register (CP0 Register 15, Select 1)	206
8.2.35: CDMMBase Register (CP0 Register 15, Select 2)	207
8.2.36: Config Register (CP0 Register 16, Select 0)	208
8.2.37: Config1 Register (CP0 Register 16, Select 1)	210
8.2.38: Config2 Register (CP0 Register 16, Select 2)	211
8.2.39: Config3 Register (CP0 Register 16, Select 3)	212
8.2.40: Config4 Register (CP0 Register 16, Select 4)	216
8.2.41: Config5 Register (CP0 Register 16, Select 5)	217
8.2.42: Config7 Register (CP0 Register 16, Select 7)	218
8.2.43: SecurityCtrl Register (CP0 Register 22, Select 0)	219
8.2.44: SecuritySwPRNG Register (CP0 Register 22, Select 1)	220
8.2.45: SecurityHwPRNG Register (CP0 Register 22, Select 2)	220
8.2.46: Debug Register (CP0 Register 23, Select 0)	221
8.2.47: Trace Control Register (CP0 Register 23, Select 1)	225
8.2.48: Trace Control2 Register (CP0 Register 23, Select 2)	227
8.2.49: User Trace Data1 Register (CP0 Register 23, Select 3)/User Trace Data2 Register (CP0 Register 24, Select 3)	229
8.2.50: TraceBPC Register (CP0 Register 23, Select 4)	230
8.2.51: Debug2 Register (CP0 Register 23, Select 6)	231
8.2.52: Debug Exception Program Counter Register (CP0 Register 24, Select 0)	232
8.2.53: Performance Counter Register (CP0 Register 25, select 0-3)	233
8.2.54: ErrCtl Register (CP0 Register 26, Select 0)	239
8.2.55: CacheErr Register (CP0 Register 27, Select 0)	240
8.2.56: ErrorEPC (CP0 Register 30, Select 0)	241
8.2.57: DeSave Register (CP0 Register 31, Select 0)	241
8.2.58: KScratchn Registers (CP0 Register 31, Selects 2 to 3)	242

Chapter 9: Hardware and Software Initialization of the M5100™ Core	243
9.1: Hardware-Initialized Processor State	243
9.1.1: Coprocessor 0 State	243
9.1.2: Bus State Machines	244
9.1.3: Static Configuration Inputs	244
9.1.4: Fetch Address	244
9.2: Software Initialized Processor State	244

9.2.1: Register File	244
9.2.2: Coprocessor 0 State	244
Chapter 10: Power Management of the M5100™ Core	246
10.1: Register-Controlled Power Management	246
10.2: Instruction-Controlled Power Management	247
Chapter 11: EJTAG Debug Support in the M5100™ Core	248
11.1: Debug Control Register	248
11.2: Hardware Breakpoints	253
11.2.1: Data Breakpoints.....	254
11.2.2: Complex Breakpoints	254
11.2.3: Conditions for Matching Breakpoints	254
11.2.4: Debug Exceptions from Breakpoints.....	258
11.2.5: Breakpoint Used as Triggerpoint.....	259
11.2.6: Instruction Breakpoint Registers	259
11.2.7: Data Breakpoint Registers	264
11.2.8: Complex Breakpoint Registers.....	271
11.3: Complex Breakpoint Usage	275
11.3.1: Checking for Presence of Complex Break Support.....	275
11.3.2: General Complex Break Behavior.....	276
11.3.3: Usage of Pass Counters	276
11.3.4: Usage of Tuple Breakpoints.....	277
11.3.5: Usage of Priming Conditions.....	277
11.3.6: Usage of Data Qualified Breakpoints	278
11.3.7: Usage of Stopwatch Timers	278
11.4: Test Access Port (TAP)	279
11.4.1: EJTAG Internal and External Interfaces.....	279
11.4.2: Test Access Port Operation	280
11.4.3: Test Access Port (TAP) Instructions	283
11.5: EJTAG TAP Registers	286
11.5.1: Instruction Register	286
11.5.2: Data Registers Overview	287
11.5.3: Processor Access Address Register.....	295
11.5.4: Fastdata Register (TAP Instruction FASTDATA).....	296
11.6: TAP Processor Accesses	297
11.6.1: Fetch/Load and Store from/to EJTAG Probe Through dmseg.....	298
11.7: SecureDebug.....	299
11.7.1: Disabling EJTAG Debugging	299
11.7.2: EJTAG Features Unmodified by SecureDebug	300
11.8: iFlowtrace™ Mechanism	300
11.8.1: A Simple Instruction-Only Tracing Scheme	301
11.8.2: Special Trace Modes	302
11.8.3: ITCB Overview	305
11.8.4: ITCB iFlowtrace Interface.....	305
11.8.5: TCB Storage Representation	306
11.8.6: ITCB Register Interface for Software Configurability	307
11.8.7: ITCB iFlowtrace Off-Chip Interface	311
11.8.8: Breakpoint-Based Enabling of Tracing.....	311
11.9: PC/Data Address Sampling.....	312
11.9.1: PC Sampling in Wait State.....	313
11.9.2: Data Address Sampling	313
11.10: Fast Debug Channel.....	313

11.10.1: Common Device Memory Map.....	314
11.10.2: Fast Debug Channel Interrupt.....	314
11.10.3: M5100™ Core FDC Buffers.....	314
11.10.4: Sleep mode.....	316
11.10.5: FDC TAP Register.....	316
11.10.6: Fast Debug Channel Registers.....	317
11.11: cJTAG Interface.....	321
Chapter 12: Instruction Set Overview.....	322
12.1: CPU Instruction Formats.....	322
12.2: Load and Store Instructions.....	323
12.2.1: Scheduling a Load Delay Slot.....	323
12.2.2: Defining Access Types.....	323
12.3: Computational Instructions.....	324
12.3.1: Cycle Timing for Multiply and Divide Instructions.....	325
12.4: Jump and Branch Instructions.....	325
12.4.1: Overview of Jump Instructions.....	325
12.4.2: Overview of Branch Instructions.....	325
12.5: Control Instructions.....	325
12.6: Coprocessor Instructions.....	325
12.7: Enhancements to the MIPS Architecture.....	325
12.7.1: CLO - Count Leading Ones.....	326
12.7.2: CLZ - Count Leading Zeros.....	326
12.7.3: MADD - Multiply and Add Word.....	326
12.7.4: MADDU - Multiply and Add Unsigned Word.....	326
12.7.5: MSUB - Multiply and Subtract Word.....	326
12.7.6: MSUBU - Multiply and Subtract Unsigned Word.....	327
12.7.7: MUL - Multiply Word.....	327
12.7.8: SSNOP- Superscalar Inhibit NOP.....	327
12.8: MCU ASE Instructions.....	327
12.8.1: ACLR.....	327
12.8.2: ASET.....	327
12.8.3: IRET.....	327
Chapter 13: M5100™ Processor Core Instructions.....	328
13.1: Understanding the Instruction Descriptions.....	328
13.2: M5100™ Core Opcode Map.....	328
13.3: MIPS32® Instruction Set for the M5100™ Core.....	333
ACLR.....	340
ACLR.....	341
ASET.....	342
IRET.....	344
IRET.....	348
LL.....	352
SC.....	353
SYNC.....	355
WAIT.....	356
Chapter 14: microMIPS™ Instruction Set Architecture.....	357
14.1: Overview.....	357
14.1.1: MIPSr3™ Architecture.....	357
14.1.2: Default ISA Mode.....	358

14.1.3: Software Detection.....	358
14.1.4: Compliance and Subsetting	358
14.1.5: Mode Switch.....	358
14.1.6: Branch and Jump Offsets.....	359
14.1.7: Coprocessor Unusable Behavior	359
14.2: Instruction Formats.....	359
14.2.1: Instruction Stream Organization and Endianness.....	362
14.3: microMIPS Re-encoded Instructions	362
14.3.1: 16-Bit Category	363
14.3.2: 16-bit Instruction Register Set.....	368
14.3.3: 32-Bit Category	370
Appendix A: References	372
Appendix B: Revision History	374

Tables

Table 2.1: High-performance MDU Stalls	45
Table 2.2: DSP Instruction Delays	46
Table 2.3: Delays for Interesting Sequences with DSPControl Dependency.....	47
Table 2.4: MDU Instruction Latencies (High-Performance MDU)	51
Table 2.5: MDU Instruction Repeat Rates (High-Performance MDU).....	52
Table 2.6: M5100™ Core Instruction Latencies (Area-Efficient MDU).....	55
Table 2.7: Pipeline Interlocks	62
Table 2.8: Instruction Interlocks	64
Table 2.9: Execution Hazards	65
Table 2.10: Instruction Hazards	65
Table 2.11: Hazard Instruction Listing	66
Table 3.1: Parameters of Floating-Point Data Types	69
Table 3.2: Value of Single or Double Floating-Point Data Type Encoding.....	70
Table 3.3: Value Supplied When a New Quiet NaN is Created	72
Table 3.4: Coprocessor 1 Register Summary	77
Table 3.5: Read/Write Properties.....	77
Table 3.6: FIR Bit Field Descriptions	78
Table 3.7: FCCR Bit Field Descriptions	80
Table 3.8: FEXR Bit Field Descriptions.....	80
Table 3.9: FENR Bit Field Descriptions	81
Table 3.10: FCSR Bit Field Descriptions	82
Table 3.11: Cause, Enables, and Flags Definitions	83
Table 3.12: Rounding Mode Definitions	84
Table 3.13: Zero Flushing for Tiny Results	85
Table 3.14: Handling of Denormalized Operand Values and Tiny Results Based on FS Bit Setting.....	85
Table 3.15: Handling of Tiny Intermediate Result Based on the FO and FS Bit Settings	85
Table 3.16: Handling of Tiny Final Result Based on FN and FS Bit Settings	86
Table 3.17: Recommended FS/FO/FN Settings	86
Table 3.18: FPU Data Transfer Instructions	89
Table 3.19: FPU Loads and Stores.....	89
Table 3.20: FPU Move To and From Instructions	89
Table 3.21: FPU IEEE Arithmetic Operations	90
Table 3.22: FPU-Approximate Arithmetic Operations	90
Table 3.23: FPU Multiply-Accumulate Arithmetic Operations	91
Table 3.24: FPU Conversion Operations Using the FCSR Rounding Mode.....	91
Table 3.25: FPU Conversion Operations Using a Directed Rounding Mode	91
Table 3.26: FPU Formatted Operand Move Instruction	92
Table 3.27: FPU Conditional Move on True/False Instructions.....	92
Table 3.28: FPU Conditional Move on Zero/Non-Zero Instructions	92
Table 3.29: FPU Conditional Branch Instructions	93
Table 3.30: Deprecated FPU Conditional Branch Likely Instructions	93
Table 3.31: CPU Conditional Move on FPU True/False Instructions	93
Table 3.32: Result for Exceptions Not Trapped	95
Table 3.33: M5100™ Core FPU Latency and Repeat Rate.....	99
Table 4.1: MIPS® DSP Module Control Register (DSPControl) Field Descriptions	102
Table 4.2: DSPControl ouflag Bits	103
Table 5.1: Virtualization Module Coprocessor 0 Registers in Numerical Order	105

Table 5.2: New and Modified Instructions	105
Table 6.1: User Mode Segments	111
Table 6.2: Kernel Mode Segments	113
Table 6.3: Physical Address and Cache Attributes for dseg, dmseg, and drseg Address Spaces	115
Table 6.4: CPU Access to drseg Address Range	115
Table 6.5: CPU Access to dmseg Address Range	116
Table 6.6: Cacheability of Segments with Block Address Translation	116
Table 7.1: Priority of Exceptions	120
Table 7.2: Interrupt Modes	122
Table 7.3: Relative Interrupt Priority for Vectored Interrupt Mode.....	125
Table 7.4: Exception Vector Offsets for Vectored Interrupts.....	130
Table 7.5: Exception Vector Base Addresses.....	134
Table 7.6: Exception Vector Offsets	134
Table 7.7: Exception Vectors	134
Table 7.8: Value Stored in EPC, ErrorEPC, or DEPC on an Exception.....	135
Table 7.9: Debug Exception Vector Addresses	138
Table 7.10: Register States an Interrupt Exception	141
Table 7.11: CP0 Register States on an Address Exception Error.....	142
Table 7.12: CP0 Register States on a SRAM Parity Error Exception	142
Table 7.13: Register States on a Coprocessor Unusable Exception	145
Table 8.1: CP0 Registers.....	151
Table 8.2: CP0 Register R/W Field Types	153
Table 8.3: Index Register Field Descriptions	154
Table 8.4: Random Register Field Descriptions.....	155
Table 8.5: EntryLo0, EntryLo1 Register Field Descriptions	156
Table 8.6: Cache Coherency Attributes	157
Table 8.7: Context Register Field Descriptions.....	157
Table 8.8: UserLocal Register Field Descriptions	158
Table 8.9: PageMask Register Field Descriptions	158
Table 8.10: Values for the Mask and MaskX ¹ Fields of the PageMask Register.....	159
Table 8.11: PageGrain Register Field Descriptions	160
Table 8.12: HWREna Register Field Descriptions	162
Table 8.13: BadVAddr Register Field Description	163
Table 8.14: BadInstr Register Field Descriptions.....	164
Table 8.15: BadInstrP Register Field Descriptions	165
Table 8.17: EntryHi Register Field Descriptions	166
Table 8.16: Count Register Field Description	166
Table 8.18: GuestCtl1 Register Field Descriptions	168
Table 8.19: Non-EIC Mode GuestCtl2 Register Field Descriptions	169
Table 8.20: EIC Mode GuestCtl2 Register Field Descriptions	171
Table 8.21: GuestCtl3 Register Field Descriptions	172
Table 8.22: Compare Register Field Description	172
Table 8.23: GuestCtl0Ext Register Field Descriptions	173
Table 8.24: Status Register Field Descriptions.....	175
Table 8.25: IntCtl Register Field Descriptions.....	180
Table 8.26: SRSCtl Register Field Descriptions	184
Table 8.27: Sources for new SRSCtl _{CSS} on an Exception or Interrupt	187
Table 8.28: SRSSMap Register Field Descriptions.....	187
Table 8.29: View_IPL Register Field Descriptions	188
Table 8.30: SRSSMap Register Field Descriptions.....	189
Table 8.31: GuestCtl0 Register Field Descriptions	190
Table 8.32: GuestCtl0 GExcCode values	195
Table 8.33: GTOffset Register Field Descriptions.....	196

Table 8.34: Cause Register Field Descriptions.....	197
Table 8.35: Cause Register ExcCode Field.....	201
Table 8.36: View_RIPL Register Field Descriptions.....	202
Table 8.37: NestedExc Register Field Descriptions.....	203
Table 8.38: EPC Register Field Description.....	204
Table 8.39: NestedEPC Register Field Descriptions.....	205
Table 8.40: PRId Register Field Descriptions.....	205
Table 8.41: EBase Register Field Descriptions.....	207
Table 8.42: CDMMBase Register Field Descriptions.....	207
Table 8.43: Config Register Field Descriptions.....	208
Table 8.44: Cache Coherency Attributes.....	209
Table 8.45: Config1 Register Field Descriptions — Select 1.....	210
Table 8.46: Config2 Register Field Descriptions — Select 1.....	211
Table 8.47: Config3 Register Field Descriptions.....	212
Table 8.48: Config4 Register Field Descriptions.....	216
Table 8.49: Config5 Register Field Descriptions.....	218
Table 8.50: Config7 Register Field Descriptions.....	219
Table 8.51: SecurityCtrl Register Field Descriptions.....	220
Table 8.52: SecuritySwPRNG Register Field Descriptions.....	220
Table 8.53: SecurityHwPRNG Register Field Descriptions.....	221
Table 8.54: Debug Register Field Descriptions.....	222
Table 8.55: TraceControl Register Field Descriptions.....	225
Table 8.56: TraceControl2 Register Field Descriptions.....	228
Table 8.57: UserTraceData1/UserTraceData2 Register Field Descriptions.....	229
Table 8.58: TraceBPC Register Field Descriptions.....	230
Table 8.59: Debug2 Register Field Descriptions.....	231
Table 8.60: DEPC Register Formats.....	232
Table 8.61: Performance Counter Register Selects.....	233
Table 8.62: Performance Counter Control Register Field Descriptions.....	233
Table 8.63: Performance Counter Events Sorted by Event Number.....	235
Table 8.64: Performance Counter Event Descriptions Sorted by Event Type.....	237
Table 8.65: Performance Counter Count Register Field Descriptions.....	239
Table 8.67: CacheErr Register Field Descriptions (Primary Caches).....	240
Table 8.66: Errctl Register Field Descriptions.....	240
Table 8.68: ErrorEPC Register Field Description.....	241
Table 8.69: DeSave Register Field Description.....	242
Table 8.70: KScratchn Register Field Descriptions.....	242
Table 11.1: DCR Register Field Descriptions.....	249
Table 11.2: Addresses for Instruction Breakpoint Registers.....	259
Table 11.3: IBS Register Field Descriptions.....	260
Table 11.4: IBAn Register Field Descriptions.....	260
Table 11.5: IBMn Register Field Descriptions.....	261
Table 11.6: IBASIDn Register Field Descriptions.....	261
Table 11.7: IBCn Register Field Descriptions.....	262
Table 11.8: IBCCn Register Field Descriptions.....	263
Table 11.9: IBPCn Register Field Descriptions.....	264
Table 11.10: Addresses for Data Breakpoint Registers.....	264
Table 11.11: DBS Register Field Descriptions.....	265
Table 11.12: DBAn Register Field Descriptions.....	265
Table 11.13: DBMn Register Field Descriptions.....	266
Table 11.14: DBASIDn Register Field Descriptions.....	266
Table 11.15: DBCn Register Field Descriptions.....	267
Table 11.16: DBVn Register Field Descriptions.....	268

Table 11.17: DBCCn Register Field Descriptions	269
Table 11.18: DBPCn Register Field Descriptions	270
Table 11.19: DVM Register Field Descriptions	270
Table 11.20: Addresses for Complex Breakpoint Registers	271
Table 11.21: CBTC Register Field Descriptions	271
Table 11.23: Priming Conditions and Register Values for 6I/2D Configuration	273
Table 11.24: Priming Conditions and Register Values for 8I/4D Configuration	273
Table 11.22: PrCndA Register Field Descriptions.....	273
Table 11.25: STCtl Register Field Descriptions	274
Table 11.26: STCtl Register Field Descriptions	275
Table 11.27: EJTAG Interface Pins	279
Table 11.28: Implemented EJTAG Instructions	284
Table 11.30: Implementation Register Descriptions	288
Table 11.29: Device Identification Register	288
Table 11.31: EJTAG Control Register Descriptions.....	290
Table 11.32: Fastdata Register Field Description	296
Table 11.33: Operation of the FASTDATA access	297
Table 11.34: EJ_DisableProbeDebug Signal Overview.....	300
Table 11.35: Data Bus Encoding	306
Table 11.36: Tag Bit Encoding.....	306
Table 11.37: Control/Status Register Field Descriptions	308
Table 11.38: ITCBTW Register Field Descriptions	309
Table 11.39: ITCBRDP Register Field Descriptions	310
Table 11.40: ITCBWRP Register Field Descriptions.....	310
Table 11.41: drseg Registers that Enable/Disable Trace from Breakpoint-Based Triggers.....	311
Table 11.42: FDC TAP Register Field Descriptions.....	316
Table 11.43: FDC Register Mapping	317
Table 11.44: FDC Access Control and Status Register Field Descriptions	317
Table 11.45: FDC Configuration Register Field Descriptions	318
Table 11.46: FDC Status Register Field Descriptions.....	319
Table 11.47: FDC Receive Register Field Descriptions.....	320
Table 11.49: FDTXn Address Decode	321
Table 11.48: FDC Transmit Register Field Descriptions.....	321
Table 12.1: Byte Access Within a Word.....	324
Table 13.1: Encoding of the Opcode Field.....	329
Table 13.2: Special Opcode Encoding of Function Field	329
Table 13.3: Special2 Opcode Encoding of Function Field	329
Table 13.4: Special3 Opcode Encoding of Function Field	330
Table 13.5: RegImm Encoding of rt Field	330
Table 13.6: COP2 Encoding of rs Field	330
Table 13.7: COP2 Encoding of rt Field When rs=BC2	330
Table 13.8: COP0 Encoding of rs Field	331
Table 13.9: COP0 Encoding of Function Field When rs=CO	331
Table 13.10: MIPS32 COP1 Encoding of rs Field.....	331
Table 13.11: MIPS32 COP1 Encoding of Function Field When rs=S	332
Table 13.12: MIPS32 COP1 Encoding of Function Field When rs=D	332
Table 13.13: MIPS32 COP1 Encoding of Function Field When rs=W or L.....	332
Table 13.14: MIPS32 COP1 Encoding of tf Bit When rs=S or D, Function=MOVCF.....	333
Table 13.15: Instruction Set.....	333
Table 14.1: 16-Bit Re-encoding of Frequent MIPS32 Instructions.....	364
Table 14.2: 16-Bit Re-encoding of Frequent MIPS32 Instruction Sequences.....	365
Table 14.3: Instruction-Specific Register Specifiers and Immediate Field Values	367
Table 14.4: 16-Bit Instruction General-Purpose Registers - \$2-\$7, \$16, \$17	368

Table 14.5: SB16, SH16, SW16 Source Registers - \$0, \$2-\$7, \$17	369
Table 14.6: 16-Bit Instruction Implicit General-Purpose Registers	369
Table 14.7: 16-Bit Instruction Special-Purpose Registers.....	370
Table 14.8: 32-bit Instructions introduced within microMIPS	370

Figures

Figure 1.1: M5100™ Processor Core Block Diagram	25
Figure 1.2: M5100™ Core Virtual Address Map	30
Figure 1.3: Address Translation During SRAM Access with FMT Implementation	31
Figure 1.4: Reference Design Block Diagram.....	34
Figure 1.5: FPU Pipeline.....	35
Figure 1.6: FDC Overview	38
Figure 1.7: cJTAG Support	39
Figure 2.1: M5100™ Core Pipeline Stages with high-performance MDU	41
Figure 2.2: M5100™ Core Pipeline Stages with area-efficient MDU	41
Figure 2.3: M5100™ Core Pipeline Stages	42
Figure 2.4: Multiply Pipeline	48
Figure 2.5: Multiply With Dependency From ALU	48
Figure 2.6: Multiply With Dependency From Load Hit	48
Figure 2.7: Multiply With Dependency From Load Miss	48
Figure 2.8: subtractMUL Bypassing Result to Integer Instructions	49
Figure 2.9: MDU Pipeline Flow During a 8-bit Divide (DIV) Operation	50
Figure 2.10: MDU Pipeline Flow During a 16-bit Divide (DIV) Operation	50
Figure 2.11: MDU Pipeline Flow During a 24-bit Divide (DIV) Operation	50
Figure 2.12: MDU Pipeline Flow During a 32-bit Divide (DIV) Operation	50
Figure 2.13: MDU Pipeline Behavior During Multiply Operations	52
Figure 2.14: MDU Pipeline Flow During a 32x16 Multiply Operation	53
Figure 2.15: MDU Pipeline Flow During a 32x32 Multiply Operation	54
Figure 2.16: High-Performance MDU Pipeline Flow During a 8-bit Divide (DIV) Operation	54
Figure 2.17: High-Performance MDU Pipeline Flow During a 16-bit Divide (DIV) Operation	54
Figure 2.18: High-Performance MDU Pipeline Flow During a 24-bit Divide (DIV) Operation	55
Figure 2.19: High-Performance MDU Pipeline Flow During a 32-bit Divide (DIV) Operation	55
Figure 2.20: M5100™ Area-Efficient MDU Pipeline Flow During a Multiply Operation	56
Figure 2.21: M5100™ Core Area-Efficient MDU Pipeline Flow During a Multiply Accumulate Operation	56
Figure 2.22: M5100™ Core Area-Efficient MDU Pipeline Flow During a Divide (DIV) Operation	57
Figure 2.23: IU Pipeline Branch Delay	58
Figure 2.24: IU Pipeline Data bypass	59
Figure 2.25: IU Pipeline M to E bypass	60
Figure 2.26: IU Pipeline A to E Data bypass	60
Figure 2.27: IU Pipeline Slip after a MFHI	61
Figure 2.28: Coprocessor 2 Interface Transactions	62
Figure 2.29: Instruction Cache Miss Slip	63
Figure 3.1: FPU Block Diagram	68
Figure 3.2: Single-Precision Floating-Point Format (S)	70
Figure 3.3: Double-Precision Floating-Point Format (D)	70
Figure 3.4: Word Fixed-Point Format (W)	73
Figure 3.5: Longword Fixed-Point Format (L)	73
Figure 3.6: Single Floating-Point or Word Fixed-Point Operand in an FPR	73
Figure 3.7: Double Floating-Point or Longword Fixed-Point Operand in an FPR	73
Figure 3.8: Effect of FPU Operations on the Format of Values Held in FPRs	75
Figure 3.9: FPU Word Load and Move-to Operations	76
Figure 3.10: FPU Doubleword Load and Move-to Operations	76
Figure 3.11: FIR Format	78

Figure 3.12: FCCR Format	79
Figure 3.13: FEXR Format	80
Figure 3.14: FENR Format	80
Figure 3.15: FCSR Format	81
Figure 3.16: FS/FO/FN Bits Influence on Multiply and Addition Results	85
Figure 3.17: Flushing to Nearest when Rounding Mode is Round to Nearest	86
Figure 3.18: FPU Pipeline	97
Figure 3.19: Arithmetic Pipeline Bypass Paths	99
Figure 4.1: MIPS32® DSP Module Control Register (DSPControl) Format	101
Figure 6.1: Address Translation During SRAM Access	108
Figure 6.2: M5100™ processor core Virtual Memory Map	110
Figure 6.3: User Mode Virtual Address Space	111
Figure 6.4: Kernel Mode Virtual Address Space	113
Figure 6.5: Debug Mode Virtual Address Space	115
Figure 6.6: FMT Memory Map (ERL=0) in the M5100™ Processor Core	117
Figure 6.7: FMT Memory Map (ERL=1) in the M5100™ Processor Core	118
Figure 7.1: Interrupt Generation for Vectored Interrupt Mode	126
Figure 7.2: Interrupt Generation for External Interrupt Controller Interrupt Mode	129
Figure 7.3: General Exception Handler (HW)	148
Figure 7.4: General Exception Servicing Guidelines (SW)	149
Figure 7.5: Reset, Soft Reset and NMI Exception Handling and Servicing Guidelines	150
Figure 8.1: Index Register Format	154
Figure 8.2: Random Register Format	155
Figure 8.3: EntryLo0, EntryLo1 Register Format	155
Figure 8.4: Context Register Format	157
Figure 8.5: UserLocal Register Format	158
Figure 8.6: PageMask Register Format	158
Figure 8.7: PageGrain Register Format	160
Figure 8.8: Wired and Random Entries in the TLB	162
Figure 8.9: Wired Register Format.....	162
Figure 8.10: HWREna Register Format	162
Figure 8.11: BadVAddr Register Format	163
Figure 8.12: BadInstr Register Format.....	164
Figure 8.13: BadInstrP Register Format	165
Figure 8.14: Count Register Format	165
Figure 8.15: EntryHi Register Format	166
Figure 8.16: GuestCtl1 Register Format	167
Figure 8.17: GuestCtl2 Register Format for non-EIC mode.....	168
Figure 8.18: GuestCtl2 Register Format for EIC mode.....	168
Figure 8.19: GuestCtl3 Register Format	172
Figure 8.20: Compare Register Format	172
Figure 8.21: GuestCtl0Ext Register Format.....	173
Figure 8.22: Status Register Format.....	175
Figure 8.23: IntCtl Register Format.....	180
Figure 8.24: SRSCtl Register Format	184
Figure 8.25: SRSSMap Register Format.....	187
Figure 8-26: View_IPL Register Format.....	188
Figure 8-27: SRSSMap Register Format	189
Figure 8.28: GuestCtl0 Register Format	189
Figure 8.29: GTOffset Register Format	196
Figure 8.30: Cause Register Format.....	197
Figure 8-31: View_RIPL Register Format	202
Figure 8-32: NestedExc Register Format	202

Figure 8.33: EPC Register Format	204
Figure 8.34: NestedEPC Register Format	204
Figure 8.35: PRId Register Format	205
Figure 8.36: EBase Register Format	207
Figure 8.37: CDMMBase Register Format.....	207
Figure 8.38: Config Register Format — Select 0	208
Figure 8.39: Config1 Register Format — Select 1	210
Figure 8.40: Config2 Register Format — Select 2	211
Figure 8-41: Config3 Register Format	212
Figure 8-42: Config4 Register Format	216
Figure 8-43: Config5 Register Format	218
Figure 8.44: Config7 Register Format	218
Figure 8.45: SecurityCtrl Register Format	220
Figure 8.46: SecuritySwPRNG Register Format	220
Figure 8.47: SecurityHwPRNG Register Format	221
Figure 8.48: Debug Register Format	221
Figure 8.49: TraceControl Register Format	225
Figure 8.50: TraceControl2 Register Format	227
Figure 8.51: User Trace Data1/User Trace Data2 Register Format	229
Figure 8.52: Trace BPC Register Format	230
Figure 8.53: Debug2 Register Format	231
Figure 8.54: DEPC Register Format	232
Figure 8.55: Performance Counter Control Register	233
Figure 8.56: Performance Counter Count Register	239
Figure 8.57: ErrCtl Register Format	239
Figure 8.58: CacheErr Register (Primary Caches)	240
Figure 8.59: ErrorEPC Register Format	241
Figure 8.60: DeSave Register Format	242
Figure 8-61: KScratchn Register Format	242
Figure 11.1: DCR Register Format	249
Figure 11.2: IBS Register Format	260
Figure 11.3: IBAn Register Format	260
Figure 11.4: IBMn Register Format	261
Figure 11.5: IBASIDn Register Format	261
Figure 11.6: IBCn Register Format	262
Figure 11.7: IBCCn Register Format	263
Figure 11.8: IBPCn Register Format	264
Figure 11.9: DBS Register Format	265
Figure 11.10: DBAn Register Format	265
Figure 11.11: DBMn Register Format	266
Figure 11.12: DBASIDn Register Format	266
Figure 11.13: DBCn Register Format	266
Figure 11.14: DBVn Register Format	268
Figure 11.15: DBCCn Register Format	269
Figure 11.16: DBPCn Register Format	270
Figure 11.17: DVM Register Format	270
Figure 11.18: CBTC Register Format	271
Figure 11.19: PrCndA Register Format	272
Figure 11.20: STCtl Register Format	274
Figure 11.21: STCnt Register Format	275
Figure 11.22: TAP Controller State Diagram	281
Figure 11.23: Concatenation of the EJTAG Address, Data and Control Registers	285
Figure 11.24: TDI to TDO Path When in Shift-DR State and FASTDATA Instruction is Selected	286

Figure 11.25: Device Identification Register Format	287
Figure 11.26: Implementation Register Format	288
Figure 11.27: EJTAG Control Register Format	289
Figure 11.28: Endian Formats for the PAD Register	296
Figure 11.29: Fastdata Register Format	296
Figure 11.30: Trace Logic Overview	305
Figure 11.31: Control/Status Register.....	307
Figure 11.32: ITCBTW Register Format	309
Figure 11.33: ITCBRDP Register Format	310
Figure 11.34: ITCBWRP Register Format.....	310
Figure 11.35: PCSAMPLE TAP Register Format (MIPS32)	312
Figure 11.36: Fast Debug Channel Buffer Organization	315
Figure 11.37: FDC TAP Register Format.....	316
Figure 11.38: FDC Access Control and Status Register.....	317
Figure 11.39: FDC Configuration Register.....	318
Figure 11.40: FDC Status Register	319
Figure 11.41: FDC Receive Register	320
Figure 11.42: FDC Transmit Register	320
Figure 11.43: cJTAG Interface	321
Figure 12.1: Instruction Formats	323
Figure 14.1: 16-Bit Instruction Formats.....	360
Figure 14.2: 32-Bit Instruction Formats.....	361
Figure 14.3: Immediate Fields within 32-Bit Instructions.....	361

Introduction to the MIPS32® M5100™ Processor Core

The MIPS32® M5100™ core from MIPS Technologies is a high-performance, low-power, 32-bit MIPS RISC processor core intended for custom system-on-silicon applications. The core is designed for semiconductor manufacturing companies, ASIC developers, and system OEMs who want to rapidly integrate their own custom logic and peripherals with a high-performance RISC processor. The M5100 core is fully synthesizable to allow maximum flexibility; it is highly portable across processes and can easily be integrated into full system-on-silicon designs. This allows developers to focus their attention on end-user specific characteristics of their product.

The M5100 core is especially well-suited for microcontrollers and applications that have real-time requirements with a high level of performance efficiency and security requirements.

The M5100 core implements the MIPS Architecture Release-5 (“R5” incorporating enhanced functionality for next generation MIPS-Based™ products) in a 5-stage pipeline. It includes support for the microMIPS™ ISA, an Instruction Set Architecture with optimized MIPS32 16-bit and 32-bit instructions that provides a significant reduction in code size with a performance equivalent to MIPS32. The M5100 core is a successor to the microAptiv™ UC, designed from the same microarchitecture, including the Microcontroller Application-Specific Extension (MCU™ ASE), enhanced interrupt handling, lower interrupt latency, a reference design of an optimized interface for flash memory and built-in native AMBA®-3 AHB-Lite Bus Interface Unit (BIU), with additional power saving, security, debug, and profiling features. In addition, the M5100 core includes the MIPS Architecture Virtualization Module that enables virtualization of operating systems, which provides a scalable, trusted, and secure execution environment.

The M5100 core has an option to include the MIPS Architecture DSP Module Revision 2 that provides digital signal processing capabilities, with support for a number of powerful data processing operations. An optional IEEE 754 compliant Floating-Point Unit (FPU) provides both single and double precision instructions.

The M5100 core is cacheless; in lieu of caches, it includes a simple interface to SRAM-style devices. This interface may be configured for independent instruction and data devices or combined into a unified interface. The SRAM interface allows deterministic latency to memory, while still maintaining high performance.

A distinguishing characteristic of the M5100 family is the inclusion of a configurable MIPS DSP Module. The MIPS DSP Module provides support for a number of powerful data processing operations. It includes instructions for executing fractional arithmetic (Q15/Q31) and saturating arithmetic. Additionally, for smaller data sizes, SIMD operations are supported, allowing 2x16b or 4x8b operations to occur simultaneously. Another feature of the DSP Module is the inclusion of additional HI/LO accumulator registers that improve the parallelization of independent accumulation routines.

The core includes one of two different Multiply/Divide Unit (MDU) implementations, selectable at build-time if the DSP Module is not configured in, allowing the user to trade-off performance and area for integer multiply and divide operations. The high-performance MDU option implements single-cycle multiply and multiply-accumulate (MAC) instructions that enable DSP algorithms to be performed efficiently. It allows 32-bit x 16-bit MAC instructions to be issued every cycle, while a 32-bit x 32-bit MAC instruction can be issued every other cycle. The area-efficient MDU option handles multiplies with a one-bit-per-clock iterative algorithm.

If the core is configured with the DSP Module, the Multiply-Divide Unit (MDU) is fully pipelined and supports a maximum issue rate of one 32x32 multiply (MUL/MULT/MULTU), multiply-add (MADD/MADDU), or multiply-subtract (MSUB/MSUBU) operations per clock.

The MMU consists of a simple Fixed Mapping Translation (FMT) mechanism, for applications that do not require the full capabilities of a Translation Lookaside Buffer- (TLB-) based MMU available on other MIPS cores.

The basic Enhanced JTAG (EJTAG) features provide CPU run control with stop, single-stepping and re-start, and with software breakpoints using the SDBBP instruction. Additional EJTAG features such as instruction and data virtual address hardware breakpoints, complex hardware breakpoints, connection to an external EJTAG probe through the Test Access Port (TAP), and PC/Data tracing, may be included as an option.

1.1 Features

- 5-stage pipeline
- 32-bit Address and Data Paths
- MIPS32 Instruction Set Architecture
- MIPS32 Enhanced Architecture Features
 - Vectored interrupts and support for external interrupt controller
 - Programmable exception vector base
 - Atomic interrupt enable/disable
 - GPR shadow registers (one, three, seven, or fifteen additional shadows can be optionally added to minimize latency for interrupt handlers)
 - Bit field manipulation instructions
- microMIPS Instruction Set Architecture
 - microMIPS ISA is a build-time configurable option that reduces code size over MIPS32, while maintaining MIPS32 performance.
 - Combining both 16-bit and 32-bit opcodes, microMIPS supports all MIPS32 instructions (except branch-likely instructions) with new optimized encoding. Frequently used MIPS32 instructions are available as 16-bit instructions.
 - Added fifteen new 32-bit instructions and thirty-nine 16-bit instructions.
 - Stack pointer implicit in instruction.
 - MIPS32 assembly and ABI-compatible.
 - Supports MIPS architecture Modules and User-defined Instructions (UDIs).
- MCU™ ASE

- Increases the number of interrupt hardware inputs from 6 to 8 for Vectored Interrupt (VI) mode, and from 63 to 255 for External Interrupt Controller (EIC) mode.
- Separate priority and vector generation. 16-bit vector address is provided.
- Hardware assist combined with the use of Shadow Register Sets to reduce interrupt latency during the prologue and epilogue of an interrupt.
- An interrupt return with automated interrupt epilogue handling instruction (IRET) improves interrupt latency.
- Supports optional interrupt chaining.
- Two memory-to-memory atomic read-modify-write instructions (ASET and ACLR) eases commonly used semaphore manipulation in microcontroller applications. Interrupts are automatically disabled during the operation to maintain coherency.
- Memory Management Unit
 - Simple Fixed Mapping Translation (FMT) mechanism
- Virtualization Module Support
 - Support for the Virtualization Module using “trap and emulate” techniques.
 - Two new operating modes: privileged (hypervisor) and unprivileged (guest)
 - System resources controlled through privileged instructions while executing in privileged mode
 - Execution of a privileged instruction in unprivileged mode will cause an exception (trap), returning control to privileged-mode software
 - Address translation is performed on the entire address space when in unprivileged mode
 - Supports 1 to 7 guests
 - Supports RPU (Root Protection Unit) memory management unit
- Simple SRAM-Style Interface
 - Cacheless operation enables deterministic response and reduces die-size
 - 32-bit address and data; input byte-enables enable simple connection to narrower devices
 - Single or multi-cycle latencies
 - Configuration option for dual or unified instruction/data interfaces
 - Redirection mechanism on dual I/D interfaces permits D-side references to be handled by I-side
 - Transactions can be aborted
- Reference Design

Introduction to the MIPS32® M5100™ Processor Core

- A typical SRAM reference design is provided.
- An AHB-Lite BIU reference design is provided between the SRAM interface and AHB-Lite Bus.
- An optimized interface for slow memory (Flash) access using prefetch buffer scheme is provided.
- Parity Support
 - The ISRAM and DSRAM support optional parity detection.
- MIPS DSP Module (Revision 2.0)
 - Support for MAC operations with 3 additional pairs of Hi/Lo accumulator registers (Ac0 - Ac3)
 - Fractional data types (Q15, Q31) with rounding support
 - Saturating arithmetic with overflow handling
 - SIMD instructions operate on 2x16-bit or 4x8-bit operands simultaneously
 - Separate MDU pipeline with full-sized hardware multiplier to support back-to-back operations
 - The DSP Module is build-time configurable.
- Multiply/Divide Unit (area-efficient configuration without DSP)
 - 32 clock latency on multiply
 - 34 clock latency on multiply-accumulate
 - 33-35 clock latency on divide (sign-dependent)
- Multiply/Divide Unit (high-performance configuration without DSP)
 - Maximum issue rate of one 32x16 multiply per clock via on-chip 32x16 hardware multiplier array.
 - Maximum issue rate of one 32x32 multiply every other clock
 - Early-in iterative divide. Minimum 11 and maximum 34 clock latency (dividend (*rs*) sign extension-dependent)
- Multiply/Divide Unit (with DSP configuration)
 - Maximum issue rate of one 32x32 multiply per clock via on-chip 32x32 hardware multiplier array
 - Maximum issue rate of one 32x32 multiply every clock
 - Early-in iterative divide. Minimum 12 and maximum 38 clock latency (dividend (*rs*) sign extension-dependent)
- Floating Point Unit (FPU)
 - 1985 IEEE-754 compliant Floating Point Unit.

- Supports single and double precision datatypes
- 2008 IEEE-754 compatibility control for NaN handling and Abs/Neg instructions
- Runs at 1:1 core/FPU clock ratio.
- The FPU is build-time configurable.
- CorExtend® User-Defined Instruction Set Extensions
 - Allows user to define and add instructions to the core at build time
 - Maintains full MIPS32 compatibility
 - Supported by industry-standard development tools
 - Single or multi-cycle instructions
- Multi-Core Support
 - External lock indication enables multi-processor semaphores based on LL/SC instructions
 - External sync indication allows memory ordering
 - Debug support includes cross-core triggers
- Coprocessor 2 interface
 - 32-bit interface to an external coprocessor
- Anti-Tamper Security Features
 - Injection of random pipeline slips controlled via CP0 register
 - Fast general-purpose 32-bit pseudo random generator accessible via CP0 register
 - Complex pseudo random generator for hardware use controlled via user-defined seed
- Power Control
 - Minimum frequency: 0 MHz
 - Power-down mode (triggered by WAIT instruction)
 - Support for software-controlled clock divider
 - Support for extensive use of local gated clocks
- EJTAG Debug/Profiling and iFlowtrace™ Mechanism
 - CPU control with start, stop, and single stepping
 - Virtual instruction and data address/value breakpoints

Introduction to the MIPS32® M5100™ Processor Core

- Hardware breakpoint supports both address match and address range triggering
- Optional simple hardware breakpoints on virtual addresses; 8I/4D, 6I/2D, 4I/2D, 2I/1D breakpoints, or no breakpoints
- Optional complex hardware breakpoints with 8I/4D, 6I/2D simple breakpoints
- TAP controller is chainable for multi-CPU debug
- Supports EJTAG (IEEE 1149.1) and compatible with cJTAG 2-wire (IEEE 1149.7) extension protocol
- Cross-CPU breakpoint support
- iFlowtrace support for real-time instruction PC and special events
- PC and/or load/store address sampling for profiling
- Performance Counters
- Support for Fast Debug Channel (FDC)
- SecureDebug
 - An optional feature that disables access via EJTAG in an untrusted environment
- Testability
 - Full scan design achieves test coverage in excess of 99% (dependent on library and configuration options)

1.2 M5100™ Core Block Diagram

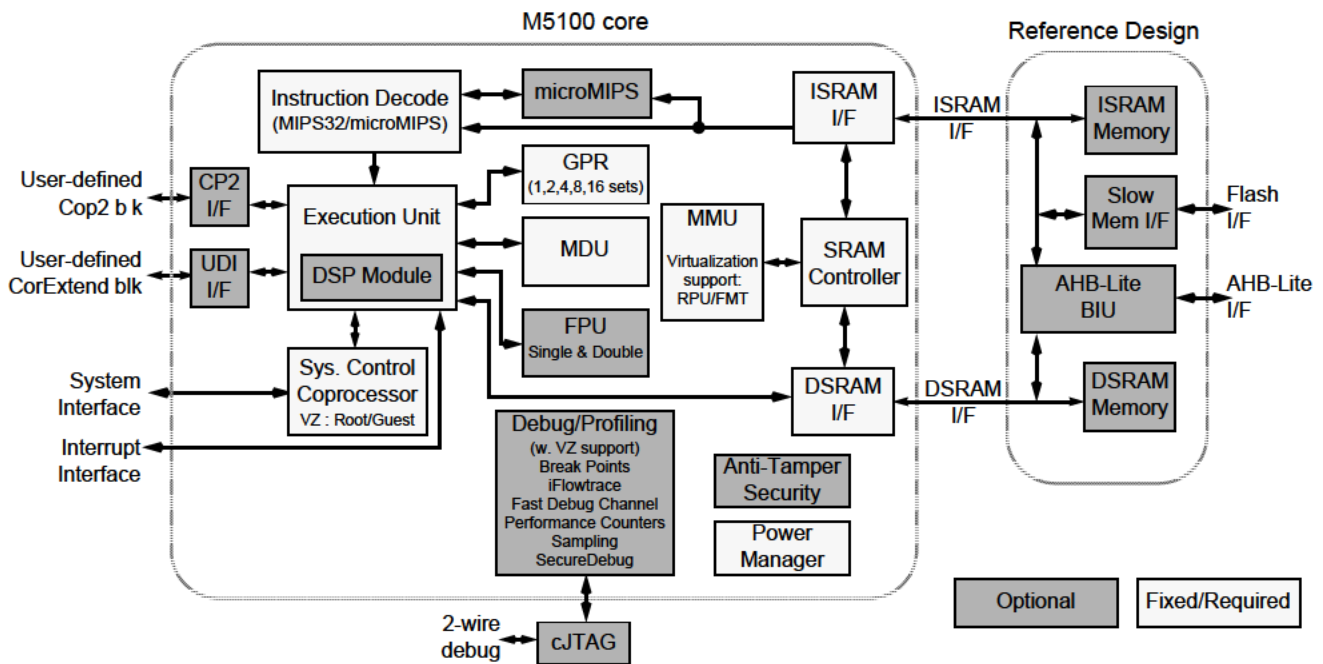
The M5100 core contains both required and optional blocks, as shown in the block diagram in [Figure 1.1](#). Required blocks are the lightly shaded areas of the block diagram and are always present in any core implementation. Optional blocks may be added to the base core, depending on the needs of a specific implementation. The required blocks are as follows:

- Instruction Decode
- Execution Unit
- General Purposed Registers (GPR)
- Multiply/Divide Unit (MDU)
- System Control Coprocessor (CP0)
- Memory Management Unit (MMU)
- I/D SRAM Interfaces
- Power Management

Optional blocks include:

- Configurable instruction decoder supporting three ISA modes: MIPS32-only, MIPS32 and microMIPS, or microMIPS-only
- DSP (integrated with MDU)
- Floating-point Unit (FPU)
- Virtualization support (integrated with MMU)
- Reference Design of I/D-SRAM, BIU, Slow Memory Interface
- Anti-Tamper Security Block
- Coprocessor 2 interface
- CorExtend® User-Defined Instruction (UDI) interface
- Debug/Profiling with Enhanced JTAG (EJTAG) Controller, Break points, Sampling, Performance counters, Fast Debug Channel, and iFlowtrace logic

Figure 1.1 M5100™ Processor Core Block Diagram



1.2.1 Required Logic Blocks

The following subsections describe the required logic blocks of the M5100 processor core.

1.2.1.1 Execution Unit

The M5100™ core execution unit implements a load/store architecture with single-cycle ALU operations (logical, shift, add, subtract) and an autonomous multiply/divide unit.

The execution unit includes:

- Arithmetic Logic Unit (ALU) for performing arithmetic and bitwise logical operations. Shared adder for arithmetic operations, load/store address calculation, and branch target calculation.
- Address unit for calculating the next PC and next fetch address selection muxes.
- Load Aligner.
- Shifter and Store Aligner.
- Branch condition comparator.
- Trap condition comparator.
- Bypass muxes to advance result between two adjacent instructions with data dependency.
- Leading Zero/One detect unit for implementing the CLZ and CLO instructions.
- Actual execution of the Atomic Instructions defined in the MCU ASE.
- A separate DSP ALU and Logic block for performing part of DSP Module instructions, such as arithmetic/shift/compare operations, if the DSP function is configured.

1.2.1.2 General Purposed Register (GPR) Shadow Registers

The M5100™ core contains thirty-two 32-bit general-purpose registers used for integer operations and address calculation. Optionally, one, three, seven or fifteen additional register file shadow sets (each containing thirty-two registers) can be added to minimize context switching overhead during interrupt/exception processing. The register file consists of two read ports and one write port and is fully bypassed to minimize operation latency in the pipeline.

1.2.1.3 Multiply/Divide Unit (MDU)

The M5100™ core includes a multiply/divide unit (MDU) that contains a separate, dedicated pipeline for integer multiply/divide operations, and DSP Module multiply instructions. This pipeline operates in parallel with the integer unit (IU) pipeline and does not stall when the IU pipeline stalls. This allows the long-running MDU operations to be partially masked by system stalls and/or other integer unit instructions.

The MIPS architecture defines that the result of a multiply or divide operation be placed in a pair (without DSP enabled) or one of 4 pairs (with DSP enabled) of *HI* and *LO* registers. Using the Move-From-HI (MFHI) and Move-From-LO (MFLO) instructions, these values can be transferred to the general-purpose register file.

There are three configuration options for the MDU: 1) a full 32x32 multiplier block; 2) a higher performance 32x16 multiplier block; 3) an area-efficient iterative multiplier block. Option 2 and 3 are available if the DSP configuration option is disabled. If the DSP configuration option is enabled, option 1 is the default. The selection of the MDU style allows the implementor to determine the appropriate performance and area trade-off for the application.

MDU with 32x32 DSP Multiplier

With the DSP configuration option enabled, the MDU supports execution of one 16x16, 32x16, or 32x32 multiply or multiply-accumulate operation every clock cycle with the built in 32x32 multiplier array. The multiplier is shared with DSP Module operations.

The MDU also implements various shift instructions operating on the HI/LO register and multiply instructions as defined in the DSP Module. It supports all the data types required for this purpose and includes three extra HI/LO registers as defined by the Module.

MDU with 32x16 High-Performance Multiplier

The high-performance MDU consists of a 32x16 Booth-recoded multiplier, a pair of result/accumulation registers (*HI* and *LO*), a divide state machine, and the necessary multiplexers and control logic. The first number shown ('32' of 32x16) represents the *rs* operand. The second number ('16' of 32x16) represents the *rt* operand. The M5100 core only checks the value of the *rt* operand to determine how many times the operation must pass through the multiplier. The 16x16 and 32x16 operations pass through the multiplier once. A 32x32 operation passes through the multiplier twice.

The MDU supports execution of one 16x16 or 32x16 multiply or multiply-accumulate operation every clock cycle; 32x32 multiply operations can be issued every other clock cycle. Appropriate interlocks are implemented to stall the issuance of back-to-back 32x32 multiply operations. The multiply operand size is automatically determined by logic built into the MDU.

MDU with Area-Efficient Option

With the area-efficient option, multiply and divide operations are implemented with a simple 1-bit-per-clock iterative algorithm. Any attempt to issue a subsequent MDU instruction while a multiply/divide is still active causes an MDU pipeline stall until the operation is completed.

Regardless of the multiplier array implementation, divide operations are implemented with a simple 1-bit-per-clock iterative algorithm. An early-in detection checks the sign extension of the dividend (*rs*) operand. If *rs* is 8 bits wide, 23 iterations are skipped. For a 16-bit-wide *rs*, 15 iterations are skipped, and for a 24-bit-wide *rs*, 7 iterations are skipped. Any attempt to issue a subsequent MDU instruction while a divide is still active causes an IU pipeline stall until the divide operation has completed.

1.2.1.4 System Control Coprocessor (CP0)

In the MIPS architecture, CP0 is responsible for the virtual-to-physical address translation, the exception control system, the processor's diagnostics capability, the operating modes (kernel, user, and debug), and whether interrupts are enabled or disabled. Configuration information, such as presence of build-time options like microMIPS, CorExtend Module or Coprocessor 2 interface, is also available by accessing the CP0 registers.

Coprocessor 0 also contains the logic for identifying and managing exceptions. Exceptions can be caused by a variety of sources, including boundary cases in data, external events, or program errors.

Interrupt Handling

The M5100 core includes support for eight hardware interrupt pins, two software interrupts, and a timer interrupt. These interrupts can be used in any of three interrupt modes, as defined by Release 2 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, and to assign a GPR shadow set for use during interrupt processing. The presence of this mode is denoted by the *VInt* bit in the *Config3* register. This mode is architecturally optional; but it is always present on the M5100 core, so the *VInt* bit will always read as a 1 for the M5100 core.
- External Interrupt Controller (EIC) mode, which redefines the way in which interrupts are handled to provide full support for an external interrupt controller handling prioritization and vectoring of interrupts. The presence of this mode denoted by the *VEIC* bit in the *Config3* register. Again, this mode is architecturally optional. On the M5100 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.

The reset state of the processor is interrupt compatibility mode, such that a processor supporting Release 2 of the Architecture, the M5100 core for example, is fully compatible with implementations of Release 1 of the Architecture.

VI or EIC interrupt modes can be combined with the optional shadow registers to specify which shadow set should be used on entry to a particular vector. The shadow registers further improve interrupt latency by avoiding the need to save context when invoking an interrupt handler.

In the M5100 core, interrupt latency is reduced by:

- Speculative interrupt vector prefetching during the pipeline flush.
- Interrupt Automated Prologue (IAP) in hardware: Shadow Register Sets remove the need to save GPRs, and IAP removes the need to save specific Control Registers when handling an interrupt.
- Interrupt Automated Epilogue (IAE) in hardware: Shadow Register Sets remove the need to restore GPRs, and IAE removes the need to restore specific Control Registers when returning from an interrupt.
- Allow interrupt chaining. When servicing an interrupt and interrupt chaining is enabled, there is no need to return from the current Interrupt Service Routine (ISR) if there is another valid interrupt pending to be serviced. The control of the processor can jump directly from the current ISR to the next ISR without IAE and IAP.

GPR Shadow Registers

The MIPS32 Architecture optionally removes the need to save and restore GPRs on entry to high-priority interrupts or exceptions, and to provide specified processor modes with the same capability. This is done by introducing multiple copies of the GPRs, called *shadow sets*, and allowing privileged software to associate a shadow set with entry to kernel mode via an interrupt vector or exception. The normal GPRs are logically considered shadow set zero.

The number of GPR shadow sets is a build-time option. The M5100 core allows 1 (the normal GPRs), 2, 4, 8, or 16 shadow sets. The highest number actually implemented is indicated by the *SRSC#HSS* field. If this field is zero, only the normal GPRs are implemented.

Shadow sets are new copies of the GPRs that can be substituted for the normal GPRs on entry to kernel mode via an interrupt or exception. When a shadow set is bound to a kernel-mode entry condition, references to GPRs operate exactly as one would expect, but they are redirected to registers that are dedicated to that condition. Privileged software may need to reference all GPRs in the register file, even specific shadow registers that are not visible in the current mode, and the RDPGPR and WRPGPR instructions are used for this purpose. The *CSS* field of the *SRSC#* register provides the number of the current shadow register set, and the *PSS* field of the *SRSC#* register provides the number of the previous shadow register set that was current before the last exception or interrupt occurred.

If the processor is operating in VI interrupt mode, binding of a vectored interrupt to a shadow set is done by writing to the *SRSMap* register. If the processor is operating in EIC interrupt mode, the binding of the interrupt to a specific

shadow set is provided by the external interrupt controller and is configured in an implementation-dependent way. Binding of an exception or non-vectored interrupt to a shadow set is done by writing to the *ESS* field of the *SRSCtl* register. When an exception or interrupt occurs, the value of *SRSCtl_{CSS}* is copied to *SRSCtl_{PSS}*, and *SRSCtl_{CSS}* is set to the value taken from the appropriate source. On an ERET, the value of *SRSCtl_{PSS}* is copied back into *SRSCtl_{CSS}* to restore the shadow set of the mode to which control returns.

Refer to [Chapter 8, “CP0 Registers of the M5100™ Core” on page 151](#) for more information on the CP0 registers. Refer to [Chapter 11, “EJTAG Debug Support in the M5100™ Core” on page 248](#) for more information on EJTAG debug registers.

1.2.1.5 Memory Management Unit (MMU)

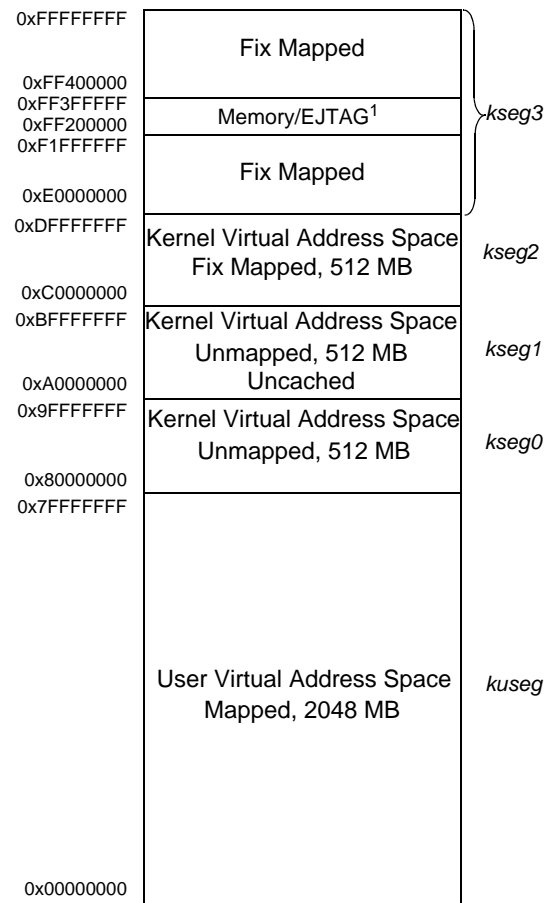
Modes of Operation

The M5100 core implements the Virtualization Module by supporting five modes of operation:

- *Guest-user mode* is most often used for applications programs.
- *Guest-kernel mode* is typically used for handling exceptions and operating-system kernel functions, including CP0 management and I/O device accesses.
- *Root-user mode* is used to run hypervisor (Virtual Machine Monitor (VMM)) user software.
- *Root-kernel mode* is used to run hypervisor (Virtual Machine Monitor (VMM)) kernel software.
- *Debug mode* is used during system bring-up and software development. Refer to the EJTAG section for more information on debug mode.

[Figure 1.2](#) shows the virtual address map of the MIPS Architecture.

Figure 1.2 M5100™ Core Virtual Address Map



1. This space is mapped to memory in user or kernel mode, and by the EJTAG module in debug mode.

Memory Management Unit (MMU) in Guest Modes

The M5100 core contains a simple Fixed Mapping Translation (FMT) MMU that interfaces between the execution unit and the SRAM controller.

- Fixed Mapping Translation (FMT)

A FMT is smaller and simpler than the full Translation Lookaside Buffer (TLB) style MMU found in other MIPS cores. Like a TLB, the FMT performs virtual-to-physical address translation and provides attributes for the different segments. Those segments that are unmapped in a TLB implementation (kseg0 and kseg1) are translated identically by the FMT.

Memory Management Unit (MMU) in Root Modes

The M5100 core provides a Root Protection Unit (RPU) MMU for guest virtual-address protection.

- Root Protection Unit (RPU)

An RPU-based MMU is similar to an RTL in that it consists of an 8, 16, or 32 dual-entry fully associative look-aside buffer, but the segments that are unmapped in a TLB implementation (kseg0 and kseg1) are translated identically by the RPU.

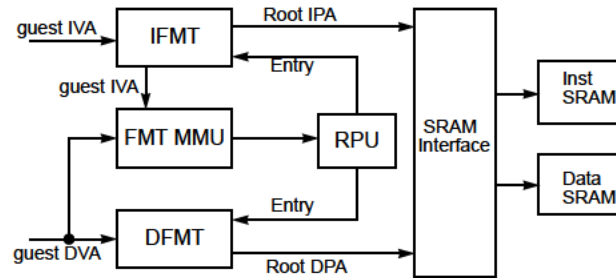
The RPU allows guest pages to be protected by a read-inhibit, execute-inhibit, and write-protection attribute.

When an instruction address is calculated, the virtual address is first compared to the contents of the 4-entry IFMT . If the address is not found, the virtual address is fix-mapped and validated by the RPU. If the address is not found in the RPU, a refill exception is taken.

When a data address is calculated, the virtual address is compared to the 4-entry DFMT . If the address is not found, the virtual address is fix-mapped and validated by the RPU. If the address is not found in the RPU, a refill exception is taken.

Figure 1.3 shows how the FMT is implemented in the M5100 core.

Figure 1.3 Address Translation During SRAM Access with FMT Implementation



1.2.1.6 SRAM Interface Controller

Instead of caches, the M5100 core contains an interface to SRAM-style memories that can be tightly coupled to the core. This permits deterministic response time with less area than is typically required for caches. The SRAM interface includes separate uni-directional 32-bit buses for address, read data, and write data.

Dual or Unified Interfaces

The SRAM interface includes a build-time option to select either dual or unified instruction and data interfaces.

The dual interface enables independent connection to instruction and data devices. It generally yields the highest performance, because the pipeline can generate simultaneous I and D requests, which are then serviced in parallel.

For simpler or cost-sensitive systems, it is also possible to combine the I and D interfaces into a common interface that services both types of requests. If I and D requests occur simultaneously, priority is given to the D side.

Back-stalling

Typically, read and write transactions will complete in a single cycle. However, if multi-cycle latency is desired, the interface can be stalled to allow connection to slower devices.

Redirection

When the dual I/D interface is present, a mechanism exists to divert D-side references to the I-side, if desired. The mechanism can be explicitly invoked for any other D-side references, as well. When the *DS_Redir* signal is asserted, a D-side request is diverted to the I-side interface in the following cycle, and the D-side will be stalled until the transaction is completed.

Transaction Abort

The core may request a transaction (fetch/load/store/sync) to be aborted. This is particularly useful in case of interrupts. Because the core does not know whether transactions are re-startable, it cannot arbitrarily interrupt a request that has been initiated on the SRAM interface. However, cycles spent waiting for a multi-cycle transaction to complete can directly impact interrupt latency. In order to minimize this effect, the interface supports an abort mechanism. The core requests an abort whenever an interrupt is detected and a transaction is pending (abort of an instruction fetch may also be requested in other cases). The external system logic can choose to acknowledge or to ignore the abort request.

Connecting to Narrower Devices

The instruction and data read buses are always 32 bits in width. To facilitate connection to narrower memories, the SRAM interface protocol includes input byte-enables that can be used by system logic to signal validity as partial read data becomes available. The input byte-enables conditionally register the incoming read data bytes within the core, and thus eliminate the need for external registers to gather the entire 32 bits of data. External muxes are required to redirect the narrower data to the appropriate byte lanes.

Lock Mechanism

The SRAM interface includes a protocol to identify a locked sequence, and is used in conjunction with the LL/SC atomic read-modify-write semaphore instructions.

Sync Mechanism

The interface includes a protocol that externalizes the execution of the SYNC instruction. External logic might choose to use this information to enforce memory ordering between various elements in the system.

External Call Indication

The instruction fetch interface contains signals that indicate that the core is fetching the target of a subroutine call-type instruction such as JAL or BAL. At some point after a call, there will typically be a return to the original code sequence. If a system prefetches instructions, it can make use of this information to save instructions that were prefetched and are likely to be executed after the return.

1.2.1.7 Power Management

The M5100 core offers a number of power management features, including low-power design, active power management, and power-down modes of operation. The core is a static design that supports slowing or halting the clocks, which reduces system power consumption during idle periods.

The M5100 core provides two mechanisms for system-level low-power support:

- Register-controlled power management
- Instruction-controlled power management

Register-Controlled Power Management

The *RP* bit in the CP0 *Status* register provides a software mechanism for placing the system into a low-power state. The state of the *RP* bit is available externally via the *SI_RP* signal. The external agent then decides whether to place the device in a low-power mode, such as reducing the system clock frequency.

Three additional bits, *StatusEXL*, *StatusERL*, and *DebugDM* support the power management function by allowing the user to change the power state if an exception or error occurs while the M5100 core is in a low-power state. Depending on what type of exception is taken, one of these three bits will be asserted and reflected on the *SI_EXL*, *SI_ERL*, or *EJ_DebugM* outputs. The external agent can look at these signals and determine whether to leave the low-power state to service the exception.

The following four power-down signals are part of the system interface and change state as the corresponding bits in the CP0 registers are set or cleared:

- The *SI_RP* signal represents the state of the *RP* bit (27) in the CP0 *Status* register.
- The *SI_EXL* signal represents the state of the *EXL* bit (1) in the CP0 *Status* register.
- The *SI_ERL* signal represents the state of the *ERL* bit (2) in the CP0 *Status* register.
- The *EJ_DebugM* signal represents the state of the *DM* bit (30) in the CP0 *Debug* register.

Instruction-Controlled Power Management

The second mechanism for invoking power-down mode is by executing the WAIT instruction. When the WAIT instruction is executed, the internal clock is suspended; however, the internal timer and some of the input pins (*SI_Int[5:0]*, *SI_NMI*, *SI_Reset*, and *SI_ColdReset*) continue to run. When the CPU is in instruction-controlled power management mode, any interrupt, NMI, or reset condition causes the CPU to exit this mode and resume normal operation.

The M5100 core asserts the *SI_Sleep* signal, which is part of the system interface bus, whenever the WAIT instruction is executed. The assertion of *SI_Sleep* indicates that the clock has stopped and the M5100 core is waiting for an interrupt.

Local clock gating

The majority of the power consumed by the M5100 core is in the clock tree and clocking registers. The core has support for extensive use of local gated clocks. Power-conscious implementors can use these gated clocks to significantly reduce power consumption within the core.

Refer to [Chapter 10, “Power Management of the M5100™ Core” on page 246](#) for more information on power management.

1.2.2 Optional Logic Blocks

The core consists of the following optional logic blocks as shown in the block diagram in [Figure 1.1](#).

1.2.2.1 Reference Design

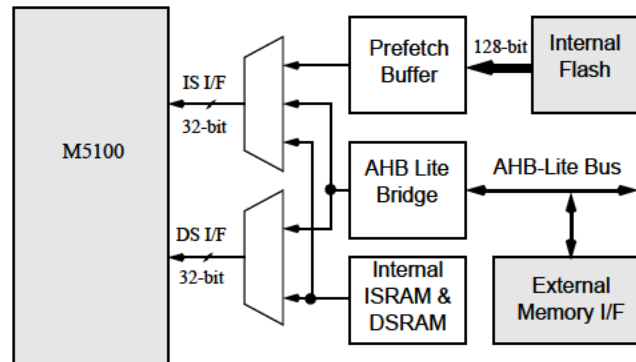
The M5100 core contains a reference design that shows a typical usage of the core with:

- Dual I-SRAM and D-SRAM interface with fast memories (i.e., SRAM) for instruction and data storage.
- Optimized interface for slow memory (i.e., Flash memory) access by having a prefetch buffer and a wider Data Read bus (i.e., *IS_RData[127:0]*) to speed up I-Fetch performance.
- AHB-lite bus interface to the system bus if the memory accesses are outside the memory map for the SRAM and Flash regions. AHB-Lite is a subset of the AHB bus protocol that supports a single bus master. The interface

shares the same 32-bit Read and Write address bus and has two unidirectional 32-bit buses for Read and Write data.

The reference design is optional and can be modified by the user to better fit the SOC design requirement.

Figure 1.4 Reference Design Block Diagram.



1.2.2.2 microMIPS™ ISA

The M5100 core supports the microMIPS ISA, which contains all MIPS32 ISA instructions (except for branch-likely instructions) in a new 32-bit encoding scheme, with some of the commonly used instructions also available in 16-bit encoded format. This ISA improves code density through the additional 16-bit instructions while maintaining a performance similar to MIPS32 mode. In microMIPS mode, 16-bit or 32-bit instructions will be fetched and recoded to legacy MIPS32 instruction opcodes in the pipeline's I stage, so that the M5100 core can have the same microAptiv UC microarchitecture. Because the microMIPS instruction stream can be intermixed with 16-bit halfword or 32-bit word size instructions on halfword or word boundaries, additional logic is in place to address the word misalignment issues, thus minimizing performance loss.

1.2.2.3 DSP Module

The M5100 core implements an optional DSP Module to benefit a wide range of DSP, Media, and DSP-like algorithms. The DSP module is highly integrated with the Execution Unit and the MDU in order to share common logic and to include support for operations on fractional data types, saturating arithmetic, and register SIMD operations. Fractional data types Q15 and Q31 are supported. Register SIMD operations can perform up to four simultaneous add, subtract, or shift operations and two simultaneous multiply operations.

In addition, the DSP Module includes some key features that efficiently address specific problems often encountered in DSP applications. These include, for example, support for complex multiply, variable-bit insert and extract, and implementation and use of virtual circular buffers. The extension also makes available three additional sets of HI-LO accumulators to better facilitate common accumulate functions such as filter operation and convolutions.

1.2.2.4 Floating Point Unit (FPU)

The M5100 core Floating Point Unit (FPU) implements the MIPS Instruction Set Architecture for floating-point computation. 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 programmed to have thirty-two 32-bit or 64-bit floating-point registers used for floating point operations.

The performance is optimized for single-precision formats. Most instructions have one FPU cycle throughput and four FPU cycle latency. The FPU implements the 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 are improved.

IEEE denormalized input operands and results are supported by hardware for some instructions. IEEE denormalized 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 *FCCR* register, and use of this mode is recommended for best performance when denormalized results are generated.

The FPU has a separate pipeline for floating point instruction execution. This pipeline operates in parallel with the integer core pipeline and does not stall when the integer pipeline stalls. This allows long-running FPU operations, such as divide or square root, to be partially masked by system stalls and/or other integer-unit instructions. Arithmetic instructions are always dispatched and completed in-order, but loads and stores can complete out-of-order. The exception model is ‘precise’ at all times. The FPU is also denoted as “Coprocessor 1”.

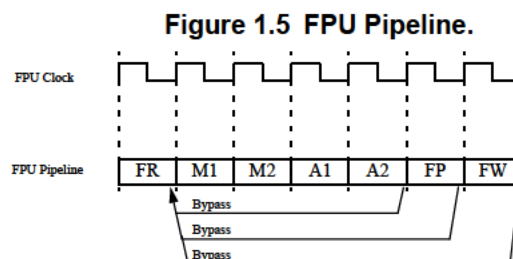
FPU Pipeline

The FPU implements a high-performance 7-stage pipeline:

- Decode, register read and unpack (FR stage)
- Multiply tree - double pumped for double (M1 stage)
- Multiply complete (M2 stage)
- Addition first step (A1 stage)
- Addition second and final step (A2 stage)
- Packing to IEEE format (FP stage)
- Register writeback (FW stage)

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.

Figure 1.5 shows the FPU pipeline.



1.2.2.5 Coprocessor 2 Interface

The M5100 core can be configured to have an interface for an on-chip coprocessor. This coprocessor can be tightly coupled to the processor core, allowing high-performance solutions integrating a graphics accelerator or DSP, for example.

The coprocessor interface is extensible and standardized on MIPS cores, allowing for design reuse. The M5100 core supports a subset of the full coprocessor interface standard: 32b data transfer, no Coprocessor 1 support, single issue in-order data transfer to coprocessor, and one out-of-order data transfer from coprocessor.

The coprocessor interface is designed to ease integration with customer IP. The interface allows high-performance communication between the core and coprocessor. There are no late or critical signals on the interface.

Refer to [Chapter 13, “M5100™ Processor Core Instructions” on page 328](#) for more information on the Coprocessor 2 supported instructions.

1.2.2.6 CorExtend® User-defined Instruction Extensions

An optional CorExtend User-defined Instruction (UDI) block enables the implementation of a small number of application-specific instructions that are tightly coupled to the core’s execution unit. The interface to the UDI block is external to the M5100 core.

Such instructions may operate on a general-purpose register, immediate data specified by the instruction word, or local state stored within the UDI block. The destination may be a general-purpose register or local UDI state. The operation may complete in one cycle or multiple cycles, if desired.

Refer to [Table 13.3 “Special2 Opcode Encoding of Function Field”](#) for a specification of the opcode map available for user-defined instructions.

1.2.2.7 Anti-Tamper Security Features

The core provides security features that counter side-channel attacks attempting to reveal the code or data running on the processor.

Random slips can be generated in order to disguise the timing and power profile of an algorithm executing on the processor. This reduces the core’s vulnerability to side channel attacks. Power management techniques are used to ensure that the power profile of a random slip is indistinguishable from the power profile of the executing instructions.

The core contains a general-purpose pseudo generator for use by software. The generator may optionally be designed by the implementor of the core, thus allowing the quality of the pseudo random number sequences to be suited for any needs.

1.2.2.8 EJTAG Debug Support

The M5100 core provides for an optional Enhanced JTAG (EJTAG) interface for use in the software debug of application and kernel code. In addition to standard user and kernel modes of operation, the M5100 core provides a Debug mode that is entered after a debug exception (derived from a hardware breakpoint, single-step exception, etc.) is taken 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 M5100 core. In addition to the standard JTAG instructions, special instructions defined in the EJTAG specification specify which registers are selected and how they are used.

Debug Registers

Four debug registers (*DEBUG*, *DEBUG2*, *DEPC*, and *DESAVE*) have been added to the MIPS Coprocessor 0 (CP0) register set. The *DEBUG* and *DEBUG2* registers show the cause of the debug exception and are used for setting up single-step operations. The *DEPC* (Debug Exception Program Counter) register holds the address on which the debug exception was taken, which is used to resume program execution after the debug operation finishes. Finally, the

DESAVE (Debug Exception Save) register enables the saving of general-purpose registers used during execution of the debug exception handler.

To exit debug mode, a Debug Exception Return (DERET) instruction is executed. When this instruction is executed, the system exits debug mode, allowing normal execution of application and system code to resume.

EJTAG Hardware Breakpoints

There are several types of *simple* hardware breakpoints defined in the EJTAG specification. These stop the normal operation of the CPU and force the system into debug mode. There are two types of simple hardware breakpoints implemented in the M5100 core: Instruction breakpoints and Data breakpoints. Additionally, *complex* hardware breakpoints can be included, which allow detection of more intricate sequences of events.

The M5100 core can be configured with the following breakpoint options:

- No data or instruction, or complex breakpoints
- One data and two instruction breakpoints, without complex breakpoints
- Two data and four instruction breakpoints, without complex breakpoints
- Two data and six instruction breakpoints, with or without complex breakpoints
- Four data and eight instruction breakpoints, with or without complex breakpoints

Instruction breakpoints occur on instruction execution operations, and the breakpoint is set on the virtual address. A mask can be applied to the virtual address to set breakpoints on a binary range of instructions.

Data breakpoints occur on load/store transactions, and the breakpoint is set on a virtual address value, with the same single address or binary address range as the Instruction breakpoint. Data breakpoints can be set on a load, a store, or both. Data breakpoints can also be set to match on the operand value of the load/store operation, with byte-granularity masking. Finally, masks can be applied to both the virtual address and the load/store value.

In addition, the M5100 core has a configurable feature to support data and instruction address-range triggered breakpoints, where a breakpoint can occur when a virtual address is either within or outside a pair of 32-bit addresses. Unlike the traditional address-mask control, address-range triggering is not restricted to a power-of-two binary boundary.

Complex breakpoints utilize the simple instruction and data breakpoints and break when combinations of events are seen. Complex break features include:

- *Pass Counters* - Each time a matching condition is seen, a counter is decremented. The break or trigger will only be enabled when the counter has counted down to 0.
- *Tuples* - A tuple is the pairing of an instruction and a data breakpoint. The tuple will match if both the virtual address of the load or store instruction matches the instruction breakpoint, and the data breakpoint of the resulting load or store address and optional data value matches.
- *Priming* - This allows a breakpoint to be enabled only after other break conditions have been met. Also called *sequential* or *armed triggering*.

- Qualified - This feature uses a data breakpoint to qualify when an instruction breakpoint can be taken. When a load matches the data address and the data value, the instruction break will be enabled. If a load matches the address, but has mis-matching data, the instruction break will be disabled.

Performance Counters

Performance counters are used to accumulate occurrences of internal predefined events/cycles/conditions for program analysis, debug, or profiling. A few examples of event types are clock cycles, instructions executed, specific instruction types executed, loads, stores, exceptions, and cycles while the CPU is stalled. There are two, 32-bit counters. Each can count an internal predefined event selected by a corresponding control register. A counter overflow can be programmed to generate an interrupt, where the interrupt-handler software can maintain larger total counts.

To allow a more flexible range of performance monitoring, two input ports, *PM_USER0* and *PM_USER1*, have been added to the core. These signals allow the core to count external events as implemented in the SOC environment.

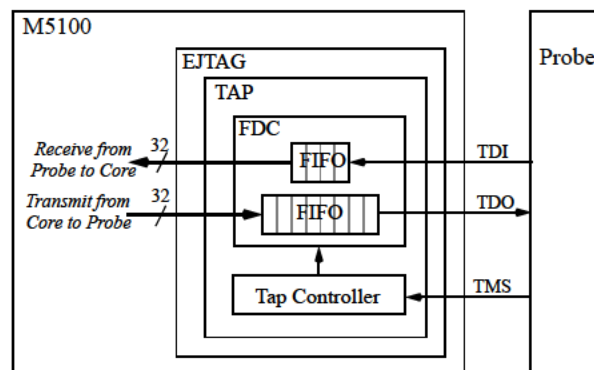
PC/Address Sampling

This sampling function is used for program profiling and hot-spots analysis. Instruction PC and/or Load/Store addresses can be sampled periodically. The result is scanned out through the EJTAG port. The *Debug Control Register (DCR)* is used to specify the sample period and the sample trigger.

Fast Debug Channel (FDC)

The M5100 core includes an optional FDC as a mechanism for high bandwidth data transfer between a debug host/probe and a target. FDC provides a FIFO buffering scheme to transfer data serially, with low CPU overhead and minimized waiting time. The data transfer occurs in the background, and the target CPU can choose either to check the status of the transfer periodically or to be interrupted at the end of the transfer.

Figure 1.6 FDC Overview



iFlowtrace™

The M5100 core has an option for a simple trace mechanism named iFlowtrace. This mechanism only traces the instruction PC, not data addresses or values. This simplification allows the trace block to be smaller and the trace compression to be more efficient. iFlowtrace memory can be configured as off-chip, on-chip, or both.

iFlowtrace also offers special-event trace modes when normal tracing is disabled, namely:

- Function Call/Return and Exception Tracing mode to trace the PC value of function calls and returns and/or exceptions and returns.

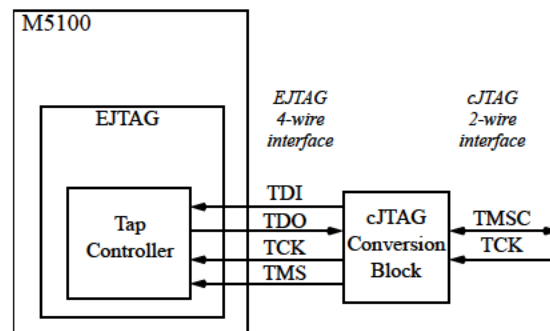
- Breakpoint Match mode traces the breakpoint ID of a matching breakpoint and, for data breakpoints, the PC value of the instruction that caused it.
- Filtered Data Tracing mode traces the ID of a matching data breakpoint, the load or store data value, access type and memory access size, and the low-order address bits of the memory access, which is useful when the data breakpoint is set up to match a binary range of addresses.
- User Trace Messages. The user can instrument their code to add their own 32-bit value messages into the trace by writing to the Cop0 UTM register.
- Delta Cycle mode works in combination with the above trace modes to provide a timestamp between stored events. It reports the number of cycles that have elapsed since the last message was generated and put into the trace.

Refer to [Chapter 11, “EJTAG Debug Support in the M5100™ Core”](#) on page 248 for more information on the EJTAG features.

cJTAG Support

The M5100 core provides an external conversion block which converts the existing EJTAG (IEEE 1149.1) 4-wire interface at the M5100 core to a cJTAG (IEEE 1149.7) 2-wire interface. cJTAG reduces the number of wires from 4 to 2 and enables the support of Star-2 scan topology in the system debug environment.

Figure 1.7 cJTAG Support



SecureDebug

SecureDebug improves security by disabling untrusted EJTAG debug access. An input signal is used to disable debug features, such as Probe Trap, Debug Interrupt Exception (EjtagBrk and DINT), EJTAGBOOT instruction, and PC Sampling.

Pipeline of the M5100™ Core

The M5100 processor core implements a 5-stage pipeline similar to the original M4K pipeline. The pipeline allows the processor to achieve high frequency while minimizing device complexity, reducing both cost and power consumption. This chapter contains the following sections:

- [Section 2.1 “Pipeline Stages”](#)
- [Section 2.2 “Multiply/Divide Operations”](#)
- [Section 2.3 “MDU Pipeline with DSP Module Enabled”](#)
- [Section 2.4 “MDU Pipeline — High-performance MDU with DSP Module Disabled”](#)
- [Section 2.5 “MDU Pipeline — Area-Efficient MDU with DSP Module Disabled”](#)
- [Section 2.6 “Branch Delay”](#)
- [Section 2.7 “Data Bypassing”](#)
- [Section 2.9 “Interlock Handling”](#)
- [Section 2.10 “Slip Conditions”](#)
- [Section 2.11 “Instruction Interlocks”](#)
- [Section 2.12 “Hazards”](#)

2.1 Pipeline Stages

The M5100 core implements a 5-stage pipeline with a performance similar to the microAptiv UC pipeline. The pipeline allows the processor to achieve high frequency while minimizing device complexity, reducing both cost and power consumption.

The M5100 core pipeline consists of five stages:

- Instruction (I Stage)
- Execution (E Stage)
- Memory (M Stage)
- Align (A Stage)
- Writeback (W stage)

The M5100 core 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 register and then read it back.

The M5100 soft core includes a build-time option that determines the type of multiply/divide unit (MDU) implemented. The MDU can be either a high-performance 32x16 multiplier array or an iterative, area-efficient array when the DSP Module configuration is not selected. The MDU choice has a significant effect on the MDU pipeline, and the latency of multiply/divide instructions executed on the core. Software can query the type of MDU present on a specific implementation of the core by querying the MDU bit in the Config register (CP0 register 16, select 0); see Chapter 8, “CP0 Registers of the M5100™ Core” on page 151 for more details. When the DSP Module configuration is selected, the multiply/divide unit (MDU) of the M5100 soft core is always implemented with a fully pipelined 32x32 multiplier array for maximum performance.

Figure 2.1 shows the operations performed in each pipeline stage of the M5100 processor core, when the high-performance multiplier is present when the DSP Module is disabled.

Figure 2.1 M5100™ Core Pipeline Stages with high-performance MDU

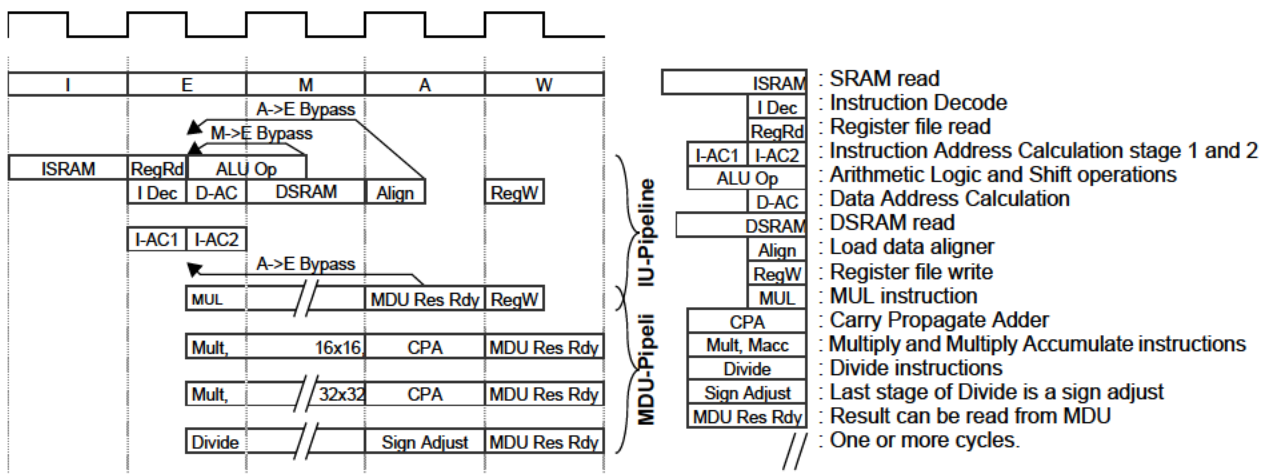


Figure 2.2 shows the operations performed in each pipeline stage of the M5100 processor core, when the area-efficient multiplier is present when the DSP Module is disabled.

Figure 2.2 M5100™ Core Pipeline Stages with area-efficient MDU

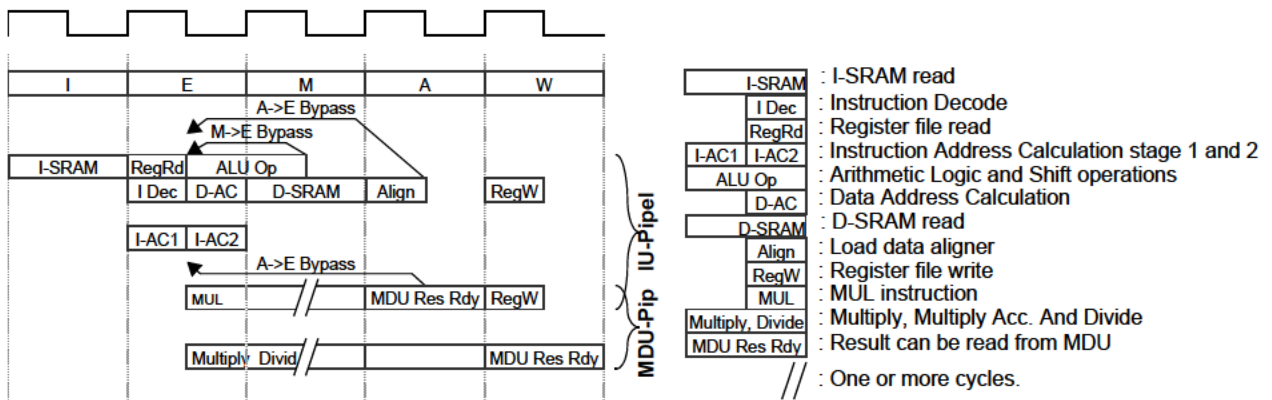
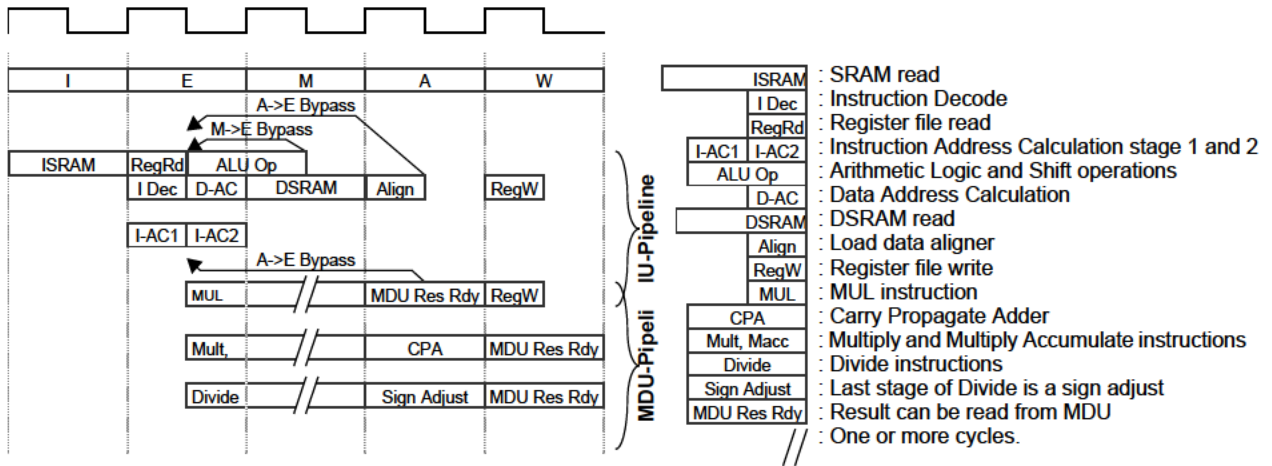


Figure 2.3 shows the operations performed in each pipeline stage of the M5100 processor core when the DSP Module is enabled.

Figure 2.3 M5100™ Core Pipeline Stages



2.1.1 I Stage: Instruction Fetch

During the Instruction fetch stage:

- An instruction is fetched from the instruction SRAM.
- If both MIPS32 and microMIPS ISAs are supported, microMIPS instructions are converted to MIPS32-like instructions. If the MIPS32 ISA is not supported, 16-bit microMIPS instructions will be first recoded into 32-bit microMIPS equivalent instructions, and then decoded in native microMIPS ISA format.

2.1.2 E Stage: Execution

During the Execution stage:

- Operands are fetched from the register file.
- Operands from the M and A stage are bypassed to this stage.
- The Arithmetic Logic Unit (ALU) begins the arithmetic or logical operation for register-to-register instructions.
- The ALU calculates the data virtual address for load and store instructions and the MMU performs the fixed virtual-to-physical address translation.
- The ALU determines whether the branch condition is true and calculates the virtual branch target address for branch instructions.
- Instruction logic selects an instruction address and the MMU performs the fixed virtual-to-physical address translation.
- All multiply and divide operations begin in this stage.

2.1.3 M Stage: Memory Fetch

During the Memory fetch stage:

- The arithmetic ALU operation completes.
- The data SRAM access is performed for load and store instructions.
- A 16x16, 32x16 or 32x32 multiply calculation completes (with DSP configuration).
- A 32x32 multiply operation stalls the MDU pipeline for one clock in the M stage (high-performance MDU option without DSP configuration).
- A multiply operation stalls the MDU pipeline for 31 clocks in the M stage (area-efficient MDU option without DSP configuration).
- A multiply-accumulate operation stalls the MDU pipeline for 33 clocks in the M stage (area-efficient MDU option without DSP configuration).
- A divide operation stalls the MDU pipeline for a maximum of 38 clocks in the M stage. Early-in sign extension detection on the dividend will skip 7, 15, or 23 stall clocks (only the divider in the fast MDU option supports early-in detection).

2.1.4 A Stage: Align

During the Align stage:

- Load data is aligned to its word boundary.
- A multiply/divide operation updates the HI/LO registers (area-efficient MDU option).
- Multiply operation performs the carry-propagate-add. The actual register writeback is performed in the W stage (high-performance MDU option).
- A MUL operation makes the result available for writeback. The actual register writeback is performed in the W stage.
- EJTAG complex break conditions are evaluated.

2.1.5 W Stage: Writeback

During the Writeback stage:

- For register-to-register or load instructions, the result is written back to the register file.

2.2 Multiply/Divide Operations

The M5100 core implements the standard MIPS II™ multiply and divide instructions. Additionally, several new instructions were standardized in the MIPS32 architecture for enhanced performance.

Pipeline of the M5100™ Core

The targeted multiply instruction, `MUL`, specifies that multiply results be placed in the general-purpose register file instead of the `HI/LO` register pair. By avoiding the explicit `MFLO` instruction, required when using the `LO` register, and by supporting multiple destination registers, the throughput of multiply-intensive operations is increased.

Four instructions, multiply-add (`MADD`), multiply-add-unsigned (`MADDU`), multiply-subtract (`MSUB`), and multiply-subtract-unsigned (`MSUBU`), are used to perform the multiply-accumulate and multiply-subtract operations. The `MADD/MADDU` instruction multiplies two numbers and then adds the product to the current contents of the `HI` and `LO` registers. Similarly, the `MSUB/MSUBU` instruction multiplies two operands and then subtracts the product from the `HI` and `LO` registers. The `MADD/MADDU` and `MSUB/MSUBU` operations are commonly used in DSP algorithms.

All multiply operations (except the `MUL` instruction) write to the `HI/LO` register pair. All integer operations write to the general purpose registers (`GPR`). Because `MDU` operations write to different registers than integer operations, integer instructions that follow can execute before the `MDU` operation has completed. The `MFLO` and `MFHI` instructions are used to move data from the `HI/LO` register pair to the `GPR` file. If an `MFLO` or `MFHI` instruction is issued before the `MDU` operation completes, it will stall to wait for the data.

2.3 MDU Pipeline with DSP Module Enabled

The M5100 processor core contains a high-performance Multiply-Divide Unit (`MDU`) and a `DSP` unit to handle integer multiply, divide, and `DSP Module` instructions.

The autonomous multiply/divide unit (`MDU`) has a separate pipeline for multiply and divide operations. This pipeline operates in parallel with the integer unit (`ALU`) pipeline and does not stall when the `ALU` pipeline stalls. This allows multi-cycle `MDU` operations, such as a divide, to be partially masked by system stalls and/or other integer unit instructions.

The following subsections describe the `MDU` pipeline in more detail.

2.3.1 MDU

The high-performance `MDU` consists of a 32x32 Booth-recoded multiplier array, separate carry-lookahead adders for multiply and divide, result/accumulation registers (`HI` and `LO`), multiply and divide state machines, and all necessary multiplexers and control logic.

Due to the multiplier array, the high-performance `MDU` supports execution of a multiply operation every clock cycle. Divide operations are implemented with a simple 1 bit-per-clock iterative algorithm with an early in detection of sign extension on the dividend (`rs`). An attempt to issue a subsequent `MDU` instruction which would access the `HI` or `LO` register before the divide completes causes a delay in starting the subsequent `MDU` instruction. Some concurrency is enabled by the separate adders for the multiply and divide data paths. The `MDU` instruction may start executing when the divide is ensured of writing to the `HI` and `LO` registers before the `MDU` instruction will access them. A `MUL` instruction, which does not access the `HI` or `LO` register, may start executing anytime relative to a previous divide instruction.

Table 2.1 lists the number of stall cycles incurred between two dependent instructions. A stall of 0 clock cycles means that the first and second instructions can be issued back-to-back in the code, without the MDU causing any stalls in the ALU pipeline.

Table 2.1 High-performance MDU Stalls

Size of Operand 1st Instruction ^[1]	Instruction Sequence		Delay Clocks
	1st Instruction	2nd Instruction	
32 bit	MULT/MULTU, MADD/MADDU, or MSUB/MSUBU	MADD/MADDU, MSUB/MSUBU, or MFHI/MFLO	0
32 bit	MUL	Integer operation ^[1]	3
8 bit	DIVU	MFHI/MFLO	8
16 bit	DIVU	MFHI/MFLO	16
24 bit	DIVU	MFHI/MFLO	24
32 bit	DIVU	MFHI/MFLO	32
8 bit	DIV	MFHI/MFLO	10 ^[2]
16 bit	DIV	MFHI/MFLO	18 ^[2]
24 bit	DIV	MFHI/MFLO	26 ^[2]
32 bit	DIV	MFHI/MFLO	34 ^[2]
any	MFHI/MFLO	Integer operation ^[1]	1
any	MTHI/MTLO	MADD/MADDU, MSUB/MSUBU	1
any	MTHI/MTLO	MFHI/MFLO	1

[1] Integer Operation refers to any integer instruction that uses the result of a previous MDU operation.
 [2] If both operands are positive, then the two Sign Adjust stages are bypassed. Delay is then the same as for DIVU.

2.3.2 DSP Module Instruction Latencies

The M5100 processor core includes support for DSP Module. Logic for these instructions is located primarily in the ALU and MDU blocks. Any DSP instructions accessing the accumulators or performing multiplication are implemented in the MDU. All others are implemented in the ALU. In addition to the “normal” MIPS32 HI/LO accumulator, the DSP Module introduces three additional HI/LO accumulator pairs.

The latency and repeat rate for the BPOSGE32 instruction is similar to those for a MIPS32 conditional branch instruction. However, unlike a MIPS32 conditional branch instruction, BPOSGE32 is dependent on *DSPControl.Pos* and not on a GPR. The LHX and LWX instructions are treated as non-blocking loads by the core; they have dependencies on the index and base registers. The delay and repeat rates for other DSP instructions are shown in the following tables. The ‘delay’ in Table 2.2 is in terms of pipeline clocks and refers to the number of cycles the pipeline must stall the second instruction in order to wait for the result of the first instruction. A delay of zero means that the first

and second instructions can be issued back-to-back without stalling the pipeline. A delay of one means that if issued back-to-back, the pipeline will stall for one cycle.

Table 2.2 DSP Instruction Delays

Dependency on ¹	Instruction Sequence		Delay Clocks
	1st Instruction	2nd Instruction	
GPR	MUL*, EXT*, MFHI, MFLO (multiplies or HI/LO reads that write to a GPR)	Instruction with GPR input	3
GPR	Other (ALU) DSP instruction with GPR result	Instruction with GPR input	0
HI/LO	DPAQ*, DPSQ*, MULSAQ*, MAQ*, MADD*, MSUB*, MULT*, MTHI, MTLO, MTTR, SHILO*, MTHLIP (HI/LO writes)	MFHI, MFLO, MFTR (HI/LO reads)	1
HI/LO	*_SA (MAC's that saturates after accumulate)	DPAQ*, DPSQ*, MULSAQ*, MAQ*, MADD*, MSUB* (MAC's)	1
HI/LO	DPAQ_S.*, DPSQ_S.*, MULSAQ_S.*, MAQ_S.*, MADD*, MSUB* (MAC's that do not saturate after accumulate)	DPAQ*, DPSQ*, MULSAQ*, MAQ*, MADD*, MSUB* (MAC's)	1
HI/LO	MTHI, MTLO, MTTR, SHILO*, MTHLIP (HI/LO writes that are not multiplies)	DPAQ*, DPSQ*, MULSAQ*, MAQ*, MADD*, MSUB* (MAC's)	1
HI/LO	DPAQ*, DPSQ*, MULSAQ*, MAQ*, MADD*, MSUB*, MULT*, MTHI, MTLO, MTTR, EXT*, SHILO*, MTHLIP (HI/LO writes)	EXT*, SHILO* (HI/LO shifts)	3
HI/LO	DPAQ*, DPSQ*, MULSAQ*, MAQ*, MADD*, MSUB*, MULT*, MTHI, MTLO, MTTR, SHILO*, MTHLIP (HI/LO writes)	MTHLIP	3

1. For dependencies on a HI/LO accumulator, the delay clocks shown assume that the 1st and 2nd instruction are operating on the same accumulator.

The delays shown in table [Table 2.2](#) with a dependency on a HI/LO accumulator pair assume that the dependent instruction sequence is operating on the *same* accumulator pair. This is the worst case situation. The delay clock value can be reduced when the second instruction operates on a different accumulator. For example, consider the following sequence:

```
MULT (writing to accumulator 0)
MADD (writing to accumulator 1)
MSUB (writing to accumulator 2)
EXTR (reading from accumulator n)
```

If the EXTR instruction is reading accumulator 2 (n=2), then a delay of 3 cycles would apply between the MSUB and EXTR operation, as indicated in Table 2.2. If the EXTR reads accumulator 1, then a delay of 2 cycles would apply between the MADD and EXTR, since there is already one unrelated instruction between the dependent ones. If the EXTR reads accumulator 0, then a delay of 1 would apply between the MULT and EXTR. Finally, if the EXTR instruction is reading accumulator 3, no delay would be incurred in the sequence.

Table 2.3 shows the repeat rates of all possible instruction sequences between two integer arithmetic, multiply, divide, or DSP instructions, with and without data dependencies.

Table 2.3 Delays for Interesting Sequences with DSPControl Dependency

MIPS32 or microMIPS	Instruction Sequence				Repeat Rate	
	1st Instruction		2nd Instruction		Without Data Dependency	With Data Dependency
	Instruction Type	Target	Instruction Type	Target		
Normal Integer Instructions	Integer Arithmetic	GPR	Integer Arithmetic	GPR	1	1
			Multiply	GPR	1	1
			Multiply	Hi/Lo	1	1
			Divide	Hi/Lo	1	1
	Multiply	GPR	Integer Arithmetic	GPR	3	4
			Multiply	GPR	1	4
			Multiply	Hi/Lo	1	4
			Divide	Hi/Lo	1	4
	Multiply	Hi/Lo	Integer Arithmetic	GPR	1	1
			Multiply	GPR	1	1
			Multiply	Hi/Lo	1	1
			Divide	Hi/Lo	1	1
	Divide	Hi/Lo	Integer Arithmetic	GPR	1	1
			Multiply	GPR	10, 18, 26, 34 ^[1]	10, 18, 26, 34 ^[1]
			Multiply	Hi/Lo	10, 18, 26, 34 ^[1]	10, 18, 26, 34 ^[1]
			Divide	Hi/Lo	10, 18, 26, 34 ^[1]	10, 18, 26, 34 ^[1]
DSP Module Instructions	Integer Arithmetic	GPR	Integer Arithmetic	GPR	1	1
			Multiply	GPR	1	1
			Multiply	Hi/Lo	1	1
	Multiply	GPR	Integer Arithmetic	GPR	3	4
			Multiply	GPR	1	4
			Multiply	Hi/Lo	1	4
	Multiply	Hi/Lo	Integer Arithmetic	GPR	1	1
			Multiply	GPR	1	1
			Multiply	Hi/Lo	1	1, 2 ^[2]

[1] : The number cycles depends on the size of input operands.

[2] : An extra cycle is needed if Saturation arithmetic is needed.

2.3.3 High-performance MDU Pipeline Stages

The multiply operation begins in stage B_{MDU} , which would be the E stage in the integer pipeline. The Booth-recoding function occurs at this time. The multiply calculation requires three clocks and occurs in the $M1_{MDU}$, $M2_{MDU}$, and $M3_{MDU}$ stages. The carry-lookahead-add (CLA) function occurs at the end of the $M3_{MDU}$ stage. In the A_{MDU} stage, the result is selected from the multiply data path, HI register, and LO register to be returned to the ALU for the MFHI, MFLO, and MUL instructions. If the MDU instruction is not one of these, the result is selected to be written into the HI/LO registers instead. The result is ready to be read from the HI/LO registers in the W_{MDU} stage.

The following figures illustrate a multiply (accumulate) instruction and the interaction with the main integer pipeline. These figures are applicable to MUL, MULT, MULTU, MADD, MADDU, MSUB, and MSUBU instructions

Figure 2.4 Multiply Pipeline



Figure 2.5 Multiply With Dependency From ALU




Figure 2.6 Multiply With Dependency From Load Hit




Figure 2.7 Multiply With Dependency From Load Miss


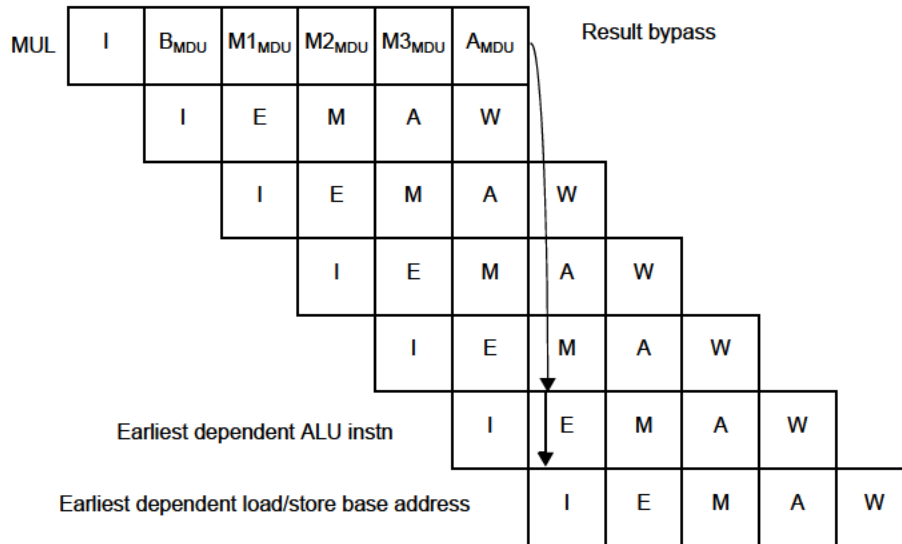


Figure 2.8 shows the results of the GPR-targeted MUL instruction being bypassed to a later instruction. Independent instructions can execute while the multiply is occurring. If a dependent instruction is found, it will stall until the result is available. When the MUL completes, it will arbitrate for access to the write port of the register file. If the integer pipe is busy with other instructions, the MDU pipeline will stall until the result can be written.

If the MUL target is being used as the base address for a load or store instruction, it needs to be bypassed by the AG stage, and thus one extra cycle will be required.

Figure 2.8 subtractMUL Bypassing Result to Integer Instructions



2.3.4 High-performance MDU Divide Operations

Divide operations are implemented using a simple non-restoring division algorithm. This algorithm works only for positive operands, and thus the first cycle of the M_{MDU} stage is used to negate the *rs* operand (RS Adjust), if needed. Note that this cycle is spent even if the adjustment is not necessary. In cycle 2, the first add/subtract iteration is executed. In cycle 3, an early-in detection is performed. The adjusted *rs* operand is detected to be zero-extended on the uppermost 8, 16, or 24 bits. If this is the case, the following 7, 15, or 23 cycles of the add/subtract iterations are skipped. During the next maximum 31 cycles (4-34), the remaining iterative add/subtract loop is executed.

The remainder adjust (Rem Adjust) cycle is required if the remainder was negative. Note that this cycle is spent even if the remainder was positive. A sign-adjust is performed on the quotient and/or remainder, if necessary. The sign adjust stages are skipped if both operands are positive.

Figure 2.16, Figure 2.17, Figure 2.11 and Figure 2.12 show the worst-case latencies for 8, 16, 24, and 32 bit divide operations respectively. The worst case repeat rate is either 14, 22, 30, or 38 cycles (two less if the *sign adjust* stage is skipped).

Figure 2.9 MDU Pipeline Flow During a 8-bit Divide (DIV) Operation




Figure 2.10 MDU Pipeline Flow During a 16-bit Divide (DIV) Operation




Figure 2.11 MDU Pipeline Flow During a 24-bit Divide (DIV) Operation





Figure 2.12 MDU Pipeline Flow During a 32-bit Divide (DIV) Operation



2.4 MDU Pipeline — High-performance MDU with DSP Module Disabled

The M5100 processor core contains an autonomous multiply/divide unit (MDU) with a separate pipeline for multiply and divide operations. This pipeline operates in parallel with the integer unit (IU) pipeline and does not stall when the IU pipeline stalls. This allows multi-cycle MDU operations, such as a divide, to be partially masked by system stalls and/or other integer unit instructions.

The MDU consists of a 32x16 Booth-encoded multiplier array, a carry propagate adder, result/accumulation registers (HI and LO), multiply and divide state machines, and all necessary multiplexers and control logic. The first number shown ('32' of 32x16) represents the *rs* operand. The second number ('16' of 32x16) represents the *rt* operand. The core only checks the latter (*rt*) operand value to determine how many times the operation must pass through the multiplier array. The 16x16 and 32x16 operations pass through the multiplier array once. A 32x32 operation passes through the multiplier array twice.

2.4 MDU Pipeline — High-performance MDU with DSP Module Disabled

The MDU supports execution of a 16x16 or 32x16 multiply operation every clock cycle; 32x32 multiply operations can be issued every other clock cycle. Appropriate interlocks are implemented to stall the issue of back-to-back 32x32 multiply operations. Multiply operand size is automatically determined by logic built into the MDU. Divide operations are implemented with a simple 1 bit per clock iterative algorithm with an early in detection of sign extension on the dividend (*rs*). Any attempt to issue a subsequent MDU instruction while a divide is still active causes an IU pipeline stall until the divide operation is completed.

Table 2.4 lists the latencies (number of cycles until a result is available) for multiply, and divide instructions. The latencies are listed in terms of pipeline clocks. In this table ‘latency’ refers to the number of cycles necessary for the first instruction to produce the result needed by the second instruction.

Table 2.4 MDU Instruction Latencies (High-Performance MDU)

Size of Operand 1st Instruction ^[1]	Instruction Sequence		Latency Clocks
	1st Instruction	2nd Instruction	
16 bit	MULT/MULTU, MADD/MADDU, MSUB/MSUBU	MADD/MADDU, MSUB/MSUBU or MFHI/MFLO	1
32 bit	MULT/MULTU, MADD/MADDU, or MSUB/MSUBU	MADD/MADDU, MSUB/MSUBU or MFHI/MFLO	2
16 bit	MUL	Integer operation ^[2]	2 ^[3]
32 bit	MUL	Integer operation ^[2]	2 ^[3]
8 bit	DIVU	MFHI/MFLO	9
16 bit	DIVU	MFHI/MFLO	17
24 bit	DIVU	MFHI/MFLO	25
32 bit	DIVU	MFHI/MFLO	33
8 bit	DIV	MFHI/MFLO	10 ^[4]
16 bit	DIV	MFHI/MFLO	18 ^[4]
24 bit	DIV	MFHI/MFLO	26 ^[4]
32 bit	DIV	MFHI/MFLO	34 ^[4]
any	MFHI/MFLO	Integer operation ^[2]	2
any	MTHI/MTLO	MADD/MADDU or MSUB/MSUBU	1

[1] For multiply operations, this is the *rt* operand. For divide operations, this is the *rs* operand.
 [2] Integer Operation refers to any integer instruction that uses the result of a previous MDU operation.
 [3] This does not include the 1 or 2 IU pipeline stalls (16 bit or 32 bit) that the MUL operation causes irrespective of the following instruction. These stalls do not add to the latency of 2.
 [4] If both operands are positive, then the Sign Adjust stage is bypassed. Latency is then the same as for DIVU.

In Table 2.4, a latency of one means that the first and second instructions can be issued back-to-back in the code, without the MDU causing any stalls in the IU pipeline. A latency of two means that if issued back-to-back, the IU pipeline will be stalled for one cycle. MUL operations are special, because the MDU needs to stall the IU pipeline in order to maintain its register file write slot. As a result, the MUL 16x16 or 32x16 operation will always force a one-cycle stall of the IU pipeline, and the MUL 32x32 will force a two-cycle stall. If the integer instruction immediately following the MUL operation uses its result, an additional stall is forced on the IU pipeline.

Table 2.5 lists the repeat rates (peak issue rate of cycles until the operation can be reissued) for multiply accumulate/subtract instructions. The repeat rates are listed in terms of pipeline clocks. In this table ‘repeat rate’ refers to the case where the first MDU instruction (in the table below) if back-to-back with the second instruction.

Table 2.5 MDU Instruction Repeat Rates (High-Performance MDU)

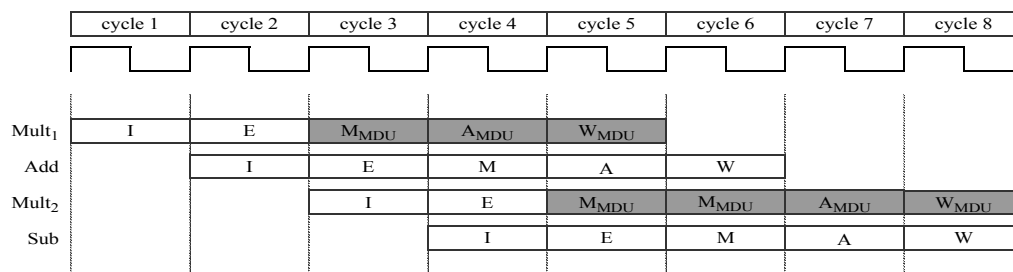
Operand Size of 1st Instruction	Instruction Sequence		Repeat Rate
	1st Instruction	2nd Instruction	
16 bit	MULT/MULTU, MADD/MADDU, MSUB/MSUBU	MADD/MADDU, MSUB/MSUBU	1
32 bit	MULT/MULTU, MADD/MADDU, MSUB/MSUBU	MADD/MADDU, MSUB/MSUBU	2

Figure 2.13 below shows the pipeline flow for the following sequence:

1. 32x16 multiply (Mult₁)
2. Add
3. 32x32 multiply (Mult₂)
4. Subtract (Sub)

The 32x16 multiply operation requires one clock of each pipeline stage to complete. The 32x32 multiply operation requires two clocks in the M_{MDU} pipe-stage. The MDU pipeline is shown as the shaded areas of Figure 2.13 and always starts a computation in the final phase of the E stage. As shown in the figure, the M_{MDU} pipe-stage of the MDU pipeline occurs in parallel with the M stage of the IU pipeline, the A_{MDU} stage occurs in parallel with the A stage, and the W_{MDU} stage occurs in parallel with the W stage. In general this need not be the case. Following the 1st cycle of the M stages, the two pipelines need not be synchronized. This does not present a problem because results in the MDU pipeline are written to the HI and LO registers, while the integer pipeline results are written to the register file.

Figure 2.13 MDU Pipeline Behavior During Multiply Operations



The following is a cycle-by-cycle analysis of Figure 2.13.

1. The first 32x16 multiply operation (Mult₁) is fetched from the instruction cache and enters the I stage.


2. An Add operation enters the I stage. The Mult₁ operation enters the E stage. The integer and MDU pipelines share the I and E pipeline stages. At the end of the E stage in cycle 2, the MDU pipeline starts processing the multiply operation (Mult₁).
3. In cycle 3, a 32x32 multiply operation (Mult₂) enters the I stage and is fetched from the instruction cache. Since the Add operation has not yet reached the M stage by cycle 3, there is no activity in the M stage of the integer pipeline at this time.
4. In cycle 4, the Subtract instruction enters I stage. The second multiply operation (Mult₂) enters the E stage. And the Add operation enters M stage of the integer pipe. Since the Mult₁ multiply is a 32x16 operation, only one clock is required for the M_{MDU} stage, hence the Mult₁ operation passes to the A_{MDU} stage of the MDU pipeline.
5. In cycle 5, the Subtract instruction enters E stage. The Mult₂ multiply enters the M_{MDU} stage. The Add operation enters the A stage of the integer pipeline. The Mult₁ operation completes and is written back in to the HI/LO register pair in the W_{MDU} stage.
6. Since a 32x32 multiply requires two passes through the multiplier, with each pass requiring one clock, the 32x32 Mult₂ remains in the M_{MDU} stage in cycle 6. The Sub instruction enters M stage in the integer pipeline. The Add operation completes and is written to the register file in the W stage of the integer pipeline.
7. The Mult₂ multiply operation progresses to the A_{MDU} stage, and the Sub instruction progress to the A stage.
8. The Mult₂ operation completes and is written to the HI/LO registers pair in the W_{MDU} stage, while the Sub instruction writes to the register file in the W stage.

2.4.1 32x16 Multiply (High-Performance MDU)

The 32x16 multiply operation begins in the last phase of the E stage, which is shared between the integer and MDU pipelines. In the latter phase of the E stage, the *rs* and *rt* operands arrive and the Booth-recoding function occurs at this time. The multiply calculation requires one clock and occurs in the M_{MDU} stage. In the A_{MDU} stage, the carry-propagate-add (CPA) function occurs and the operation is completed. The result is ready to be read from the HI/LO registers in the W_{MDU} stage.

Figure 2.14 shows a diagram of a 32x16 multiply operation.

Figure 2.14 MDU Pipeline Flow During a 32x16 Multiply Operation




2.4.2 32x32 Multiply (High-Performance MDU)

The 32x32 multiply operation begins in the last phase of the E stage, which is shared between the integer and MDU pipelines. In the latter phase of the E stage, the *rs* and *rt* operands arrive and the Booth-recoding function occurs at this time. The multiply calculation requires two clocks and occurs in the M_{MDU} stage. In the A_{MDU} stage, the CPA function occurs and the operation is completed.

Figure 2.15 shows a diagram of a 32x32 multiply operation.

Figure 2.15 MDU Pipeline Flow During a 32x32 Multiply Operation



2.4.3 Divide (High-Performance MDU)

Divide operations are implemented using a simple non-restoring division algorithm. This algorithm works only for positive operands, hence the first cycle of the M_{MDU} stage is used to negate the rs operand (RS Adjust) if needed. Note that this cycle is spent even if the adjustment is not necessary. During the next maximum 32 cycles (3-34) an iterative add/subtract loop is executed. In cycle 3 an early-in detection is performed in parallel with the add/subtract. The adjusted rs operand is detected to be zero extended on the upper most 8, 16 or 24 bits. If this is the case the following 7, 15 or 23 cycles of the add/subtract iterations are skipped.

The remainder adjust (Rem Adjust) cycle is required if the remainder was negative. Note that this cycle is spent even if the remainder was positive. A sign adjust is performed on the quotient and/or remainder if necessary. The sign adjust stage is skipped if both operands are positive. In this case the Rem Adjust is moved to the A_{MDU} stage.

Figure 2.16, Figure 2.17, Figure 2.18 and Figure 2.19 show the latency for 8, 16, 24 and 32 bit divide operations, respectively. The repeat rate is either 11, 19, 27 or 35 cycles (one less if the *sign adjust* stage is skipped) as a second divide can be in the *RS Adjust* stage when the first divide is in the *Reg WR* stage.

Figure 2.16 High-Performance MDU Pipeline Flow During a 8-bit Divide (DIV) Operation




Figure 2.17 High-Performance MDU Pipeline Flow During a 16-bit Divide (DIV) Operation




Figure 2.18 High-Performance MDU Pipeline Flow During a 24-bit Divide (DIV) Operation





Figure 2.19 High-Performance MDU Pipeline Flow During a 32-bit Divide (DIV) Operation



2.5 MDU Pipeline — Area-Efficient MDU with DSP Module Disabled

The area-efficient multiply/divide unit (MDU) is a separate autonomous block for multiply and divide operations. The MDU is not pipelined, but rather performs the computations iteratively in parallel with the integer unit (IU) pipeline and does not stall when the IU pipeline stalls. This allows the long-running MDU operations to be partially masked by system stalls and/or other integer unit instructions.

The MDU consists of one 32-bit adder result-accumulate registers (HI and LO), a combined multiply/divide state machine, and all multiplexers and control logic. A simple 1-bit-per-clock recursive algorithm is used for both multiply and divide operations. Using Booth’s algorithm all multiply operations complete in 32 clocks. Two extra clocks are needed for multiply-accumulate. The non-restoring algorithm used for divide operations will not work with negative numbers. Adjustment before and after are thus required depending on the sign of the operands. All divide operations complete in 33 to 35 clocks.

Table 2.6 lists the latencies (number of cycles until a result is available) for multiply and divide instructions. The latencies are listed in terms of pipeline clocks. In this table ‘latency’ refers to the number of cycles necessary for the second instruction to use the results of the first.

Table 2.6 M5100™ Core Instruction Latencies (Area-Efficient MDU)

Operand Signs of 1st Instruction (Rs,Rt)	Instruction Sequence		Latency Clocks
	1st Instruction	2nd Instruction	
any, any	MULT/MULTU	MADD/MADDU, MSUB/MSUBU, or MFHI/MFLO	32
any, any	MADD/MADDU, MSUB/MSUBU	MADD/MADDU, MSUB/MSUBU, or MFHI/MFLO	34
any, any	MUL	Integer operation ^[1]	32
any, any	DIVU	MFHI/MFLO	33

Table 2.6 M5100™ Core Instruction Latencies (Area-Efficient MDU)

Operand Signs of 1st Instruction (Rs,Rt)	Instruction Sequence		Latency Clocks
	1st Instruction	2nd Instruction	
pos, pos	DIV	MFHI/MFLO	33
any, neg	DIV	MFHI/MFLO	34
neg, pos	DIV	MFHI/MFLO	35
any, any	MFHI/MFLO	Integer operation ^[1]	2
any, any	MTHI/MTLO	MADD/MADDU, MSUB/MSUBU	1


[1] Integer Operation refers to any integer instruction that uses the result of a previous MDU operation.

2.5.1 Multiply (Area-Efficient MDU)

Multiply operations are executed using a simple iterative multiply algorithm. Using Booth’s approach, this algorithm works for both positive and negative operands. The operation uses 32 cycles in M_{MDU} stage to complete a multiplication. The register writeback to HI and LO are done in the A stage. For MUL operations, the register file writeback is done in the W_{MDU} stage.

Figure 2.20 shows the latency for a multiply operation. The repeat rate is 33 cycles as a second multiply can be in the first M_{MDU} stage when the first multiply is in A_{MDU} stage.

Figure 2.20 M5100™ Area-Efficient MDU Pipeline Flow During a Multiply Operation




2.5.2 Multiply Accumulate (Area-Efficient MDU)

Multiply-accumulate operations use the same multiply machine as used for multiply only. Two extra stages are needed to perform the addition/subtraction. The operations uses 34 cycles in M_{MDU} stage to complete the multiply-accumulate. The register writeback to HI and LO are done in the A stage.

Figure 2.21 shows the latency for a multiply-accumulate operation. The repeat rate is 35 cycles as a second multiply-accumulate can be in the E stage when the first multiply is in the last M_{MDU} stage.

Figure 2.21 M5100™ Core Area-Efficient MDU Pipeline Flow During a Multiply Accumulate Operation




2.5.3 Divide (Area-Efficient MDU)

Divide operations also implement a simple non-restoring algorithm. This algorithm works only for positive operands, hence the first cycle of the M_{MDU} stage is used to negate the rs operand (RS Adjust) if needed. Note that this cycle is executed even if negation is not needed. The next 32 cycle (3-34) executes an interactive add/subtract-shift function.

Two sign adjust (Sign Adjust 1/2) cycles are used to change the sign of one or both the quotient and the remainder. Note that one or both of these cycles are skipped if they are not needed. The rule is, if both operands were positive or if this is an unsigned division; both of the sign adjust cycles are skipped. If the rs operand was negative, one of the sign adjust cycles is skipped. If only the rs operand was negative, none of the sign adjust cycles are skipped. Register writeback to HI and LO are done in the A stage.

Figure 2.22 shows the pipeline flow for a divide operation. The repeat rate is either 34, 35 or 36 cycles (depending on how many sign adjust cycles are skipped) as a second divide can be in the E stage when the first divide is in the last M_{MDU} stage.

Figure 2.22 M5100™ Core Area-Efficient MDU Pipeline Flow During a Divide (DIV) Operation



2.6 Branch Delay

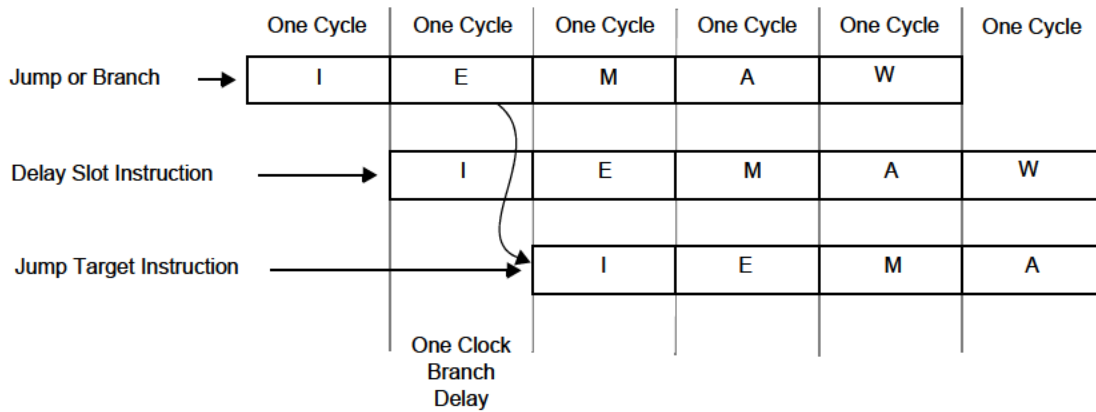
The pipeline has a branch delay of one cycle. The one-cycle branch delay is a result of the branch decision logic operating during the E pipeline stage. This allows the branch target address to be used in the I stage of the instruction following 2 cycles after the branch instruction. By executing the 1st instruction following the branch instruction sequentially before switching to the branch target, the intervening branch delay slot is utilized. This avoids bubbles being injected into the pipeline on branch instructions. Both the address calculation and the branch condition check are performed in the E stage.

The pipeline begins the fetch of either the branch path or the fall-through path in the cycle following the delay slot. After the branch decision is made, the processor continues with the fetch of either the branch path (for a taken branch) or the fall-through path (for the non-taken branch).

The branch delay means that the instruction immediately following a branch is always executed, regardless of the branch direction. If no useful instruction can be placed after the branch, then the compiler or assembler must insert a NOP instruction in the delay slot.

Figure 2.23 illustrates the branch delay.

Figure 2.23 IU Pipeline Branch Delay



2.7 Data Bypassing

Most MIPS32 instructions use one or two register values as source operands. These operands are fetched from the register file in the first part of E stage. The ALU straddles the E-to-M boundary, and can present the result early in the M stage. However, the result is not written to the register file before the W stage. If no precautions were taken, it would take 3 cycles before the result was available for the following instructions. To avoid this, data bypassing is implemented.

Between the register file and the ALU a data-bypass multiplexer is placed on both operands (see figure below). This enables the M5100 core to forward data from a preceding instruction whose target is a source register of a following instruction. An M to E bypass and an A to E bypass feed the bypass multiplexers. A W to E bypass is not needed, as the register file is capable of making an internal bypass of Rd write data directly to the Rs and Rt read ports.

Figure 2.24 IU Pipeline Data bypass

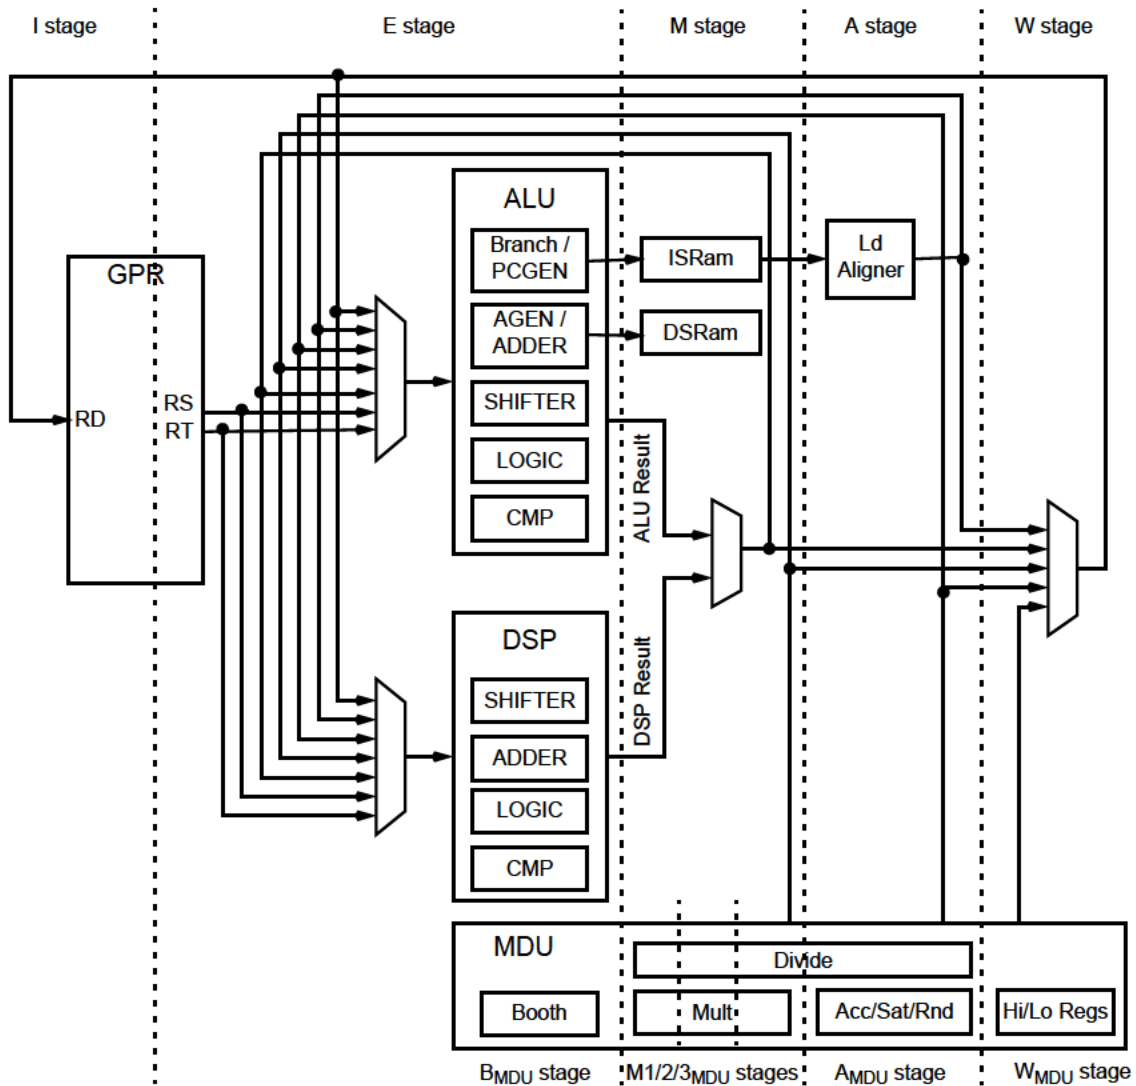
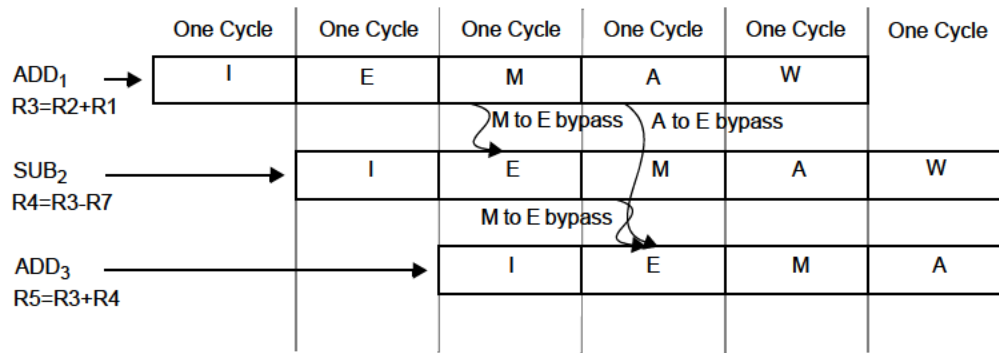


Figure 2.25 shows the data bypass for an Add_1 instruction followed by a Sub_2 and another Add_3 instruction. The Sub_2 instruction uses the output from the Add_1 instruction as one of the operands, and thus the M to E bypass is used. The following Add_3 uses the result from both the first Add_1 instruction and the Sub_2 instruction. Since the Add_1 data is now in A stage, the A to E bypass is used, and the M to E bypass is used to bypass the Sub_2 data to the Add_2 instruction.

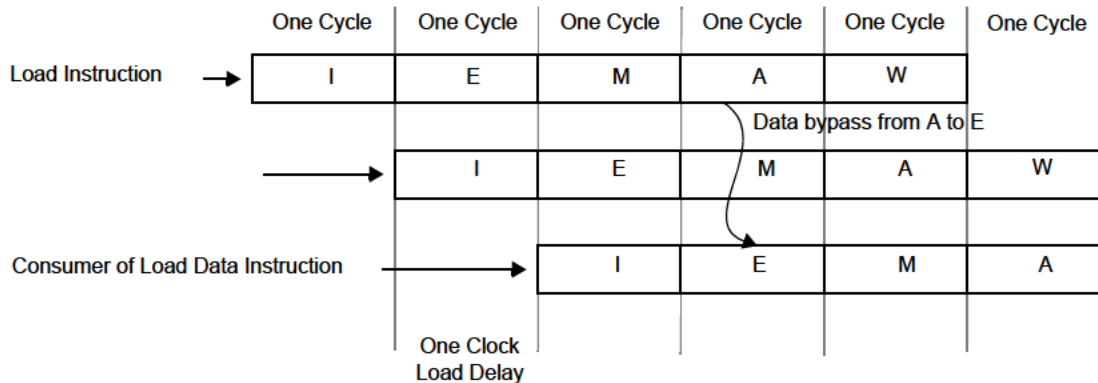
Figure 2.25 IU Pipeline M to E bypass



2.7.1 Load Delay

Load delay refers to the fact that data fetched by a load instruction is not available in the integer pipeline until after the load aligner in A stage. All instructions need the source operands available in the E stage. An instruction immediately following a load instruction will, if it has the same source register as was the target of the load, cause an instruction interlock pipeline slip in the E stage (see 2.11 “Instruction Interlocks” on page 63). If an instruction following the load by 1 or 2 cycles uses the data from the load, the A to E bypass (see Figure 2.24) serves to reduce or avoid stall cycles. An instruction flow of this is shown in Figure 2.26.


Figure 2.26 IU Pipeline A to E Data bypass



2.7.2 Move from HI/LO and CP0 Delay

As indicated in Figure 2.24, not only load data, but also data moved from the HI or LO registers (MFHI/MFLO) and data moved from CP0 (MFC0) enters the IU-Pipeline in the A stage. That is, data is not available in the integer pipeline until early in the A stage. The A to E bypass is available for this data. But as for Loads, an instruction following immediately after one of these move instructions must be paused for one cycle if the target of the move is among the sources of the following instruction and this causes an interlock slip in the E stage (see 2.11 “Instruction Interlocks” on page 63). An interlock slip after a MFHI is illustrated in Figure 2.27.

Figure 2.27 IU Pipeline Slip after a MFHI



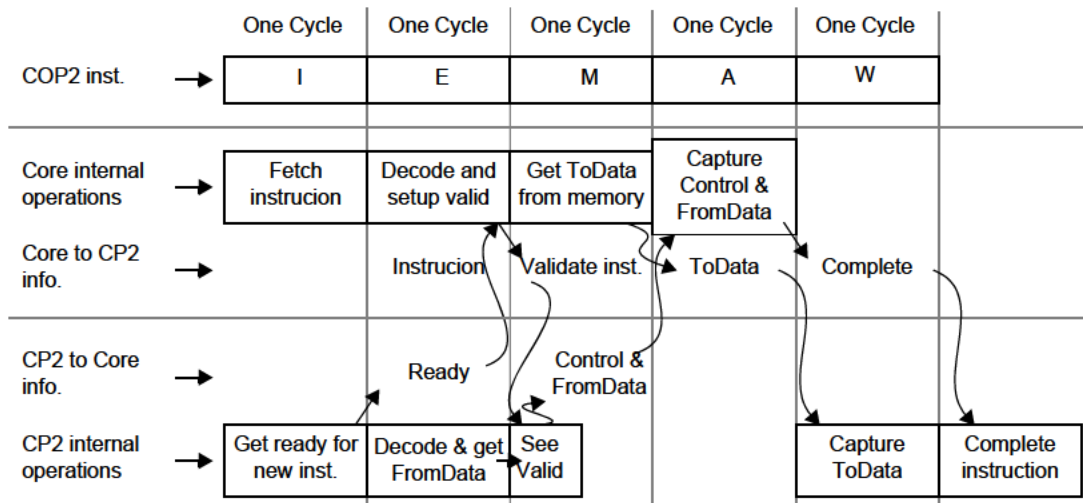
2.8 Coprocessor 2 Instructions

If a coprocessor 2 is attached to the M5100 core, a number of transactions must take place on the CP2 Interface for each coprocessor 2 instruction. First, if the CU[2] bit in the CP0 *Status* register is not set, then no coprocessor 2 related instruction will start a transaction on the CP2 Interface; instead, a Coprocessor Unusable exception will be signaled. If the CU[2] bit is set, and a coprocessor 2 instruction is fetched, the following transactions will occur on the CP2 Interface:

1. The Instruction is presented on the instructions bus in E stage. Coprocessor 2 can do a decode in the same cycle.
2. The Instruction is validated from the core in M stage. From this point, the core will accept control and data signals back from coprocessor 2. All control and data signals from coprocessor 2 are captured on input latches to the core.
3. If all the expected control and data signals were presented to the core in the previous M stage, the core will proceed to execute the A stage. If some return information is missing, the A stage will not advance and cause a slip in all I, E, and M stages (see 2.10 “Slip Conditions” on page 63).
If this instruction sent data from the core to coprocessor 2, this data is sent in the A stage.
4. The instruction completion is signaled to coprocessor 2 in the W stage. Potential data from the coprocessor is written to the register file.

Figure 2.28 shows the timing relationship between the M5100 core and coprocessor 2 for all coprocessor 2 instructions.

Figure 2.28 Coprocessor 2 Interface Transactions



As can be seen in the Figure, all control and data from the coprocessor must occur in the M stage. If this is not the case, the A stage will start slipping in the following cycle and thus stall the I, E, M, and A stages; but if all expected control and data is available in the M stage, coprocessor 2 instructions can execute with no pipeline stalls. The only exception to this is the Branch on Coprocessor conditions (BC2) instruction. All branch instructions, including the regular BEQ, BNE, etc., must be resolved in the E stage. The M5100 core does not have branch prediction logic, and thus the target address must be available before the end of the E stage. The BC2 instruction has to follow the same protocol as all other coprocessor 2 instructions on the CP2 Interface. All core interface operations belonging to the E, M, and A stages will have to occur in the E stage for BC2 instructions. This means that a BC2 instruction always slips for a minimum of 2 cycles into the E stage, and any delay in the return of branch information from coprocessor 2 will add to the number of slip cycles. All other Coprocessor 2 instructions can operate without slips, provided that all control and data information from coprocessor 2 is transferred in the M stage.

2.9 Interlock Handling

Smooth pipeline flow is interrupted when cache misses occur or when data dependencies are detected. Interruptions handled entirely in hardware, such as cache misses, are referred to as *interlocks*. At each cycle, interlock conditions are checked for all active instructions.

Table 2.7 lists the types of pipeline interlocks for the M5100 processor core.

Table 2.7 Pipeline Interlocks

Interlock Type	Sources	Slip Stage
I-side SRAM Stall	SRAM Access not complete	E Stage
Instruction	Producer-consumer hazards	E/M Stage
	Hardware Dependencies (MDU)	E Stage
	BC2 waiting for COP2 Condition Check	E Stage
D-side SRAM Stall	SRAM Access not complete	A Stage
Coprocessor 2 completion slip	Coprocessor 2 control and/or data delay from coprocessor	A Stage

In general, MIPS processors support two types of hardware interlocks:

- Stalls, which are resolved by halting the pipeline
- Slips, which allow one part of the pipeline to advance while another part of the pipeline is held static

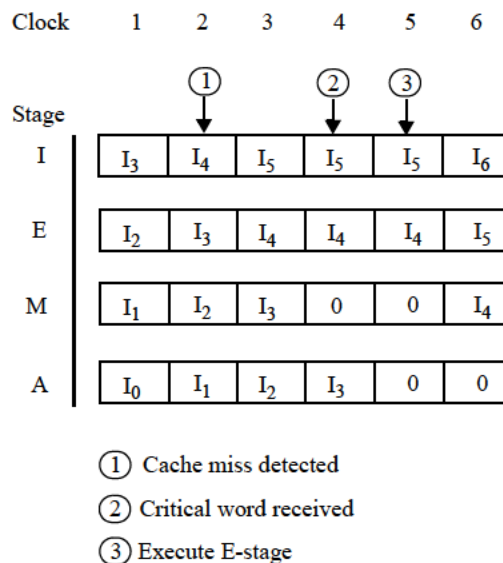
In the M5100 processor core, all interlocks are handled as slips.

2.10 Slip Conditions

On every clock, internal logic determines whether each pipe stage is allowed to advance. These slip conditions propagate backwards down the pipe. For example, if the M stage does not advance, neither does the E or I stage.

Slipped instructions are retried on subsequent cycles until they issue. The back end of the pipeline advances normally during slips. This resolves the conflict when the slip was caused by a missing result. NOPs are inserted into the bubble in the pipeline. [Figure 2.29](#) shows an instruction cache miss that causes a two-cycle slip.

Figure 2.29 Instruction Cache Miss Slip



In the first clock cycle in [Figure 2.29](#), the pipeline is full and the cache miss is detected. Instruction I₀ is in the A stage, instruction I₁ is in the M stage, instruction I₂ is in the E stage, and instruction I₃ is in the I stage. The cache miss occurs in clock 2 when the I₄ instruction fetch is attempted. I₄ advances to the E stage and waits for the instruction to be fetched from main memory. In this example, two clocks (3 and 4) are required to fetch the I₄ instruction from memory. After the cache miss has been resolved in clock 4 and the instruction is bypassed to the E stage, the pipeline is restarted, causing I₄ to finally execute its E-stage operations.

2.11 Instruction Interlocks

Most instructions can be issued at a rate of one per clock cycle. In order to adhere to the sequential programming model, the issue of an instruction must sometimes be delayed to ensure that the result of a prior instruction is avail-

able. Table 2.8 details the instruction interactions that prevent an instruction from advancing in the processor pipeline.

Table 2.8 Instruction Interlocks

Instruction Interlocks				
First Instruction		Second Instruction	Issue Delay (in Clock Cycles)	Slip Stage
LB/LBU/LH/LHU/LL/LW/LWL/LWR		Consumer of load data	1	E stage
MFC0		Consumer of destination register	1	E stage
MULTx/MADDx/MSUBx (high-performance MDU with DSP Module disabled)	16bx32b	MFLO/MFHI	0	
	32bx32b		1	M stage
MUL (high-performance MDU with DSP Module disabled)	16bx32b	Consumer of target data	2	E stage
	32bx32b		3	E stage
MUL (high-performance MDU with DSP Module disabled)	16bx32b	Non-Consumer of target data	1	E stage
	32bx32b		2	E stage
MFHI/MFLO		Consumer of target data	1	E stage
MULTx/MADDx/MSUBx (high-performance MDU with DSP Module disabled)	16bx32b	MULT/MUL/MADD/MSUB MTHI/MTLO/DIV	0 ^[1]	E stage
	32bx32b		1 ^[1]	E stage
DIV		MUL/MULTx/MADDx/ MSUBx/MTHI/MTLO/ MFHI/MFLO/DIV	Until DIV completes	E stage
MULT/MUL/MADD/MSUB/MTHI/MTLO/MFHI /MFLO/DIV (area-efficient MDU with DSP Module disabled)		MULT/MUL/MADD/MSUB/ MTHI/MTLO/MFHI/MFLO/ DIV	Until 1st MDU op completes	E stage
MUL (area-efficient MDU with DSP Module disabled)		Any Instruction	Until MUL completes	E stage
MFC0/MFC2/CFC2		Consumer of target data	1	E stage

2.12 Hazards

In general, the M5100 core ensures that instructions are executed following a fully sequential program model in which each instruction in the program sees the results of the previous instruction. There are some deviations to this model, referred to as *hazards*.

Prior to Release 2 of the MIPS Architecture, hazards (primarily CP0 hazards) were relegated to implementation-dependent cycle-based solutions, primarily based on the SSNOP instruction. This has been an insufficient and error-prone practice that must be addressed with a firm compact between hardware and software. As such, new instructions have been added to Release 2 of the architecture which act as explicit barriers that eliminate hazards. To the extent that it was possible to do so, the new instructions have been added in such a way that they are backward-compatible with existing MIPS processors.

2.12.1 Types of Hazards

With one exception, all hazards were eliminated in Release 1 of the Architecture for unprivileged software. The exception occurs when unprivileged software writes a new instruction sequence and then wishes to jump to it. Such an operation remained a hazard, and is addressed by the capabilities of Release 2.

In privileged software, there are two types of hazards: *execution hazards* and *instruction hazards*.

Execution hazards are those created by the execution of one instruction, and seen by the execution of another instruction. [Table 2.9](#) lists execution hazards.

Table 2.9 Execution Hazards

Producer	→	Consumer	Hazard On	Spacing (Instructions)
MTC0	→	Coprocessor instruction execution depends on the new value of Status _{CU}	Status _{CU}	1
MTC0	→	ERET	EPC DEPC ErrorEPC	1
MTC0	→	ERET	Status	0
MTC0, EI, DI	→	Interrupted Instruction	Status _{IE}	1
MTC0	→	Interrupted Instruction	Cause _{IP}	3
MTC0	→	RDPGPR WRPGPR	SRSCtl _{PSS}	1
MTC0	→	Instruction not seeing a Timer Interrupt	Compare update that clears Timer Interrupt	4 ¹
MTC0	→	Instruction affected by change	Any other CP0 register	2

1. This is the minimum value. Actual value is system-dependent since it is a function of the sequential logic between the *SL_TimerInt* output and the external logic which feeds *SL_TimerInt* back into one of the *SL_Int* inputs, or a function of the method for handling *SL_TimerInt* in an external interrupt controller.

Instruction hazards are those created by the execution of one instruction, and seen by the instruction fetch of another instruction. [Table 2.10](#) lists instruction hazards.

Table 2.10 Instruction Hazards

Producer	→	Consumer	Hazard On	Spacing (Instructions)
MTC0	→	Instruction fetch seeing the new value (including a change to ERL followed by an instruction fetch from the useg segment)	Status	
Instruction stream write via redirected store	→	Instruction fetch seeing the new instruction stream	Cache entries	3

2.12.2 Instruction Listing

Table 2.11 lists the instructions designed to eliminate hazards. See the document titled *MIPS32® Architecture for Programmers Volume II: The MIPS32® Instruction Set* (MD00086) for a more detailed description of these instructions.

Table 2.11 Hazard Instruction Listing

Mnemonic	Function
EHB	Clear execution hazard
JALR.HB	Clear both execution and instruction hazards
JR.HB	Clear both execution and instruction hazards
SYNCI	Synchronize caches after instruction stream write

2.12.2.1 Instruction Encoding

The EHB instruction is encoded using a variant of the NOP/SSNOP encoding. This encoding was chosen for compatibility with the Release 1 SSNOP instruction, such that existing software may be modified to be compatible with both Release 1 and Release 2 implementations. See the EHB instruction description for additional information.

The JALR.HB and JR.HB instructions are encoding using bit 10 of the *hint* field of the JALR and JR instructions. These encodings were chosen for compatibility with existing MIPS implementations, including many which pre-date the MIPS architecture. Because a pipeline flush clears hazards on most early implementations, the JALR.HB or JR.HB instructions can be included in existing software for backward and forward compatibility. See the JALR.HB and JR.HB instructions for additional information.

The SYNCI instruction is encoded using a new encoding of the REGIMM opcode. This encoding was chosen because it causes a Reserved Instruction exception on all Release 1 implementations. As such, kernel software running on processors that don't implement Release 2 can emulate the function using the CACHE instruction.

2.12.3 Eliminating Hazards

The Spacing column shown in Table 2.9 and Table 2.10 indicates the number of unrelated instructions (such as NOPs or SSNOPs) that, prior to the capabilities of Release 2, would need to be placed between the producer and consumer of the hazard in order to ensure that the effects of the first instruction are seen by the second instruction. Entries in the table that are listed as 0 are traditional MIPS hazards which are not hazards on the M5100 core.

With the hazard elimination instructions available in Release 2, the preferred method to eliminate hazards is to place one of the instructions listed in Table 2.11 between the producer and consumer of the hazard. Execution hazards can be removed by using the EHB, JALR.HB, or JR.HB instructions. Instruction hazards can be removed by using the JALR.HB or JR.HB instructions, in conjunction with the SYNCI instruction. Since the M5100 core does not contain caches, the SYNCI instruction is not strictly necessary, but is still recommended to create portable code that can be run on other MIPS processors that may contain caches.

Floating-Point Unit of the M5100™ Core

This chapter describes the MIPS64® Floating-Point Unit (FPU) included in the M5100 core. This chapter contains the following sections:

- [Section 3.1 “Features Overview”](#)
- [Section 3.2 “Enabling the Floating-Point Coprocessor”](#)
- [Section 3.3 “Data Formats”](#)
- [Section 3.4 “Floating-Point General Registers”](#)
- [Section 3.5 “Floating-Point Control Registers”](#)
- [Section 3.6 “Instruction Overview”](#)
- [Section 3.7 “Exceptions”](#)
- [Section 3.8 “Pipeline and Performance”](#)
- [Section 3.9 “2008 FPU Support”](#)

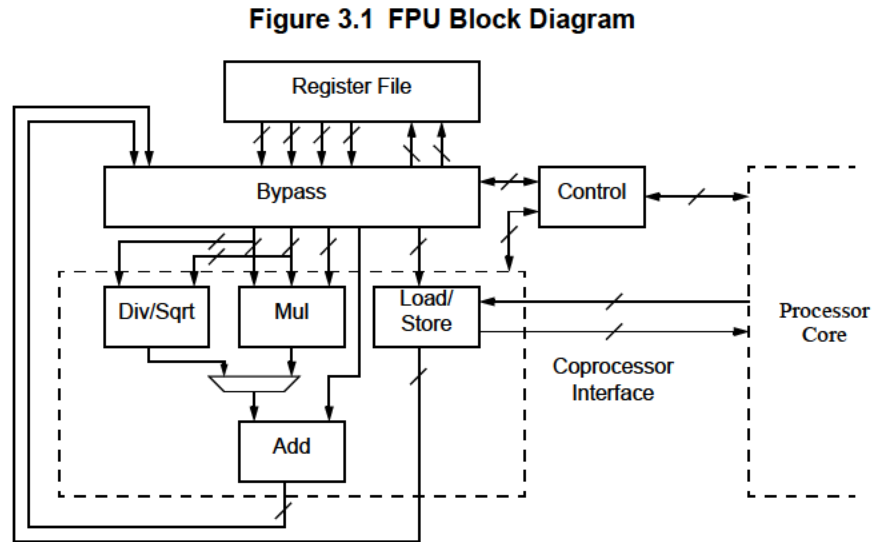
3.1 Features Overview

The FPU is provided via Coprocessor 1. 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 negate instructions.

Figure 3.1 depicts a block diagram of the FPU.



The MIPS architecture is designed such that a combination of hardware and software can be used to implement the architecture. The M5100 core FPU can operate on numbers within a specific range (in general, the IEEE normalized numbers), but it relies on a software handler to operate on numbers not handled by the FPU hardware (in general, the IEEE denormalized numbers). Supported number ranges for different instructions are described later in this chapter. A fast Flush To Zero mode is provided to optimize performance for cases where IEEE denormalized operands and results are not supported by hardware. The fast Flush to Zero mode is enabled through the CP1 *FCSR* register; use of this mode is recommended for best performance.

3.1.1 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.

For more information about this standard, see the IEEE web page at <http://stdsbbs.ieee.org/>.

3.2 Enabling the Floating-Point Coprocessor

3.3 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 MIPS architecture.

3.3.1 Floating-Point Formats

The FPU provides the following two floating-point formats:

- a 32-bit single-precision floating point (type S, shown in [Figure 3.2](#))
- a 64-bit double-precision floating point (type D, shown in [Figure 3.3](#))

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 3.1](#).

Table 3.1 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

Table 3.1 Parameters of Floating-Point Data Types (Continued)

Parameter	Single	Double
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 [Figure 3.2](#) and [Figure 3.3](#) 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 3.2 Single-Precision Floating-Point Format (S)





Figure 3.3 Double-Precision Floating-Point Format (D)



Values are encoded in the specified format using the unbiased exponent, fraction, and sign values listed in [Table 3.2](#). The high-order bit of the Fraction field, identified as b_1 , is also important for NaNs.

Table 3.2 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 ¹	Typical Double Bit Pattern ¹
$E_{max} + 1$	$\neq 0$		1	SNaN	Signaling NaN ($FCSR_{NAN2008} = 0$)	0x7fffffff	0x7fffffff ffffffff
			0	QNaN	Quiet NaN ($FCSR_{NAN2008} = 0$)	0x7fbfffff	0x7ff7ffff ffffffff
$E_{max} + 1$	$\neq 0$		1	QNaN	Quiet NaN ($FCSR_{NAN2008} = 1$)	0x7fffffff	0x7fffffff ffffffff
			0	SNaN	Signaling NaN ($FCSR_{NAN2008} = 1$)	0x7fbfffff	0x7ff7ffff ffffffff
$E_{max} + 1$	0	1		$-\infty$	Minus infinity	0xff800000	0xfff00000 00000000
		0		$+\infty$	Plus infinity	0x7f800000	0x7ff00000 00000000

Table 3.2 Value of Single or Double Floating-Point Data Type Encoding (Continued)

Unbiased E	f	s	b ₁	Value V	Type of Value	Typical Single Bit Pattern ¹	Typical Double Bit Pattern ¹
E_{min} to E_{max}		1		$-(2^E)(1.f)$	Negative normalized number	0x80800000 through 0xff7fffff	0x80100000 00000000 through 0xffefffff 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.

3.3.1.1 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.

3.3.1.2 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 Table 3.2) for plus infinity ($+\infty$), minus infinity ($-\infty$), quiet non-numbers (QNaN), and signaling non-numbers (SNaN).

3.3.1.3 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 ∞ is generated as the default result in division by zero operations and some cases of overflow as described in Section 3.7.2 “Exception Conditions”.

When created as a default result, ∞ can become an operand in a subsequent operation. The infinities are interpreted such that $-\infty < (\text{every finite number}) < +\infty$. Arithmetic with ∞ is the limiting case of real arithmetic with operands of arbitrarily large magnitude, when such limits exist. In these cases, arithmetic on ∞ is regarded as exact, and exception conditions do not arise. The out-of-range indication represented by ∞ is propagated through subsequent computations. For some cases, there is no meaningful limiting case in real arithmetic for operands of ∞ . These cases raise the Invalid Operation exception condition as described in Section 3.7.2.1 “Invalid Operation Exception”.

3.3.1.4 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) non-arithmetic; they do not signal IEEE 754 exceptions.

3.3.1.5 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 one¹ 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 3.3 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 3.3 Value Supplied When a New Quiet NaN is Created

Format	New QNaN value (FCSR _{NAN2008} = 0)	New QNaN value (FCSR _{NAN2008} = 1)
Single floating-point	0x7fbf ffff	0x7fff ffff
Double floating-point	0x7ff7 ffff ffff ffff	0x7fff ffff ffff ffff
Word fixed-point	0x7fff ffff	0x7fff ffff
Longword fixed-point	0x7fff ffff ffff ffff	0x7fff ffff ffff ffff



3.3.2 Fixed-Point Formats

The FPU provides two fixed-point data types:

- a 32-bit Word fixed point (type W), shown in Figure 3.4
- a 64-bit Longword fixed point (type L), shown in Figure 3.5

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.

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.

Figure 3.4 Word Fixed-Point Format (W)**Figure 3.5 Longword Fixed-Point Format (L)**

3.4 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 FPU is in FR64 mode and the 64b register model is used, which defines 32 64-bit registers with all formats supported in a register.
- When the FR bit is a 0, the FPU is in FR32 mode and the 32b register model is used, which defines 32 32-bit registers with D-format values stored in even-odd pairs of registers; thus the register file can also be viewed as having 16 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.

3.4.1 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.

[Figure 3.6](#) and [Figure 3.7](#) show the FPR organization and the way that operand data is stored in them.

Figure 3.6 Single Floating-Point or Word Fixed-Point Operand in an FPR**Figure 3.7 Double Floating-Point or Longword Fixed-Point Operand in an FPR**

3.4.2 Formats of Values Used in FP 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

Floating-Point Unit of the M5100™ Core

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.

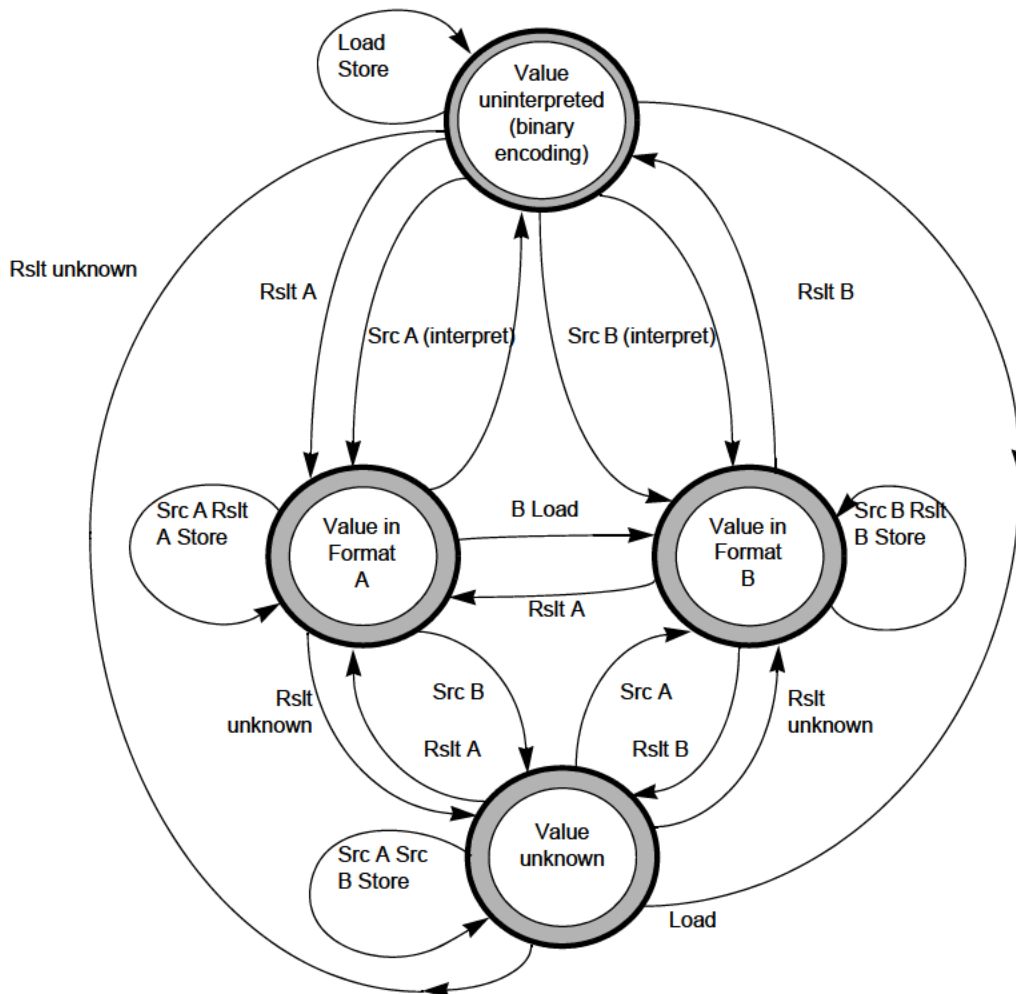
When an FPR with an *uninterpreted* value is used as a source operand by an instruction that requires a value of format *fmt*, the binary contents are interpreted as an encoded value in format *fmt*, and the value in the FPR changes to a value of format *fmt*. The binary contents cannot be reinterpreted 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 3.8](#) illustrates the manner in which the formatted value in an FPR is set and changed.

Figure 3.8 Effect of FPU Operations on the Format of Values Held in FPRs



A, B: Example formats

Load: Destination of LWC1, LDC1, or MTC1 instructions.

Store: Source operand of SWC1, SDC1, or MFC1 instructions.

Src fmt: Source operand of computational instruction expecting format "fmt."

Rslt fmt: Result of computational instruction producing value of format "fmt."

3.4.3 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 3.9](#) and [Figure 3.10](#), 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 3.9 FPU Word Load and Move-to Operations

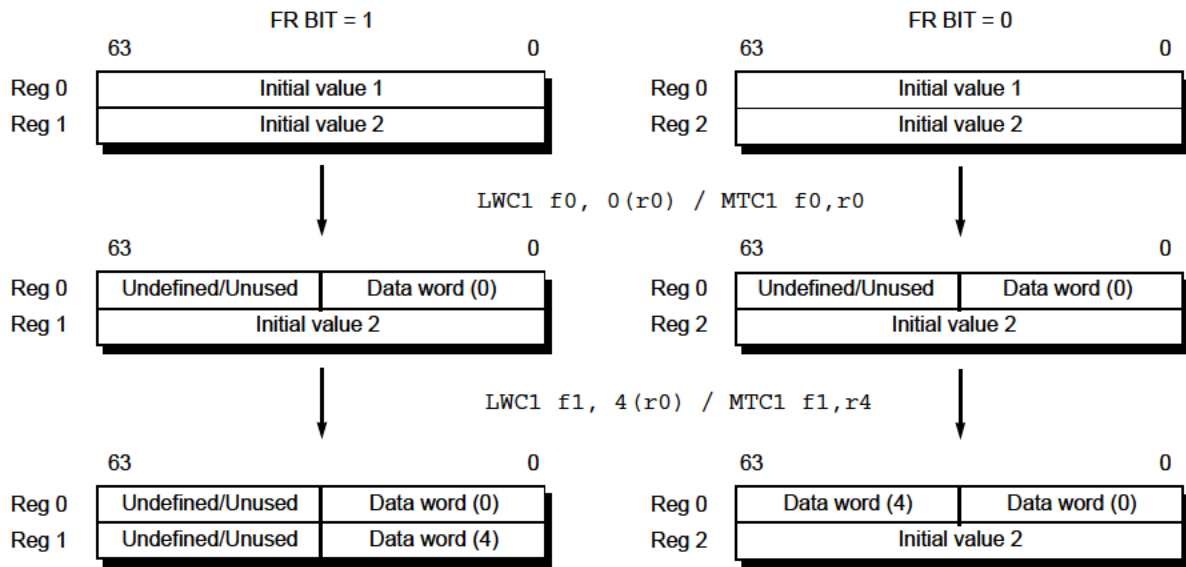



Figure 3.10 FPU Doubleword Load and Move-to Operations



3.5 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.

CP1 control registers are summarized in [Table 3.4](#) 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 3.4 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 3.5](#) defines the notation used for the read/write properties of the register bit fields.

Table 3.5 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 “Pre-set”, 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.

3.5.1 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 3.11](#) shows the format of the *FIR*; [Table 3.6](#) describes the *FIR* bit fields.

Figure 3.11 FIR Format




Table 3.6 FIR Bit Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
UFRP	28	See Release 5 definition of CFC1 and CTC1 <table border="1" style="width: 100%; border-collapse: collapse; margin-left: 20px;"> <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>User mode FR switching instructions not supported.</td> </tr> <tr> <td style="text-align: center;">1</td> <td>User mode FR switching instructions supported.</td> </tr> </tbody> </table>	Encoding	Meaning	0	User mode FR switching instructions not supported.	1	User mode FR switching instructions supported.	R	1
Encoding	Meaning									
0	User mode FR switching instructions not supported.									
1	User mode FR switching instructions supported.									
FC	24	Indicates that full convert ranges are implemented: 0: Full convert ranges not implemented 1: Full convert ranges implemented 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 occur on denormal operands though).	R	1						
Has2008	23	Indicates that one or more IEEE-754-2008 features are implemented. This bit is always set in the M5100 to indicate that the MAC2008, ABS2008, NAN2008 bits within the FCSR register exist. For more information, refer to Section 3.5.5 “Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)” .	R	1						
F64	22	Indicates that this is a 64-bit FPU: 0: Not a 64-bit FPU 1: A 64-bit FPU. 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: 0: Long type not implemented 1: Long implemented 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: 0: Word type not implemented 1: Word implemented This bit is always 1 to indicate that word fixed point data types are implemented.	R	1						

Table 3.6 FIR Bit Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
3D	19	Indicates that the MIPS-3D ASE is implemented: 0: MIPS-3D not implemented 1: MIPS-3D implemented 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: 0: PS floating-point not implemented 1: PS floating-point implemented 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: 0: D floating-point not implemented 1: D floating-point implemented This bit is always 1 to indicate that double-precision floating-point data types are implemented.	R	1
S	16	Indicates that the single-precision (S) floating-point data type and instructions are implemented: 0: S floating-point not implemented 1: S floating-point implemented This bit is always 1 to indicate that single-precision floating-point data types are implemented.	R	1
Processor ID	15:8	This value matches the corresponding field of the CP0 PRId register.	R	0xA6
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	Preset
0	31:25, 23	These bits must be written as zeros; they return zeros on reads.	0	0

3.5.2 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 3.12](#) shows the format of the *FCCR*; [Table 3.7](#) describes the *FCCR* bit fields.

Figure 3.12 FCCR Format




Table 3.7 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 Section 3.5.5 “Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)” .	R/W	Undefined
0	31:8	These bits must be written as zeros; they return zeros on reads.	0	0

3.5.3 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 3.13](#) shows the format of the *FEXR*; [Table 3.8](#) describes the *FEXR* bit fields.

Figure 3.13 FEXR Format




Table 3.8 FEXR Bit Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
Cause	17:12	Cause bits. Refer to the description of this field in Section 3.5.5, "Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)" .	R/W	Undefined
Flags	6:2	Flag bits. Refer to the description of this field in Section 3.5.5 “Floating-Point Control and Status Register (FCSR, CP1 Control Register 31)” .	R/W	Undefined
0	31:18, 11:7, 1:0	These bits must be written as zeros; they return zeros on reads.	0	0

3.5.4 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 3.14](#) shows the format of the *FENR*; [Table 3.9](#) describes the *FENR* bit fields.

Figure 3.14 FENR Format




Table 3.9 FENR Bit Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
Enables	11:7	Enable bits. Refer to the description of this field in Section 3.5.5 “Floating-Point Control and Status Register (FCSR, CPI Control Register 31)” .	R/W	Undefined
FS	2	Flush to Zero bit. Refer to the description of this field in Section 3.5.5, “Floating-Point Control and Status Register (FCSR, CPI Control Register 31)” .	R/W	Undefined
RM	1:0	Rounding mode. Refer to the description of this field in Section 3.5.5, “Floating-Point Control and Status Register (FCSR, CPI Control Register 31)” .	R/W	Undefined
0	31:12, 6:3	These bits must be written as zeros; they return zeros on reads.	0	0

3.5.5 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 3.15](#) shows the format of the *FCSR*; [Table 3.10](#) describes the *FCSR* bit fields.

Figure 3.15 FCSR Format




Table 3.10 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). Refer to Section 3.5.6 “Operation of the FS/FO/FN Bits” for more details on this bit.	R/W	Undefined
FO	22	Flush Override (FO). Refer to Section 3.5.6 “Operation of the FS/FO/FN Bits” for more details on this bit.	R/W	Undefined
FN	21	Flush to Nearest (FN). Refer to Section 3.5.6 “Operation of the FS/FO/FN Bits” for more details on this bit.	R/W	Undefined
MAC2008	20	Fused multiply-add mode, compliant with IEEE Standard 754-2008. The fused multiply-add operation multiplies and adds with unbounded range and precision, rounding only once to the destination format. The fused multiply-add is not supported in the M5100 core. M5100 implements the unfused multiply-add, which rounds the intermediary multiplication result to the destination format. This field applies to the MADD fmt, NMADD fmt, MSUB.fmt, and NMSUB fmt instructions. 0: Unfused multiply-add 1: IEEE 754-2008 fused multiply-add	R	0
ABS2008	19	ABS.fmt & NEG.fmt instructions compliant with IEEE Standard 754-2008. The IEEE 754-2008 standard requires that the ABS and NEG functions accept QNaN inputs without trapping. This bit is always set in the M5100 core to indicate support for the IEEE 754-2008 standard. 0: ABS & NEG trap for QNaN input 1: ABS & NEG accept QNaN input without trapping. IEEE 754-2008 behavior.	R	1
NAN2008	18	Quiet and signaling NaN encodings recommended by the IEEE Standard 754-2008, i.e. a quiet NaN is encoded with the first bit of the fraction being 1 and a signaling NaN is encoded with the first bit of the fraction field being 0. In the M5100 core, this bit is always set to indicate support for the IEEE Standard 754-2008 encoding. 0: MIPS NaN encoding 1: IEEE 754-2008 NaN encoding	R	1

Table 3.10 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 Table 3.11 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 Table 3.11 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 Table 3.11 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 Table 3.12 for the encoding of this field.	R/W	Undefined
0	20:18	These bits must be written as zeros; they return zeros on reads.	0	0

Table 3.11 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

Table 3.11 Cause, Enables, and Flags Definitions (Continued)

Bit Name	Bit Meaning
I	Inexact

Table 3.12 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.

3.5.6 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 NEG¹

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 3.16 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 3.16 FS/FO/FN Bits Influence on Multiply and Addition Results

3.5.6.1 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. Table 3.13 lists the flushing behavior for tiny results..

Table 3.13 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 3.14.

Table 3.14 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.

3.5.6.2 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 3.15.

Table 3.15 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 3.15 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 3.32) 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.

3.5.6.3 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 3.17). 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 3.17 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 3.16.

Table 3.16 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 3.32) 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.

3.5.6.4 Recommended FS/FO/FN Settings

Table 3.17 summarizes the recommended FS/FO/FN settings.

Table 3.17 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 3.17 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. ¹

1. Note that in this mode, MADD might return a different result other than the equivalent MUL and ADD operation sequence.

3.5.7 FCSR Cause Bit Update Flow

3.5.7.1 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.

3.5.7.2 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.

3.5.7.3 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 (no flushing happens on tiny results if FO = 1).
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 [Section 3.5.7.2 “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.

3.5.7.4 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.

3.5.7.5 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.

3.6 Instruction Overview

The functional groups into which the FPU instructions are divided are described in the following subsections:

- [Section 3.6.1 “Data Transfer Instructions”](#)
- [Section 3.6.2 “Arithmetic Instructions”](#)
- [Section 3.6.3 “Conversion Instructions”](#)
- [Section 3.6.4 “Formatted Operand-Value Move Instructions”](#)
- [Section 3.6.5 “Conditional Branch Instructions”](#)
- [Section 3.6.6 “Miscellaneous Instructions”](#)

The instructions are described in detail in [Chapter 13, “M5100™ Processor Core Instructions”](#) on page 328, including descriptions of supported formats (fmt).

3.6.1 Data Transfer Instructions

The FPU has two separate register sets: coprocessor general registers (FPRs) and 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 3.18 lists the supported transfer operations.

Table 3.18 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

3.6.1.1 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), 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.

3.6.1.2 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 3.19 through 3.20 list the FPU data transfer instructions.

Table 3.19 FPU Loads and Stores

Mnemonic	Instruction	Addressing Mode
LDC1	Load Doubleword to Floating Point	Register+offset
LWC1	Load Word to Floating Point	Register+offset
SDC1	Store Doubleword from Floating Point	Register+offset
SWC1	Store Word from Floating Point	Register+offset
LDXC1	Load Doubleword Indexed to Floating Point	Register+Register
LUXC1	Load Doubleword Indexed Unaligned to Floating Point	Register+Register
LWXC1	Load Word Indexed to Floating Point	Register+Register
SDXC1	Store Doubleword Indexed from Floating Point	Register+Register
SUXC1	Store Doubleword Indexed Unaligned from Floating Point	Register+Register
SWXC1	Store Word Indexed from Floating Point	Register+Register

Table 3.20 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
MFHC1	Move Word From High Half of Floating Point
MTC1	Move Word To Floating Point

Table 3.20 FPU Move To and From Instructions (Continued)

Mnemonic	Instruction
MTHC1	Move Word to High Half of Floating Point

3.6.2 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).

In general, the arithmetic instructions take an Unimplemented Operation exception for denormalized numbers, except for the ABS, C, and NEG instructions, which can handle denormalized numbers. The FS, FO, and FN bits in the CP1 *FCSR* register can override this behavior as described in Section 3.5.6 “Operation of the FS/FO/FN Bits”.

Table 3.21 lists the FPU IEEE compliant arithmetic operations.

Table 3.21 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 3.22 lists the FPU-approximate arithmetic operations.

Table 3.22 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 3.23. 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, subtract, or negate instructions.

Table 3.23 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

3.6.3 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.

In general, the conversion instructions only take an Unimplemented Operation exception for denormalized numbers. The FS and FN bits in the CP1 *FCSR* register can override this behavior as described in [Section 3.5.6 “Operation of the FS/FO/FN Bits”](#).

[Table 3.24](#) and [Table 3.25](#) list the FPU conversion instructions according to their rounding mode.

Table 3.24 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 3.25 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

3.6.4 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* [7] and *microMIPS32™ Architecture Reference Manual, Volume II* [8].)

Table 3.26 through Table 3.28 list the formatted operand-value move instructions.

Table 3.26 FPU Formatted Operand Move Instruction

Mnemonic	Instruction
MOV fmt	Floating-Point Move

Table 3.27 FPU Conditional Move on True/False Instructions

Mnemonic	Instruction
MOVf fmt	Floating-Point Move Conditional on FP False
MOVt fmt	Floating-Point Move Conditional on FP True

Table 3.28 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

3.6.5 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.

The MIPS64 architecture defines eight condition codes for use in compare and branch instructions. For backward compatibility with previous revisions of the ISA, condition code bit 0 and condition code bits 1 through 7 are in discontinuous fields in the *FCSR*.

Table 3.29 lists the conditional branch (branch and branch likely) FPU instructions; Table 3.30 lists the deprecated conditional branch likely instructions.

Table 3.29 FPU Conditional Branch Instructions

Mnemonic	Instruction
BC1F	Branch on FP False
BC1T	Branch on FP True

Table 3.30 Deprecated FPU Conditional Branch Likely Instructions

Mnemonic	Instruction
BC1FL	Branch on FP False Likely
BC1TL	Branch on FP True Likely

3.6.6 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 3.31 lists these conditional move instructions.

Table 3.31 CPU Conditional Move on FPU True/False Instructions

Mnemonic	Instruction
MOVN	Move Conditional on FP False
MOVZ	Move Conditional on FP True

3.7 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.

3.7.1 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 3.32](#)). 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*.

3.7.2 Exception Conditions

The subsections below describe the following five exception conditions defined by IEEE Standard 754:

- [Section 3.7.2.1 “Invalid Operation Exception”](#)
- [Section 3.7.2.2 “Division By Zero Exception”](#)
- [Section 3.7.2.3 “Underflow Exception”](#)
- [Section 3.7.2.4 “Overflow Exception”](#)
- [Section 3.7.2.5 “Inexact Exception”](#)

[Section 3.7.2.6 “Unimplemented Operation Exception”](#) also describes a MIPS-specific exception condition, Unimplemented Operation Exception, that 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.

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 excep-

tion 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 3.32 summarizes the default results.

Table 3.32 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: 0 (RN) and 1 (RZ): Supplies a zero with the sign of the exact result. 2 (RP): For positive underflow values, supplies $2^{E_{\min}}$ (MinNorm). For negative underflow values, supplies a positive zero. 3 (RM): For positive underflow values, supplies a negative zero. For negative underflow values, supplies a negative $2^{E_{\min}}$ (MinNorm). Note that this behavior is only valid if the $FCSR_{FN}$ 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: 0 (RN): Supplies an infinity with the sign of the exact result. 1 (RZ): Supplies the format's largest finite number with the sign of the exact result. 2 (RP): For positive overflow values, supplies positive infinity. For negative overflow values, supplies the format's most negative finite number. 3 (RM): For positive overflow values, supplies the format's largest finite number. For negative overflow values, supplies minus infinity.

3.7.2.1 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 ∞/∞ , 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.

3.7.2.2 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.

3.7.2.3 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.

3.7.2.4 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.

3.7.2.5 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.

3.7.2.6 Unimplemented Operation Exception

The Unimplemented Operation exception is a MIPS-defined exception that provides software emulation support. This exception is not IEEE-compliant.

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 handed by the FS/FO/FN bits.

3.8 Pipeline and Performance

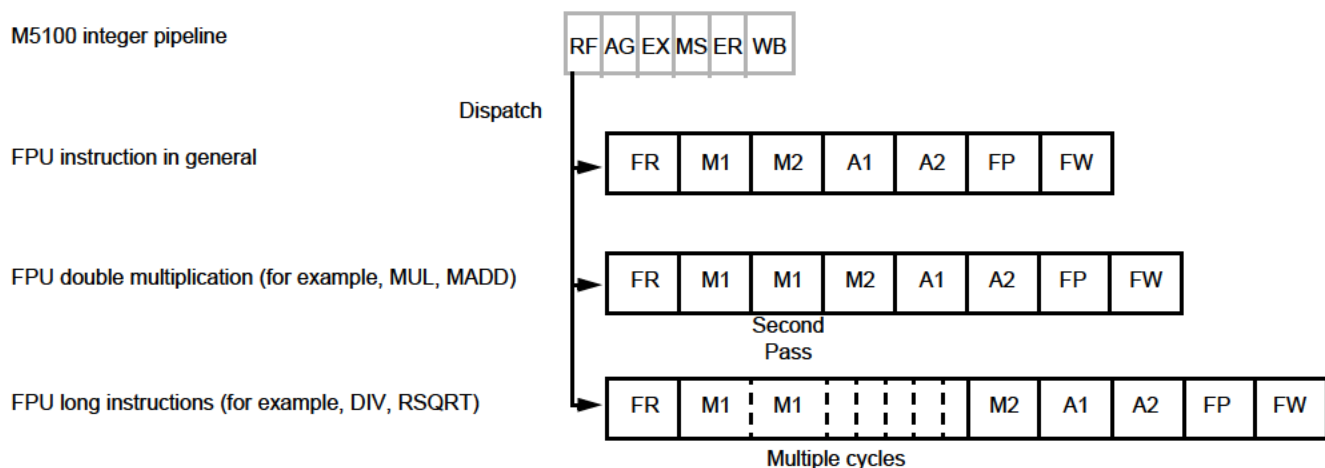
This section describes the structure and operation of the FPU pipeline.

3.8.1 Pipeline Overview

The FPU has a seven stage pipeline to which the integer pipeline dispatches instructions. The FPU pipeline runs in parallel with the M5100 integer pipeline. The FPU can be built to run at either the same frequency as the integer core or at one-half the frequency of the integer core.

The FPU pipe is optimized for single-precision instructions, such that the basic multiply, ADD/SUB, and MADD/MSUB instructions can be performed with single-cycle throughput and low latency. Executing double-precision multiply and MADD/MSUB instructions requires a second pass through the M1 stage to generate all 64 bits of the product. Executing long latency instructions, such as DIV and RSQRT, extends the M1 stage. Figure 3.18 shows the FPU pipeline.

Figure 3.18 FPU Pipeline



3.8.1.1 FR Stage — Decode, Register Read, and Unpack

The FR stage has the following functionality:

- The dispatched instruction is decoded for register accesses.

Floating-Point Unit of the M5100™ Core

- Data is read from the register file.
- The operands are unpacked into an internal format.

3.8.1.2 M1 Stage — Multiply Tree

The M1 stage has the following functionality:

- A single-cycle multiply array is provided for single-precision data format multiplication, and two cycles are provided for double-precision data format multiplication.
- The long instructions, such as divide and square root, iterate for several cycles in this stage.
- Sum of exponents is calculated.

3.8.1.3 M2 Stage — Multiply Complete

The M2 stage has the following functionality:

- Multiplication is complete when the carry-save encoded product is compressed into binary.
- Rounding is performed.
- Exponent difference for addition path is calculated.

3.8.1.4 A1 Stage — Addition First Step

This stage performs the first step of the addition.

3.8.1.5 A2 Stage — Addition Second and Final Step

This stage performs the second and final step of the addition.

3.8.1.6 FP Stage — Result Pack

The FP stage has the following functionality:

- The result coming from the datapath is packed into IEEE 754 Standard format for the FPR register file.
- Overflow and underflow exceptional conditions are resolved.


3.8.1.7 FW Stage — Register Write

The result is written to the FPR register file.

3.8.2 Bypassing

The FPU pipeline implements extensive bypassing, as shown in [Figure 3.19](#). Results do not need to be written into the register file and read back before they can be used, but can be forwarded directly to an instruction already in the pipe. Some bypassing is disabled when operating in 32-bit register file mode, the FP bit in the CP0 *Status* register is 0, due to the paired even-odd 32-bit registers that provide 64-bit registers.

Figure 3.19 Arithmetic Pipeline Bypass Paths



3.8.3 Repeat Rate and Latency

Table 3.33 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 3.33 M5100™ Core FPU Latency and Repeat Rate

Opcode ¹	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 fnt and MOVT fnt instruction / MOVT, MOVN, BC1 instruction	1 / 2	1
CVT.D.S, CVT.[S,D].[W,L]	4	1
CVT.S.D	6	1
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.

3.9 2008 FPU Support

The M5100 implements the following new status/control bits to provide greater compatibility with the new IEEE Standard 754 Floating-point released in 2008, and to maintain compatibility with previous FPU implementations.

Floating-Point Unit of the M5100™ Core

- The *Has2008* bit in *FIR* will always read as 1 to signify that 2008 FPU is implemented.
- The *MAC2008* bit in *FCSR* will always read as 0 to signify that Fused Multiply Add operation is not yet implemented.
- The *ABS2008* bit in *FCSR* can be written as 1 (2008-IEEE). When *ABS2008* = 1, which makes ABS and NEG instructions non-arithmetic instructions. All floating-point exceptions will be disabled.
- The *NAN2008* bit in *FCSR* can be written as 1 (2008-IEEE). When *NAN2008* = 1, it flips SNaN and QNaN decoding from the former implementation. In addition, the following behaviors are implemented:
 - In the case of one or more QNaN operands (no SNaN), the QNaN operand is propagated from one of the input operands (in order of priority): fs, ft, and fr.
 - When SNaN is used as an input, and exceptions are disabled, QNaN is the expected output.
 - The QNaN output will not be a fixed value. To comply with IEEE, an input NaN should produce a NaN with the payload of the input NaN if representable in the destination format, where the payload is defined as the Mantissa field less its most-significant bit.
 - If *ABS2008*=1 and *MAC2008*=0 (as it always is in MUP/MUC), the sign of NMADD and NMSUB do not flip the sign of any QNaN input, and the sign is retained and propagated to the output.
 - When a NaN is an input, the output will be one of the input NaNs with as much of the mantissa preserved as possible.
 - SNaN inputs have higher priority than QNaN inputs and then fs has higher priority than ft which has higher priority than fr.
 - The sign of the selected NaN input is preserved. If the input that is selected for the output is already a QNaN, then the entire mantissa is preserved. However, if the input that is selected for the output is an SNaN, then the most significant bit of the SNaN mantissa is complemented to convert the SNaN into a QNaN. If this conversion to a QNaN would result in an infinity, then the next most significant bit of the mantissa is set.
 - For CVT.s.d, the NaN mantissa msbs are preserved. For CVT.d.s, the NaN mantissa is padded with 0's in the lsbs.
 - For mult-add, if both fs/ft and fr are QNaNs, then the multiply produces a QNaN based upon fs/ft, and this QNaN has priority over fr in the add operation. However, if both fs/ft and fr are SNaNs and the invalid trap is not enabled, then the multiply generates a QNaN based upon fs/ft, which is then added to the signaling fr and the signaling fr has priority
 - When a NaN is needed for output but there is no NaN input, a positive QNaN is created that has all other mantissa bits set.

The MIPS® DSP Module

The M5100 includes support for the MIPS DSP Module 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.

Refer to *MIPS® Architecture For Programmers Volume IV-e* [16] or [16] for a general description of the DSP Module and detailed descriptions of the DSP instructions. Additional programming information is contained in *Five Methods of Utilizing the MIPS® DSP Module* [18], *Efficient DSP Module Programming in C: Tips and Tricks* [19], and *Accelerating DSP Filter Loops with MIPS® CorExtend® Instructions* [20].

4.1 Additional Register State for the DSP Module

The DSP Module 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 Module and to include these additional registers in the context save and restore operations.

4.1.1 HI-LO Registers

The DSP Module 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 instructions that target the accumulators use two instruction bits to specify the destination accumulator, with the zero value referring to the original accumulator.

4.1.2 DSPControl Register

The *DSPControl* register contains control and status information used by DSP instructions. [Figure 4.1](#) illustrates the bits in this register, and [Table 4.1](#) describes their usage.

Figure 4.1 MIPS32® DSP Module Control Register (DSPControl) Format

31	28	27	24	23	16	15	14	13	12	7	6	5	0
0	ccond		ouflag		0	EFI	c	scount		0	pos		

Table 4.1 MIPS® DSP Module Control Register (DSPControl) Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
0	31:28	Reserved. Used in the MIPS64 architecture but not used in the MIPS32 architecture. Must be written as zero; returns zero on read.	0	0	Required
<i>ccond</i>	27:24	Condition code bits set by compare instructions. The compare instruction sets the right-most bits as required by the number of elements in the vector compare. Bits not set by the instruction remain unchanged.	R/W	0	Required
<i>ouflag</i>	23:16	This field is written by hardware when certain instructions overflow or underflow and may have been saturated. See Table 4.2 for a full list of which bits are set by what instructions.	R/W	0	Required
<i>EFI</i>	14	Extract Fail Indicator. This bit is set to 1 when an EXTP, EXTPV, EXTPDP, or EXTPDPV instruction fails. These instructions fail when there are insufficient bits to extract, that is, when the value of <i>pos</i> in <i>DSPControl</i> is less than the value of <i>size</i> specified in the instruction. 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	0	Required
<i>c</i>	13	Carry bit. This bit is set and used by special add instructions that implement a 64-bit add across two GPRs. The ADDSC instruction sets the bit and the ADDWC instruction uses this bit.	R/W	0	Required
<i>scount</i>	12:7	This field is for use by the INSV instruction. The value of this field is used to specify the size of the bit field to be inserted.	R/W	0	Required
<i>pos</i>	5:0	This field is used by the variable insert instructions INSV to specify the insert position. It is also used to indicate the extract position for the EXTP, EXTPV, EXTPDP, and EXTPDPV instructions. The decrement <i>pos</i> (DP) variants of these instructions on completion will have decremented the value of <i>pos</i> (by the size amount). The MTHLIP instruction will increment the <i>pos</i> value by 32 after copying the value of <i>LO</i> to <i>HI</i> .	R/W	0	Required
0	15:13	Must be written as zero; returns zero on read.	0	0	Reserved

The bits of the overflow flag *ouflag* field in the *DSPControl* register are set by a number of instructions, as described in Table 4.2. These bits are sticky and can be reset only by an explicit write to these bits in the register (using the *WRDSP* instruction).

Table 4.2 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: <i>DPAQ_S</i> , <i>DPAQ_SA</i> , <i>DPSQ_S</i> , <i>DPSQ_SA</i> , <i>DPAQX_S</i> , <i>DPAQX_SA</i> , <i>DPSQX_S</i> , <i>DPSQX_SA</i> , <i>MAQ_S</i> , <i>MAQ_SA</i> and <i>MULSAQ_S</i> .
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: <i>ABSQ_S</i> , <i>ADDQ</i> , <i>ADDQ_S</i> , <i>ADDU</i> , <i>ADDU_S</i> , <i>ADDWC</i> , <i>SUBQ</i> , <i>SUBQ_S</i> , <i>SUBU</i> and <i>SUBU_S</i> .
21	Instructions that set this bit on an overflow/underflow: <i>MUL</i> , <i>MUL_S</i> , <i>MULEQ_S</i> , <i>MULEU_S</i> , <i>MULQ_RS</i> , and <i>MULQ_S</i> .
22	Instructions that set this bit on an overflow/underflow: <i>PRECRC_RS</i> , <i>SHLL</i> , <i>SHLL_S</i> , <i>SHLLV</i> , and <i>SHLLV_S</i> .
23	Instructions that set this bit on an overflow/underflow: <i>EXTR</i> , <i>EXTR_S</i> , <i>EXTR_RS</i> , <i>EXTRV</i> , and <i>EXTRV_RS</i> .

4.2 Software Detection of the DSP Module Revision 2

The presence of the MIPS DSP Module in the M5100 is indicated by two static bits in the *Config3* register: the *DSPP* (*DSP Present*) bit indicates the presence of the DSP Module, and the *DSP2P* (*DSP Rev2 Present*) bit indicates the presence of the MIPS DSP Module Rev2. Because the DSP Module is configurable in the M5100 processor core, and it always comes with the DSP Module Rev2 if the DSP Module is configured, therefore the *DSPP* and *DSP2P* are always preset to 0's or 1's.

The *MX* (*DSP Module Enable*) read/write bit in the *CP0 Status* register must be set to enable access to the additional instructions defined by the DSP Module, as well as to the *MTLO/HI*, *MFLO/HI* instructions that access accumulators *ac1*, *ac2*, and *ac3*. Executing a DSP Module instruction or the *MTLO/HI*, *MFLO/HI* instructions with this bit set to zero causes a DSP State Disabled Exception (exception code 26 in the *CP0 Cause* register). This exception can be used by system software to do lazy context switching.

The MIPS® Virtualization Module

The M5100™ core includes the MIPS Architecture Virtualization (VZ) Module that enables the virtualization of operating systems, thus providing a scalable, trusted, and secure execution environment. M5100 Virtualization defines a set of extensions to the MIPS32 Architecture for efficient implementation of virtualized systems.

Virtualization is enabled by software—the key element is a control program known as a Virtual Machine Monitor (VMM) or hypervisor. The hypervisor is in full control of machine resources at all times. When an operating system (OS) kernel runs within a virtual machine (VM), it becomes a guest of the hypervisor. All operations performed by a guest must be explicitly permitted by the hypervisor. To ensure that it remains in control, the hypervisor always runs at a higher level of privilege than a guest operating system kernel. The hypervisor is responsible for managing access to sensitive resources, maintaining the expected behavior for each VM, and sharing resources between multiple VMs.

In a traditional operating system, the kernel (or supervisor) typically runs at a higher level of privilege than user applications. The kernel provides a protected virtual-memory environment for each user application, inter-process communications, IO device sharing and transparent context switching. The hypervisor performs the same basic functions in a virtualized system, except that the hypervisor's clients are full operating systems rather than user applications.

The virtual machine execution environment created and managed by the hypervisor consists of the full Instruction Set Architecture, including all Privileged Resource Architecture facilities, plus any device-specific or board-specific peripherals and associated registers. It appears to each guest operating system as if it is running on a real machine with full and exclusive control.

The Virtualization Module enables full virtualization, and is intended to allow VM scheduling to take place while meeting real-time requirements, and to minimize costs of context switching between VMs

Refer to the *MIPS® Architecture for Programmers Volume IV-i: Virtualization Module of the MIPS32® Architecture* [14] or *MIPS® Architecture for Programmers Volume IV-i: Virtualization Module of the microMIPS32® Architecture* [15]. for a general description of the VZ Module and detailed descriptions of the VZ instructions. Refer also to [Chapter 6, “Memory Management of the M5100™ Core” on page 107](#) for a detailed description of memory protection features provided by the Virtualization Module.

5.1 Additional Register State for the VZ Module

The VZ Module defines six additional CP0 registers and modifications to bit fields in the *Cause*, *Config3*, *WatchHi*, and *PerfCnt* registers. These registers require the operating system to recognize the presence of the VZ Module and to include these additional registers in the context save and restore operation.

[Table 5.1](#) lists the new and modified CP0 registers in numerical order. The individual registers are described in [Chapter 8, “CP0 Registers of the M5100™ Core” on page 151](#).

Table 5.1 Virtualization Module Coprocessor 0 Registers in Numerical Order

Register Number	Sel	Register Name	Modification	Reference
12	6	<i>GuestCtl0</i>	New Register. Controls guest mode behavior.	Section 8.2.15
10	4	<i>GuestCtl1</i>	New Register. Guest ID	Section 8.2.15
10	5	<i>GuestCtl2</i>	New Register. Interrupt related	Section 8.2.16
10	6	<i>GuestCtl3</i>	New Register. GPR Shadow Set related.	Section 8.2.17
11	4	<i>GuestCtl0Ext</i>	Extension to GuestCtl0	Section 8.2.19
12	7	<i>GTOffset</i>	New Register. Guest timer offset.	Section 8.2.27
13	0	<i>Cause</i>	Addition of hypervisor cause code.	Section 8.2.28
16	3	<i>Config3</i>	Identifies Virtualization Module feature set.	Section 8.2.39
19	0	<i>WatchHi</i>	Watch Debug.	Section 8.2.45
25	0	<i>PerfCnt</i>	Performance Counter, adds virtualization support.	Section 8.2.53

5.2 Additional Instructions for the Virtualization Module

The Virtualization Module adds new instructions and modifies existing instructions to allow root-mode access to the guest Coprocessor 0 context and the guest TLB. These instructions are listed in alphabetical order in [Table 5.2](#) and described in detail in [Chapter 13, “M5100™ Processor Core Instructions”](#) on page 328.

Table 5.2 New and Modified Instructions

Mnemonic	Instruction	Description
HYPCALL	Hypercall	Trigger Hypercall exception.
MFGC0	Move from Guest Coprocessor 0	Read guest coprocessor 0 into GPR.
MTGC0	Move from Guest Coprocessor 0	Write guest coprocessor 0 from GPR.
TLBGINV	Guest TLB Invalidate	Trigger guest TLB invalidate from root mode.
TLBGINVF	Guest TLB Invalidate Flush	Trigger guest TLB invalidate from root mode.
TLBGP	Probe Guest TLB	Trigger guest TLB probe from root mode.
TLBGR	Read Guest TLB	Trigger guest TLB read from root mode.
TLBGWI	Write Guest TLB	Trigger guest TLB write from root mode.
TLBGWR	Write Guest TLB	Trigger guest TLB write from root mode.
TLBINV	TLB Invalidate	Modified TLB Invalidate behavior.
TLBINVF	TLB Invalidate Flush	Modified TLB Invalidate Flush behavior.
TLBP	TLB Probe	Modified TLB probe behavior.
TLBR	Read TLB	Modified TLB read behavior.
TLBWI	Write TLB, Indexed	Modified indexed TLB write behavior.
TLBWR	Write TLB, Random	Modified random TLB write behavior.

5.3 Software Detection of the VZ Module

Software can determine if the Virtualization Module is implemented by checking the state of the *VZ (Virtualization Module Enable)* bit in the *Config3* CP0 register.

Memory Management of the M5100™ Core

The M5100™ processor core includes a Memory Management Unit (MMU) that interfaces between the execution unit and the cache controller. The core implements a simple Fixed Mapping Translation (FMT) style MMU.

This chapter contains the following sections:

- [Section 6.1 “Introduction”](#)
- [Section 6.2 “Modes of Operation”](#)
- [Section 6.3 “Fixed Mapping MMU”](#)
- [Section 6.4 “System Control Coprocessor”](#)

6.1 Introduction

The MMU in a M5100 processor core translates a virtual address to a physical address before the request is sent to the SRAM interface for an external memory reference.

In the M5100 processor core, the MMU is based on a simple algorithm to translate virtual addresses to physical addresses via a Fixed Mapping Translation (FMT) mechanism. These translations are different for various regions of the virtual address space (useg/kuseg, kseg0, kseg1, kseg2/3).

6.1.1 Memory Management Unit (MMU) in Guest Modes

The M5100 core contains a simple Fixed Mapping Translation (FMT) MMU that interfaces between the execution unit and the SRAM controller.

6.1.1.1 Fixed Mapping Translation (FMT)

An FMT is smaller and simpler than the full Translation Lookaside Buffer (TLB) style MMU found in other MIPS cores. Like a TLB, the FMT performs virtual-to-physical address translation and provides attributes for the different segments. Those segments that are unmapped in a TLB implementation (kseg0 and kseg1) are translated identically by the FMT.

6.1.2 Memory Management Unit (MMU) in Root Modes

The M5100 core provides a Root Protection Unit (RPU) MMU for guest virtual-address protection. This document treats TLB or MMU the equivalent of an RPU. An example of this is CP0 register descriptions referring to a TLB entry is also referring the RPU entry.

6.1.2.1 Root Protection Unit (RPU)

An RPU-based MMU is similar to an RTLb in that it consists of an 8, 16, or 32 dual-entry fully associative look-aside buffer, but the segments that are unmapped in a TLB implementation (kseg0 and kseg1) are translated identically by the RPU.

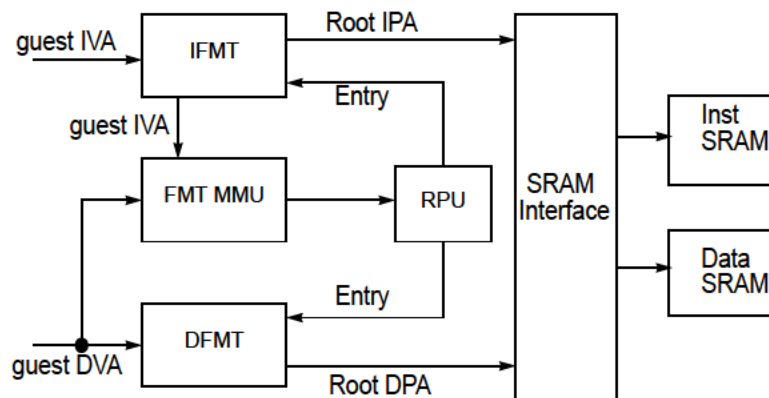
The RPU allows guest pages to be protected by a read-inhibit, execute-inhibit, and write-protection attribute.

When an instruction address is calculated, the virtual address is first compared to the contents of the 4-entry IFMT. If the address is not found, the virtual address is fixed mapped and validated by the RPU. If the address is not found in the RPU, a refill exception is taken.

When a data address is calculated, the virtual address is compared to the 4-entry DFMT. If the address is not found, the virtual address is fixed mapped and validated by the RPU. If the address is not found in the RPU, a refill exception is taken.

Figure 6.1 shows how the memory management unit interacts with the SRAM access in the M5100 core.

Figure 6.1 Address Translation During SRAM Access



6.2 Modes of Operation

The M5100 core implements the Virtualization Module by supporting five modes of operation:

- *Guest-user mode* is most often used for applications programs.
- *Guest-kernel mode* is typically used for handling exceptions and operating-system kernel functions, including CP0 management and I/O device accesses.
- *Root-user mode* is used to run hypervisor (Virtual Machine Monitor (VMM)) user software.
- *Root-kernel mode* is used to run hypervisor (Virtual Machine Monitor (VMM)) kernel software.
- *Debug mode* is used during system bring-up and software development. Refer to the EJTAG section for more information on debug mode.

User mode is most often used for application programs. 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 most likely occurs within a software development tool.

The address translation performed by the MMU depends on the mode in which the processor is operating.

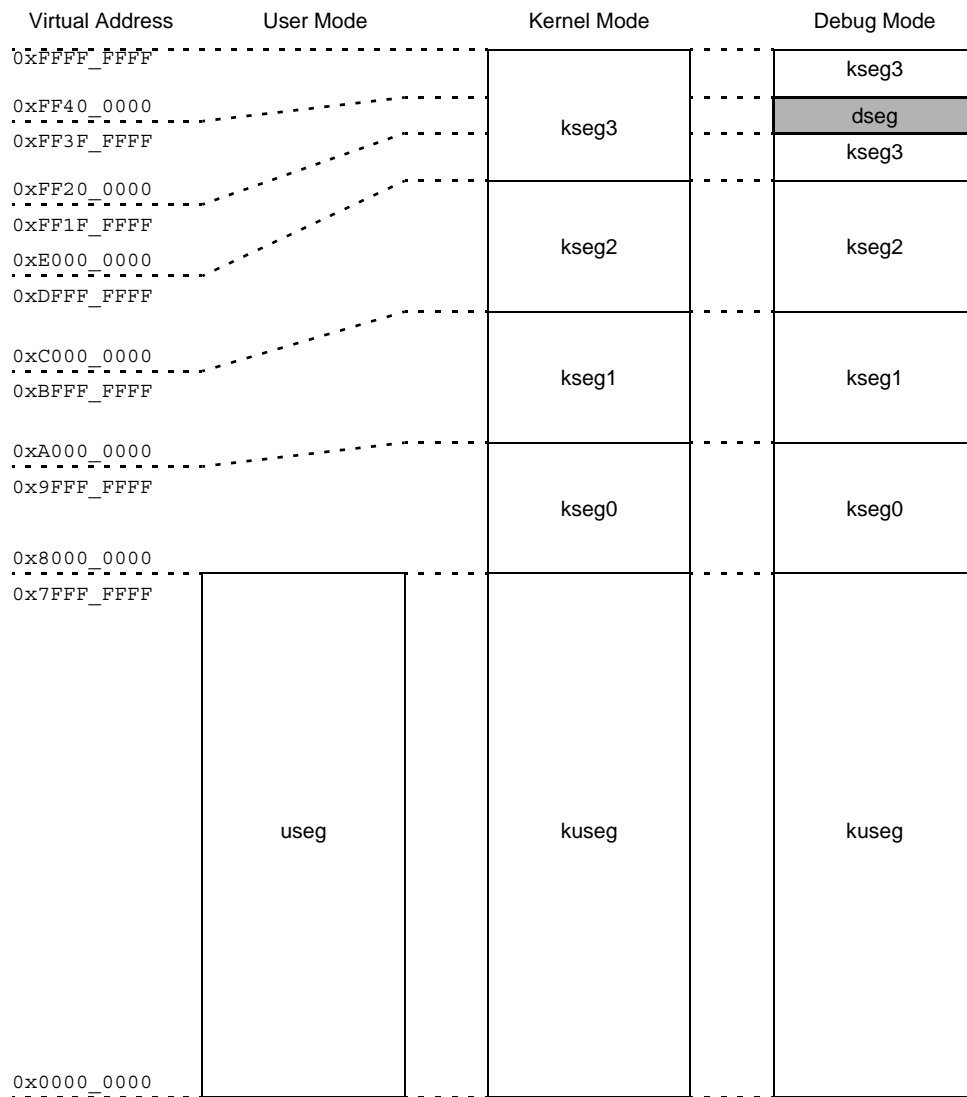
6.2.1 Virtual Memory Segments

The Virtual memory segments differ depending on the mode of operation. [Figure 6.2](#) shows the segmentation for the 4 GByte (2^{32} bytes) virtual memory space addressed by a 32-bit virtual address, for the three modes of operation.

The core enters Kernel mode both at reset and when an exception is recognized. While in Kernel mode, software has access to the entire address space, as well as all CP0 registers. 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.

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 for Kernel mode. In addition, while in Debug mode the core has access to the debug segment dseg. This area overlays part of the kernel segment kseg3. dseg access in Debug mode can be turned on or off, allowing full access to the entire kseg3 in Debug mode, if so desired.

Figure 6.2 M5100™ processor core Virtual Memory Map



Each of the segments shown in Figure 6.2 are either mapped or unmapped. The following two sub-sections explain the distinction. Then sections 6.2.2 “User Mode”, 6.2.3 “Kernel Mode” and 6.2.4 “Debug Mode” specify which segments are actually mapped and unmapped.

6.2.1.1 Unmapped Segments

An unmapped segment does not use the FMT to translate from virtual-to-physical addresses.

Unmapped segments have a fixed simple translation from virtual to physical address. This is much like the translations the FMT provides for the M5100 core, but we will still make the distinction.

All segments are treated as uncached within the M5100 core. Cache coherency attributes of cached or uncached can be specified and this information will be sent with the request to allow the system to make a distinction between the two.

6.2.1.2 Mapped Segments


A mapped segment does use the FMT to translate from virtual-to-physical addresses.

For the M5100 core, the mapped segments have a fixed translation from virtual to physical address. The cacheability of the segment is defined in the CP0 *Config* register fields *K23* and *KU* (see 8.2.36 “Config Register (CP0 Register 16, Select 0)”). Write protection of segments is not possible during FMT translation.

6.2.2 User Mode

In user mode, a single 2 GByte (2^{31} bytes) uniform virtual address space called the user segment (useg) is available. Figure 6.3 shows the location of user mode virtual address space.

Figure 6.3 User Mode 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.

The processor operates in User mode when the *Status* register contains the following bit values:

- *UM* = 1
- *EXL* = 0
- *ERL* = 0

In addition to the above values, the *DM* bit in the *Debug* register must be 0.

Table 6.1 lists the characteristics of the useg User mode segments.

Table 6.1 User Mode Segments

Address Bit Value	Status Register			Segment Name	Address Range	Segment Size
	Bit Value					
	EXL	ERL	UM			
32-bit A(31) = 0	0	0	1	useg	0x0000_0000 --> 0x7FFF_FFFF	2 GByte (2^{31} bytes)

Memory Management of the M5100™ Core

All valid user mode virtual addresses have their most significant bit cleared to 0, indicating that user mode can only access the lower half of the virtual memory map. Any attempt to reference an address with the most significant bit set while in user mode causes an address error exception.

The system maps all references to *useg* through the FMT.

6.2.3 Kernel Mode

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:

- *UM* = 0
- *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.

Kernel mode virtual address space is divided into regions differentiated by the high-order bits of the virtual address, as shown in [Figure 6.4](#). Also, [Table 6.2](#) lists the characteristics of the Kernel mode segments.

Figure 6.4 Kernel Mode Virtual Address Space

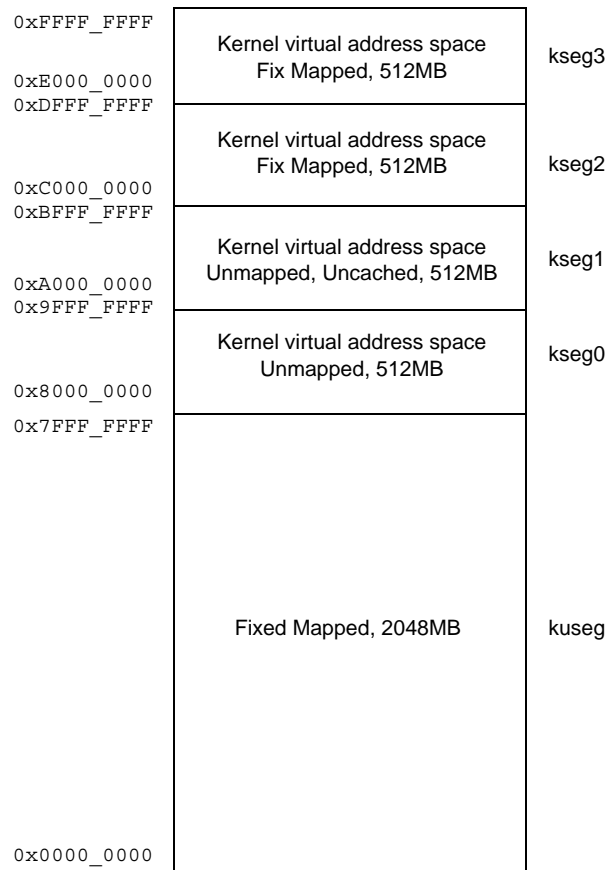


Table 6.2 Kernel Mode Segments

Address Bit Values	Status Register Is One of These Values			Segment Name	Address Range	Segment Size
	UM	EXL	ERL			
A(31) = 0	(UM = 0 or EXL = 1 or ERL = 1) and DM = 0			kuseg	0x0000_0000 through 0x7FFF_FFFF	2 GBytes (2^{31} bytes)
A(31:29) = 100 ₂				kseg0	0x8000_0000 through 0x9FFF_FFFF	512 MBytes (2^{29} bytes)
A(31:29) = 101 ₂				kseg1	0xA000_0000 through 0xBFFF_FFFF	512 MBytes (2^{29} bytes)
A(31:29) = 110 ₂				kseg2	0xC000_0000 through 0xDFFF_FFFF	512 MBytes (2^{29} bytes)
A(31:29) = 111 ₂				kseg3	0xE000_0000 through 0xFFFF_FFFF	512 MBytes (2^{29} bytes)

6.2.3.1 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.

When the *Status* register's *ERL* = 1, the user address region becomes a 2^{29} -byte unmapped and uncached address space. While in this setting, the kuseg virtual address maps directly to the same physical address.

6.2.3.2 Kernel Mode, Kernel Space 0 (kseg0)

In Kernel mode, when the most-significant three bits of the virtual address are 100_2 , 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.

6.2.3.3 Kernel Mode, Kernel Space 1 (kseg1)

In Kernel mode, when the most-significant three bits of the 32-bit virtual address are 101_2 , 32-bit 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.

6.2.3.4 Kernel Mode, Kernel Space 2 (kseg2)

In Kernel mode, when *UM* = 0, *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 110_2 , 32-bit kseg2 virtual address space is selected. In the M5100 core, this 2^{29} -byte (512-MByte) kernel virtual space is located at physical addresses 0xC000_0000 - 0xDFFF_FFFF.


6.2.3.5 Kernel Mode, Kernel Space 3 (kseg3)

In Kernel mode, when the most-significant three bits of the 32-bit virtual address are 111_2 , the kseg3 virtual address space is selected. In the M5100 core, this 2^{29} -byte (512-MByte) kernel virtual space is located at physical addresses 0xE000_0000 - 0xFFFF_FFFF.

6.2.4 Debug Mode

Debug mode address space is identical to Kernel mode address space with respect to mapped and unmapped areas, except for *kseg3*. In *kseg3*, a debug segment *dseg* co-exists in the virtual address range 0xFF20_0000 to 0xFF3F_FFFF. The layout is shown in [Figure 6.5](#).

Figure 6.5 Debug Mode Virtual Address Space



The dseg is sub-divided into the dmseg segment at 0xFF20_0000 to 0xFF2F_FFFF which is used when the 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 for the segments are shown in Table 6.3.

Accesses to memory that would normally cause an exception if tried from kernel mode cause the core to re-enter debug mode via a debug mode exception.

The unmapped kseg0 and kseg1 segments from kernel mode address space are available from debug mode, which allows the debug handler to be executed from uncached and unmapped memory.

Table 6.3 Physical Address and Cache Attributes for dseg, dmseg, and drseg Address Spaces

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	

6.2.4.1 Conditions and Behavior for Access to drseg, EJTAG Registers

The behavior of CPU access to the drseg address range at 0xFF30_0000 to 0xFF3F_FFFF is determined as shown in Table 6.4

Table 6.4 CPU Access to drseg Address Range

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

Memory Management of the M5100™ Core

unpredictable, and writes are ignored to any unimplemented register in the drseg. Refer to [Chapter 11, “EJTAG Debug Support in the M5100™ Core” on page 248](#) for more information on the *DCR*.

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.

6.2.4.2 Conditions and Behavior for Access to dmseg, EJTAG Memory

The behavior of CPU access to the dmseg address range at 0xFF20_0000 to 0xFF2F_FFFF is determined by the table shown in [Table 6.5](#).

Table 6.5 CPU Access to dmseg Address Range

Transaction	ProbEn bit in DCR register	LSNM bit in Debug register	Access
Load / Store	Don't care	1	Kernel mode address space (kseg3)
Fetch	1	Don't care	dmseg
Load / Store	1	0	
Fetch	0	Don't care	See comments below
Load / Store	0	0	

The case with access to the dmseg when the *ProbEn* bit in the *DCR* register is 0 is not expected to happen. 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 happen, the reference hangs until it is satisfied by the probe. The probe can not assume that there will never be a reference to dmseg if 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.

6.3 Fixed Mapping MMU

The M5100 core implements a simple Fixed Mapping (FM) memory management unit that is smaller than the a full translation lookaside buffer (TLB) and more easily synthesized. Like a TLB, the FMT performs virtual-to-physical address translation and provides attributes for the different memory segments. Those memory segments which are unmapped in a TLB implementation (kseg0 and kseg1) are translated identically by the FMT MMU.

The FMT also determines the cacheability of each segment. These attributes are controlled via bits in the *Config* register. [Table 6.6](#) shows the encoding for the K23 (bits 30:28), KU (bits 27:25) and K0 (bits 2:0) of the *Config* register.

The M5100 core does not contain caches and will treat all references as uncached, but these *Config* fields will be sent out to the system with the request and it can choose to use them to control any external caching that may be present..

Table 6.6 Cacheability of Segments with Block Address Translation

Segment	Virtual Address Range	Cacheability
useg/kuseg	0x0000_0000-0x7FFF_FFFF	Controlled by the KU field (bits 27:25) of the <i>Config</i> register.
kseg0	0x8000_0000-0x9FFF_FFFF	Controlled by the K0 field (bits 2:0) of the <i>Config</i> register.

Table 6.6 Cacheability of Segments with Block Address Translation (Continued)

Segment	Virtual Address Range	Cacheability
kseg1	0xA000_0000- 0xBFFF_FFFF	Always uncacheable.
kseg2	0xC000_0000- 0xDFFF_FFFF	Controlled by the K23 field (bits 30:28) of the <i>Config</i> register.
kseg3	0xE000_0000- 0xFFFF_FFFF	Controlled by K23 field (bits 30:28) of the <i>Config</i> register.

The FMT performs a simple translation to map from virtual addresses to physical addresses. This mapping is shown in Figure 6.6. When *ERL*=1, useg and kuseg become unmapped and uncached. The *ERL* behavior is the same as if there was a TLB. The *ERL* mapping is shown in Figure 6.7.

The *ERL* bit is usually never asserted by software. It is asserted by hardware after a Reset, SoftReset or NMI. See 7.8 “Exception Descriptions” on page 138 for further information on exceptions.

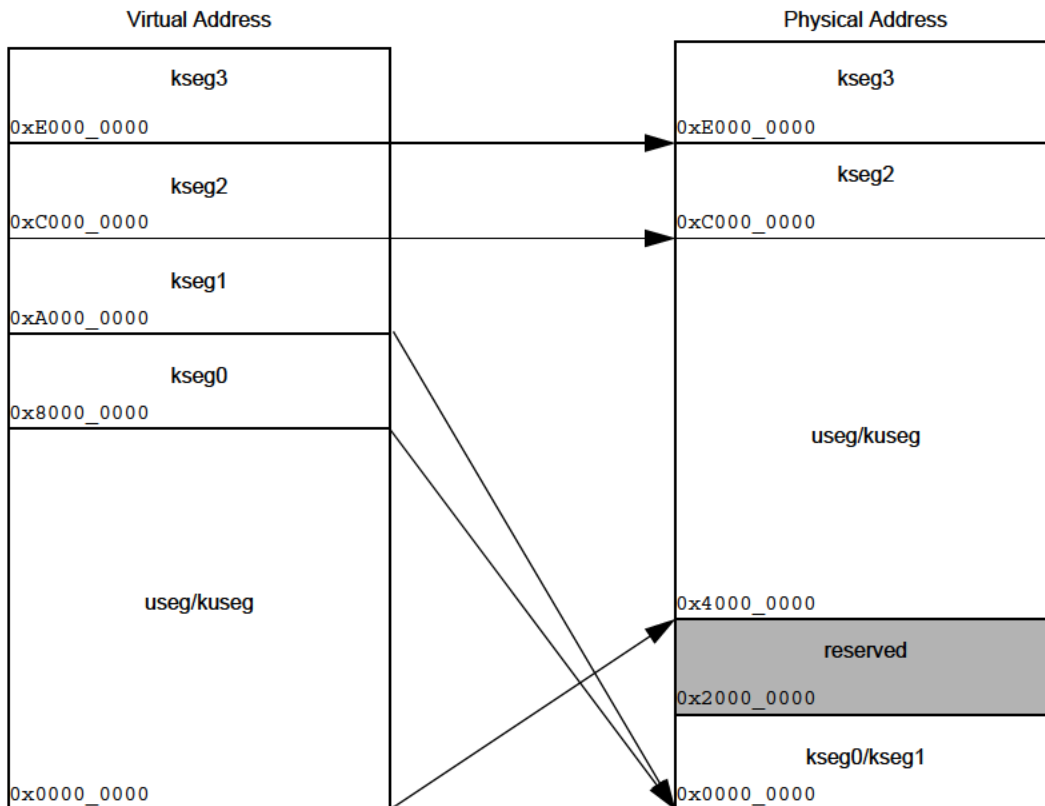
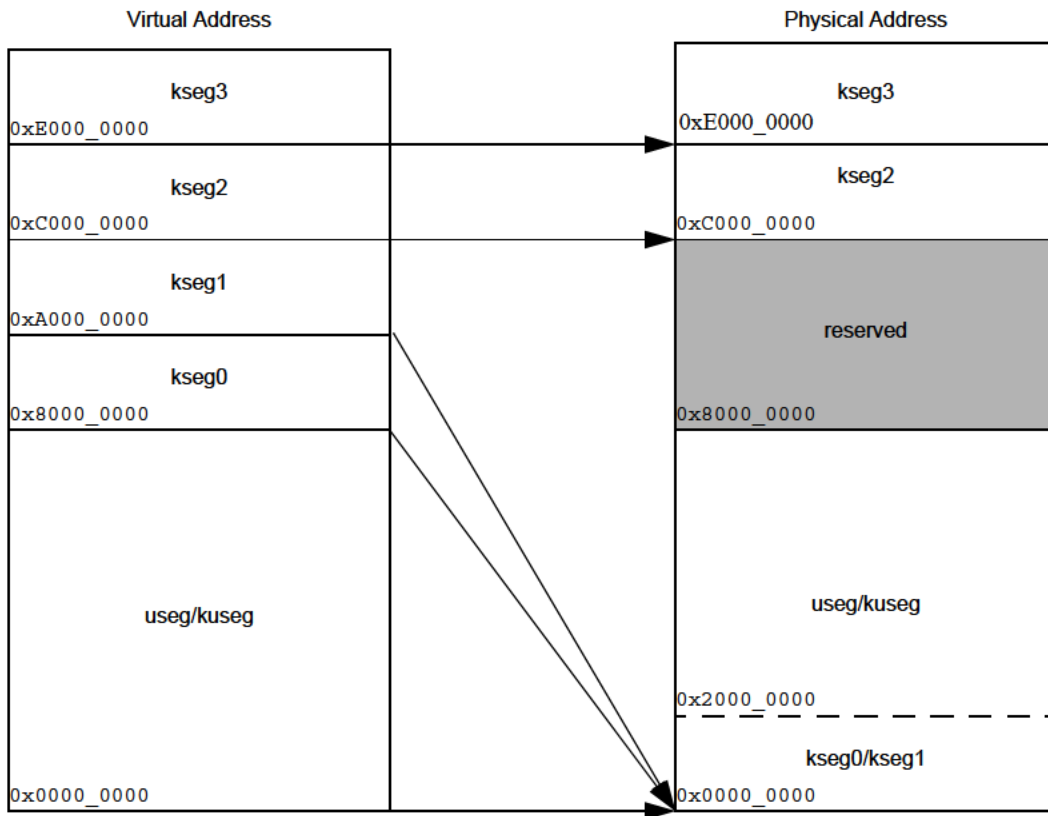
Figure 6.6 FMT Memory Map (ERL=0) in the M5100™ Processor Core

Figure 6.7 FMT Memory Map (ERL=1) in the M5100™ Processor Core



6.4 System Control Coprocessor

The System Control Coprocessor (CP0) is implemented as an integral part of M5100 processor core and supports memory management, address translation, exception handling, and other privileged operations. Certain CP0 registers are used to support memory management. Refer to [Chapter 8, “CP0 Registers of the M5100™ Core”](#) on page 151 for more information on the CP0 register set.

Exceptions and Interrupts in the M5100™ Core

The M5100™ processor core receives exceptions from a number of sources, including arithmetic overflows, I/O interrupts, and system calls. When the CPU detects one of these exceptions, the normal sequence of instruction execution is suspended and the processor enters kernel mode.

In kernel mode the core disables interrupts and forces execution of a software exception processor (called a handler) located at a specific address. The 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.

When an exception occurs, the core loads the *Exception Program Counter (EPC)* register with a location where execution can restart after the exception has been serviced. Most exceptions are *precise*, which mean that *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, the address of the branch instruction immediately preceding the delay slot. To distinguish between the two, software must read the BD bit in the CP0 *Cause* register. Bus error exceptions and CP2 exceptions may be imprecise. For imprecise exceptions the instruction that caused the exception cannot be identified.

This chapter contains the following sections:

- [Section 7.1 “Exception Conditions”](#)
- [Section 7.2 “Exception Priority”](#)
- [Section 7.3 “Interrupts”](#)
- [Section 7.4 “GPR Shadow Registers”](#)
- [Section 7.5 “Exception Vector Locations”](#)
- [Section 7.6 “General Exception Processing”](#)
- [Section 7.7 “Debug Exception Processing”](#)
- [Section 7.8 “Exception Descriptions”](#)
- [Section 7.9 “Exception Handling and Servicing Flowcharts”](#)

7.1 Exception Conditions

When an exception condition occurs, the instruction causing the exception and all those that follow it in the pipeline are cancelled (“flushed”). Accordingly, any stall conditions and any later exception conditions that might have referenced this instruction are inhibited—obviously there is no benefit in servicing stalls for a cancelled instruction.

Exceptions and Interrupts in the M5100™ Core

When an exception condition is detected on an instruction fetch, the core aborts that instruction and all instructions that follow. When this instruction reaches the W stage, various CP0 registers are written with the exception state, change the current program counter (PC) to the appropriate exception vector address, and clearing the exception bits of earlier pipeline stages.

This implementation allows all preceding instructions to complete execution and prevents all subsequent instructions from completing. Thus, the value in the *EPC* (*ErrorEPC* for errors, or *DEPC* for debug exceptions) is sufficient to restart execution. It also ensures that exceptions are taken in the order of execution; an instruction taking an exception may itself be killed by an instruction further down the pipeline that takes an exception in a later cycle.

7.2 Exception Priority

Table 7.1 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. When several exceptions occur simultaneously, the exception with the highest priority is taken.

Table 7.1 Priority of Exceptions

Exception	Description
Reset	Assertion of <i>SI_ColdReset</i> signal.
Soft 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.
NMI	Asserting edge of <i>SI_NMI</i> signal.
Interrupt	Assertion of unmasked hardware or software interrupt signal.
Protection - Instruction fetch	Instruction fetch access to a protected memory region was attempted.
DIB	EJTAG debug hardware instruction break matched.
AdEL	Fetch address alignment error. User-mode fetch reference to kernel address.
ISRAM Parity Error	Parity error on I-SRAM access
IBE	Instruction fetch bus error.
Instruction Validity Exceptions	An instruction could not be completed because it was not allowed access to the required resources (Coprocessor Unusable) or was illegal (Reserved Instruction). If both exceptions occur on the same instruction, the Coprocessor Unusable Exception takes priority over the Reserved Instruction Exception.
Guest Reserved Instruction Redirect	A guest-mode instruction will trigger a Reserved Instruction Exception. When <i>GuestCtl0_R</i> =1, this root-mode exception is raised before the guest-mode exception can be taken.
Guest Privileged Sensitive Instruction Exception	An instruction executing in guest-mode could not be completed because it was denied access to the required resources by the <i>Root.GuestCtl0</i> register.
Hypercall	A HYPCALL hypercall instruction was executed.
Guest Software Field Change	During guest execution, a software initiated change to certain CP0 register fields occurred. Refer to Section titled “Guest Software Field Change Exception” in the <i>MIPS® Architecture Reference Manual Volume IV-i: Virtualization Module of the MIPS32 Architecture [14]</i> or <i>MIPS® Architecture Reference Manual Volume IV-i: Virtualization Module of the microMIPS32™ Architecture [15]</i> .

Table 7.1 Priority of Exceptions (Continued)

Exception	Description
Guest Hardware Field Change	During guest execution, a hardware initiated set of <i>Status_{EXL/TS}</i> occurred. Refer to Section titled “Guest Hardware Field Change Exception” in the <i>MIPS® Architecture Reference Manual Volume IV-i: Virtualization Module of the MIPS32 Architecture [14]</i> or <i>MIPS® Architecture Reference Manual Volume IV-i: Virtualization Module of the microMIPS32™ Architecture [15]</i> .
Execution Exception	An instruction-based exception occurred: Integer overflow, trap, system call, breakpoint, floating point, or Coprocessor 2 exception.
Tr	Execution of a trap (when trap condition is true).
Protection - Data access	Data access to a protected memory region was attempted.
DDBL / DDBS	EJTAG Data Address Break (address only) or EJTAG Data Value Break on Store (address and value).
AdEL	Load address alignment error. User mode load reference to kernel address.
AdES	Store address alignment error. User mode store to kernel address.
DSRAM Parity Error	Parity error on D-SRAM access.
DBE	Load or store bus error.
DDBL	EJTAG data hardware breakpoint matched in load data compare.
CBrk	EJTAG complex breakpoint.

7.3 Interrupts

In the MIPS32® Release 1 architecture, support for exceptions included two software interrupts, six hardware interrupts, and a special-purpose timer interrupt. The timer interrupt was provided external to the core and was typically combined with hardware interrupt 5 in a system-dependent manner. Interrupts were handled either through the general exception vector (offset 0x180) or the special interrupt vector (0x200), based on the value of *CauseIV*. Software was required to prioritize interrupts as a function of the *CauseIV* bits in the interrupt handler prologue.

Release 2 of the Architecture, implemented by the M5100 core, adds a number of upward-compatible extensions to the Release 1 interrupt architecture, including support for vectored interrupts and the implementation of a new interrupt mode that permits the use of an external interrupt controller.

The M5100 core also includes the Microcontroller Application-Specific Extension (MCU ASE) that provides enhanced interrupt delivery and interrupt-latency reduction.

7.3.1 Interrupt Modes

The M5100 core includes support for three interrupt modes, as defined by Release 2 of the Architecture:

- Interrupt Compatibility mode, in which the behavior of the M5100 is identical to the behavior of a Release 1 implementations.
- Vectored Interrupt (VI) mode, which adds the ability to prioritize and vector interrupts to a handler dedicated to that interrupt, and to assign a GPR shadow set for use during interrupt processing. 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 M5100 processor, so the *VInt* bit will always read as a 1.

Exceptions and Interrupts in the M5100™ Core

- 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 M5100 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 M5100 processor defaults to Compatibility mode, which is fully compatible with all implementations of Release 1 of the Architecture.

Table 7.2 shows the current interrupt mode of the processor as a function of the Coprocessor 0 register fields that can affect the mode.

Table 7.2 Interrupt Modes

Status _{BEV}	Cause _{IV}	IntCtl _{VS}	Config3 _{VINT}	Config3 _{VEIC}	Interrupt Mode
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	Can't happen - IntCtl _{VS} can not be non-zero if neither Vectored Interrupt nor External Interrupt Controller mode is implemented.
"x" denotes don't care					

7.3.1.1 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 16#180 (if *Cause_{IV}* = 0) or vector offset 16#200 (if *Cause_{IV}* = 1). This mode is in effect if any of the following conditions are true:

- Cause_{IV}* = 0
- Status_{BEV}* = 1
- IntCtl_{VS}* = 0, which would be the case if vectored interrupts are not implemented, or have been disabled.

Here is a typical software handler for interrupt compatibility mode:

```

/*
 * Assumptions:
 * - CauseIV = 1 (if it were zero, the interrupt exception would have to
 *   be isolated from the general exception vector before getting
 *   here)
 * - GPRs k0 and k1 are available (no shadow register switches invoked in
 *   compatibility mode)
 * - The software priority is IP9..IP0 (HW7..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, IP9..IP0; k0 = 14..23 */
    xori    k0, k0, 0x17         /* 14..23 => 9..0 */
    sll     k0, k0, VS           /* Shift to emulate software IntCtl_VS */
    la      k1, VectorBase      /* Get base of 10 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 simply 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 simple return to the interrupted code.
*/
    eret                        /* Return to interrupted code */

NestedException:
/*
* Nested exceptions typically require saving the EPC and Status registers,
* 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 can not 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 */

```

Exceptions and Interrupts in the M5100™ Core

```
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.
 * In some environments this may be done with a thread 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, C0_Status      /* Restore the original value */
mtc0    k1, C0_EPC        /* and EPC */
/* Restore GPRs and software state */
eret    /* Dismiss the interrupt */
```

7.3.1.2 Vectored Interrupt (VI) 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. This mode also allows each interrupt to be mapped to a GPR shadow register set for use by the interrupt handler. 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 eight hardware interrupts are interpreted as individual hardware interrupt requests. The timer interrupt is combined in a system-dependent way (external to the core) with the hardware interrupts (the interrupt with which they are combined is indicated by the *PTI* field in *IntCtl*) to provide the appropriate relative priority of the timer interrupt 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,

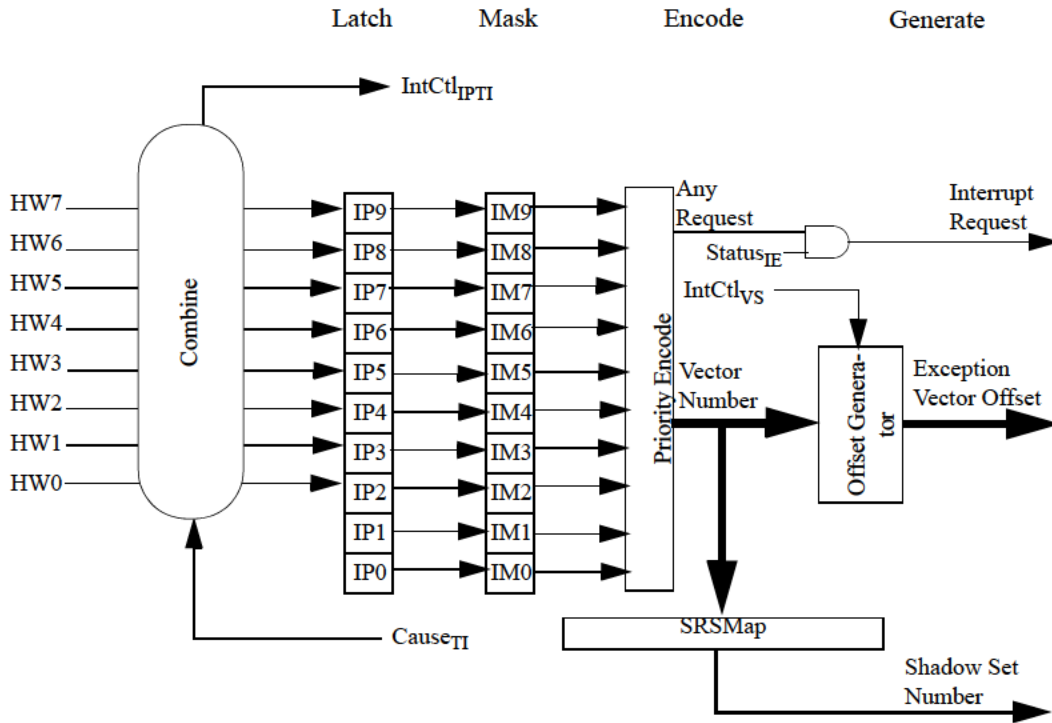
($Status_{EXL} = 0$, and $Status_{ERL} = 0$), an interrupt is signaled and a priority encoder scans the values in the order shown in [Table 7.3](#).

Table 7.3 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	HW7	IP9 and IM9	9
		HW6	IP8 and IM8	8
		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

The priority order places a relative priority on each hardware interrupt and places the software interrupts at a priority lower than all hardware interrupts. When the priority encoder finds the highest priority pending interrupt, it outputs an encoded vector number that is used in the calculation of the handler for that interrupt, as described below. This is shown pictorially in [Figure 7.1](#).

Figure 7.1 Interrupt Generation for Vectored Interrupt Mode



A typical software handler for vectored interrupt 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. Unlike the compatibility mode examples, a vectored interrupt handler may take advantage of a dedicated GPR shadow set to avoid saving any registers. As such, the Simple Interrupt code shown above need not save the GPRs.

A nested interrupt is similar to that shown for compatibility mode, but may also take advantage of running the nested exception routine in the GPR shadow set dedicated to the interrupt or in another shadow set. Such a routine might look as follows:

```

NestedException:
/*
 * Nested exceptions typically require saving the EPC, Status and SRSSctl registers,
 * setting up the appropriate GPR shadow set for the 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 can not cover all nuances of this processing and is intended only
 * to demonstrate the concepts.
 */

/* Use the current GPR shadow set, 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 */
mfc0 k0, CO_SRSSctl     /* Save SRSSctl if changing shadow sets */
sw   k0, SRSSctlSave
    
```

```

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 */
/* If switching shadow sets, write new value to SRSCtlpss here */
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 */

/*
 * If switching shadow sets, clear only KSU above, write target
 * address to EPC, and do execute an eret to clear EXL, switch
 * shadow sets, and jump to routine
 */

/* 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, CO_Status        /* Restore the original value */
lw      k0, SRSCtlSave      /* Get saved SRSCtl */
mtc0   k1, CO_EPC          /* and EPC */
mtc0   k0, CO_SRSCtl       /* Restore shadow sets */
ehb     /* Clear hazard */
eret    /* Dismiss the interrupt */

```

7.3.1.3 External Interrupt Controller Mode

External Internal Interrupt Controller 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, and performance counter interrupts, and directly supplying to the processor the priority level and 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 interrupt mode, the processor sends the state of the software interrupt requests ($Cause_{IP1..IP0}$), the timer interrupt request ($Cause_{TI}$), the performance counter interrupt request ($Cause_{PCI}$) and Fast Debug Channel Interrupt ($Cause_{FDCI}$) to the external interrupt controller, where it prioritizes these interrupts in a system-dependent way with other hardware interrupts. The interrupt controller can be a hard-wired logic block, or it can be configurable based on 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.

Exceptions and Interrupts in the M5100™ Core

The external interrupt controller prioritizes its interrupt requests and produces the priority level and the vector number of the highest priority interrupt to be serviced. The priority level, called the Requested Interrupt Priority Level (RIPL), is an 8-bit encoded value in the range 0..255, inclusive. A value of 0 indicates that no interrupt requests are pending. The values 1..255 represent the lowest (1) to highest (255) RIPL for the interrupt to be serviced. The interrupt controller passes this value on the 8 hardware interrupt lines, which are treated as an encoded value in EIC interrupt mode. There are two implementation options available for the vector offset:

1. The first option is to send a separate vector number along with the RIPL to the processor.
2. A second option is to send an entire vector offset along with the RIPL to the processor. This option is enabled through the core's configuration GUI, and it is not affected by software.

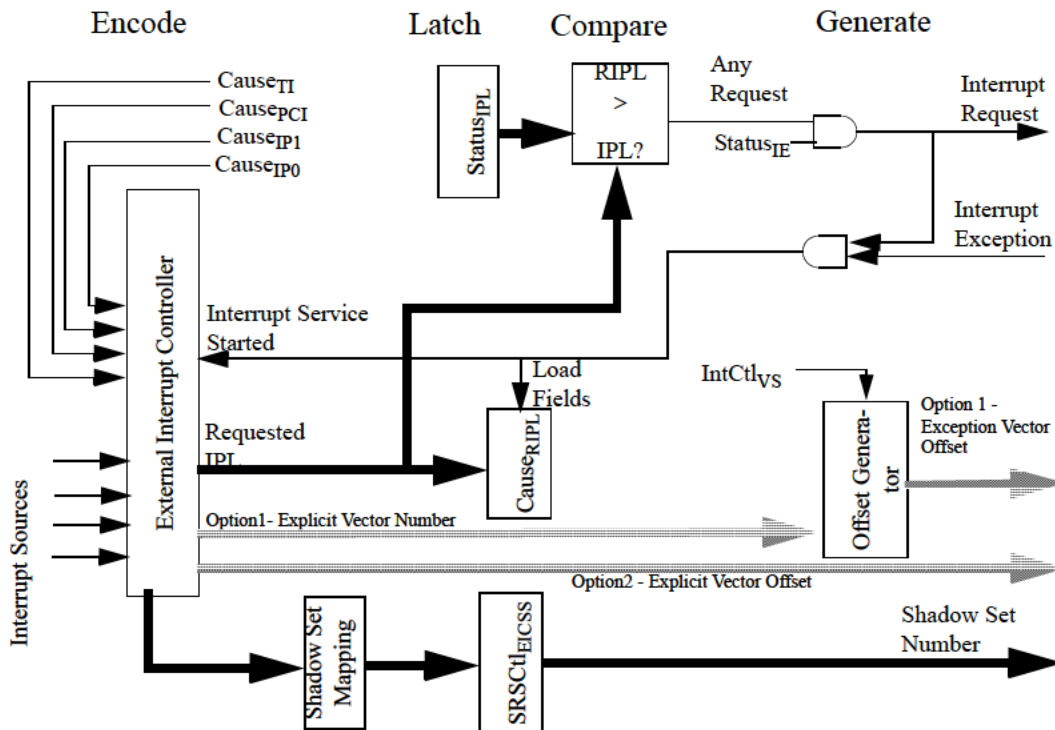
The M5100 core does not support the option to treat the RIPL value as the vector number for the processor.

$Status_{IPL}$ (which overlays $Status_{IM9..IM2}$) is interpreted as the Interrupt Priority Level (IPL) at which the processor is currently operating (with a value of zero indicating 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 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 $Cause_{IP9..IP2}$) and signals the external interrupt controller to notify it that the request is being serviced. Because $Cause_{RIPL}$ is only loaded by the processor when an interrupt exception is signaled, it is available to software during interrupt processing. The vector number that the EIC passes to the core is combined with the $IntCtl_{VS}$ to determine where the interrupt service routine is located. The vector number is not stored in any software-visible registers.

In EIC interrupt mode, the external interrupt controller is also responsible for supplying the GPR shadow set number to use when servicing the interrupt. As such, the $SRSSMap$ register is not used in this mode, and the mapping of the vectored interrupt to a GPR shadow set is done by programming (or designing) the interrupt controller to provide the correct GPR shadow set number when an interrupt is requested. When the processor loads an interrupt request into $Cause_{RIPL}$, it also loads the GPR shadow set number into $SRSCtl_{EICSS}$, which is copied to $SRSCtl_{CSS}$ when the interrupt is serviced.

The operation of EIC interrupt mode is shown pictorially in [Figure 7.2](#).

Figure 7.2 Interrupt Generation for External Interrupt Controller Interrupt Mode



A typical software handler for EIC interrupt 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. Unlike the compatibility mode examples, an EIC interrupt handler may take advantage of a dedicated GPR shadow set to avoid saving any registers. As such, the Simple Interrupt code shown above need not save the GPRs.

A nested interrupt is similar to that shown for compatibility mode, but may also take advantage of running the nested exception routine in the GPR shadow set dedicated to the interrupt or in another shadow set. It also need only copy $Cause_{RIPL}$ to $Status_{IPL}$ to prevent lower priority interrupts from interrupting the handler. Such a routine might look as follows:

```

NestedException:
/*
 * Nested exceptions typically require saving the EPC, Status, and SRSCtl registers,
 * setting up the appropriate GPR shadow set for the 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 can not cover all nuances of this processing and is intended only
 * to demonstrate the concepts.
 */

/* Use the current GPR shadow set, and setup software context */
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 */

```

Exceptions and Interrupts in the M5100™ Core

```

sw      k0, StatusSave      /* Save in memory */
ins     k0, k1, S_StatusIPL, 6 /* Set IPL to RIPL in copy of Status */
mfc0   k1, C0_SRSCtl       /* Save SRSCtl if changing shadow sets */
sw      k1, SRSCtlSave
/* If switching shadow sets, write new value to SRSCtlpss here */
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 */

/*
 * If switching shadow sets, clear only KSU above, write target
 * address to EPC, and do execute an eret to clear EXL, switch
 * shadow sets, and jump to routine
 */

/* Process interrupt here, including clearing device interrupt */

/*
 * The interrupt completion code is identical to that shown for VI mode above.
 */

```

7.3.2 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 16#200 to create the exception vector offset. For VI interrupt mode, the vector number is in the range 0..9, inclusive. For EIC interrupt mode, the vector number is in the range 0..63, inclusive. 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, and [Table 7.4](#) shows the exception vector offset for a representative subset of the vector numbers and values of the *IntCtl_{VS}* field.

Table 7.4 Exception Vector Offsets for Vectored Interrupts

Vector Number	Value of IntCtl _{VS} Field				
	2#00001	2#00010	2#00100	2#01000	2#10000
0	16#0200	16#0200	16#0200	16#0200	16#0200
1	16#0220	16#0240	16#0280	16#0300	16#0400
2	16#0240	16#0280	16#0300	16#0400	16#0600
3	16#0260	16#02C0	16#0380	16#0500	16#0800
4	16#0280	16#0300	16#0400	16#0600	16#0A00
5	16#02A0	16#0340	16#0480	16#0700	16#0C00
6	16#02C0	16#0380	16#0500	16#0800	16#0E00
7	16#02E0	16#03C0	16#0580	16#0900	16#1000
		• • •			
61	16#09A0	16#1140	16#2080	16#3F00	16#7C00
62	16#09C0	16#1180	16#2100	16#4000	16#7E00
63	16#09E0	16#11C0	16#2180	16#4100	16#8000

The general equation for the exception vector offset for a vectored interrupt is:

$$\text{vectorOffset} \leftarrow 16\#200 + (\text{vectorNumber} \times (\text{IntCtlVs} \parallel 2\#00000))$$

When using large vector spacing and EIC mode, the offset value can overlap with bits that are specified in the EBase register. Software must ensure that any overlapping bits are specified as 0 in EBase. This implementation ORs together the offset and base registers, but it is architecturally undefined and software should not rely on this behavior.

Although there are 255 EIC priority interrupts, only 64 vectors are provided. There is no one-to-one mapping for each EIC interrupt to its interrupt vector. The 255 priority interrupts will share the 64 interrupt vectors as specified by the *SL_EICVector*[5:0] input pins. However, as mentioned in option 2 of [Section 7.3.1.3 “External Interrupt Controller Mode”](#), the *SL_Offse*[17:1] input pins can be used to provide each EIC interrupt with a unique interrupt handler location.

7.3.3 MCU ASE Enhancement for Interrupt Handling

The MCU ASE extends the MIPS/microMIPS Architecture with a set of new features designed for the microcontroller market. The MCU ASE contains enhancements in two key areas: interrupt delivery and interrupt latency. For more details, refer to the *The MCU Privileged Resource Architecture* chapter of the *MIPS® Architecture for Programmers Volume IV-h: The MCU Application-Specific Extension to the MIPS32 Architecture* [10] or *MIPS® Architecture for Programmers Volume IV-h: The MCU Application-Specific Extension to the microMIPS32™ Architecture* [11].

7.3.3.1 Interrupt Delivery

The MCU ASE extends the number of hardware interrupt sources from 6 to 8. For legacy and vectored-interrupt mode, this represents 8 external interrupt sources. For EIC mode, the widened *IPL* and *RIPL* fields can now represent 256 external interrupt sources.

7.3.3.2 Interrupt Latency Reduction

The MCU ASE includes a package of extensions to MIPS/microMIPS that decrease the latency of the processor’s response to a signalled interrupt.

Interrupt Vector Prefetching

Normally on MIPS architecture processors, when an interrupt or exception is signalled, execution pipelines must be flushed before the interrupt/exception handler is fetched. This is necessary to avoid mixing the contexts of the interrupted/faulting program and the exception handler. The MCU ASE introduces a hardware mechanism in which the interrupt exception vector is prefetched whenever the interrupt input signals change. The prefetch memory transaction occurs in parallel with the pipeline flush and exception prioritization. This decreases the overall latency of the execution of the interrupt handler’s first instruction.

Automated Interrupt Prologue

The use of Shadow Register Sets avoids the software steps of having to save general-purpose registers before handling an interrupt.

The MCU ASE adds additional hardware logic that automatically saves some of the COP0 state in the stack and automatically updates some of the COP0 registers in preparation for interrupt handling.

Automated Interrupt Epilogue

A mirror to the Automated Prologue, this feature automates the restoration of some of the COP0 registers from the stack and the preparation of some of the COP0 registers for returning to non-exception mode. This feature is implemented within the IRET instruction, which is introduced in this ASE.

Interrupt Chaining

An optional feature of the Automated Interrupt Epilogue, this feature allows handling a second interrupt after a primary interrupt is handled, without returning to non-exception mode (and the related pipeline flushes that would normally be necessary).

7.4 GPR Shadow Registers

Release 2 of the Architecture optionally removes the need to save and restore GPRs on entry to high priority interrupts or exceptions, and to provide specified processor modes with the same capability. This is done by introducing multiple copies of the GPRs, called *shadow sets*, and allowing privileged software to associate a shadow set with entry to kernel mode via an interrupt vector or exception. The normal GPRs are logically considered shadow set zero.

The number of GPR shadow sets is a build-time option on the M5100 core. Although Release 2 of the Architecture defines a maximum of 16 shadow sets, the core allows one (the normal GPRs), two, four, eight or sixteen shadow sets. The highest number actually implemented is indicated by the *SRSCtl_{HSS}* field. If this field is zero, only the normal GPRs are implemented.

Shadow sets are new copies of the GPRs that can be substituted for the normal GPRs on entry to kernel mode via an interrupt or exception. When a shadow set is bound to a kernel mode entry condition, reference to GPRs work exactly as one would expect, but they are redirected to registers that are dedicated to that condition. Privileged software may need to reference all GPRs in the register file, even specific shadow registers that are not visible in the current mode. The RDPGPR and WRPGPR instructions are used for this purpose. The *CSS* field of the *SRSCtl* register provides the number of the current shadow register set, and the *PSS* field of the *SRSCtl* register provides the number of the previous shadow register set (that which was current before the last exception or interrupt occurred).

If the processor is operating in VI interrupt mode, binding of a vectored interrupt to a shadow set is done by writing to the *SRSMap* register. If the processor is operating in EIC interrupt mode, the binding of the interrupt to a specific shadow set is provided by the external interrupt controller, and is configured in an implementation-dependent way. Binding of an exception or non-vectored interrupt to a shadow set is done by writing to the *ESS* field of the *SRSCtl* register. When an exception or interrupt occurs, the value of *SRSCtl_{CSS}* is copied to *SRSCtl_{PSS}*, and *SRSCtl_{CSS}* is set to the value taken from the appropriate source. On an ERET, the value of *SRSCtl_{PSS}* is copied back into *SRSCtl_{CSS}* to restore the shadow set of the mode to which control returns. More precisely, the rules for updating the fields in the *SRSCtl* register on an interrupt or exception are as follows:

1. No field in the *SRSCtl* register is updated if any of the following conditions is true. In this case, steps 2 and 3 are skipped.
 - The exception is one that sets *Status_{ERL}*: Reset, Soft Reset, or NMI.
 - The exception causes entry into EJTAG Debug Mode.
 - *Status_{BEV}* = 1
 - *Status_{EXL}* = 1

2. $SRSCtl_{CSS}$ is copied to $SRSCtl_{PSS}$.
3. $SRSCtl_{CSS}$ is updated from one of the following sources:
 - The appropriate field of the $SRSSMap$ register, based on IPL, if the exception is an interrupt, $Cause_{IV} = 1$, $Config3_{VEIC} = 0$, and $Config3_{VInt} = 1$. These are the conditions for a vectored interrupt.
 - The $EICSS$ field of the $SRSCtl$ register if the exception is an interrupt, $Cause_{IV} = 1$, and $Config3_{VEIC} = 1$. These are the conditions for a vectored EIC interrupt.
 - The ESS field of the $SRSCtl$ register in any other case. This is the condition for a non-interrupt exception, or a non-vectored interrupt.

Similarly, the rules for updating the fields in the $SRSCtl$ register at the end of an exception or interrupt are as follows:

1. No field in the $SRSCtl$ register is updated if any of the following conditions is true. In this case, step 2 is skipped.
 - A DERET is executed.
 - An ERET is executed with $Status_{ERL} = 1$.
2. $SRSCtl_{PSS}$ is copied to $SRSCtl_{CSS}$.

These rules have the effect of preserving the $SRSCtl$ register in any case of a nested exception or one which occurs before the processor has been fully initialize ($Status_{BEV} = 1$).

Privileged software may switch the current shadow set by writing a new value into $SRSCtl_{PSS}$, loading EPC with a target address, and doing an ERET.

7.5 Exception Vector Locations

The Reset, Soft Reset, and NMI exceptions are always vectored to location 16#BFC0.0000. The vector location for EJTAG Debug exceptions depends on the value of the ProbTrap bit in the $EJTAG_Control_register$. Debug exceptions are vectored to location 16#BFC0.0480 if the ProbTrap bit is zero, or to location 16#FF20.0200 if the ProbTrap bit is one. Vector addresses for all other exceptions are a combination of a vector offset and a vector base address. In Release 1 of the architecture, the vector base address was fixed. In Release 2 of the architecture, software is allowed to specify the vector base address via the $EBase$ register for exceptions that occur when $Status_{BEV}$ equals 0. [Table 7.5](#) gives the vector base address as a function of the exception and whether the BEV bit is set in the $Status$ register. [Table 7.6](#) gives 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. For implementations of Release 2 of the Architecture,

[Table 7.4](#) shows the offset from the base address in the case where $Status_{BEV} = 0$ and $Cause_{IV} = 1$. For implementations of Release 1 of the architecture in which $Cause_{IV} = 1$, the vector offset is as if $IntCtl_{IVS}$ were 0. [Table 7.7](#) combines these two 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, the vector address value assumes that the *EBase* register, as implemented in Release 2 devices, is not changed from its reset state and that *IntCtl_{VS}* is 0.

Table 7.5 Exception Vector Base Addresses

Exception	Status _{BEV}	
	0	1
Reset, Soft Reset, NMI	16#BFC0.0000	
EJTAG Debug (with <i>ProbEn</i> = 0 in the <i>EJTAG Control Register</i>)	16#BFC0.0480	
EJTAG Debug (with <i>ProbEn</i> = 1 in the <i>EJTAG Control Register</i>)	16#FF20.0200	
SRAM Parity Error	EBase _{31:30} 1 EBase _{28:12} 16#000 Note that EBase _{31:30} has the fixed value 2#10	16#BFC0.0300
Other	<i>For Release 1 of the architecture:</i> 16#8000.0000 <i>For Release 2 of the architecture:</i> EBase _{31:12} 16#000 Note that EBase _{31:30} has the fixed value 2#10	16#BFC0.0200

Table 7.6 Exception Vector Offsets

Exception	Vector Offset
General Exception	16#180
Interrupt, <i>Cause_{IV}</i> = 1	16#200 (In Release 2 implementations, this is the base of the vectored interrupt table when <i>Status_{BEV}</i> = 0)
Reset, Soft Reset, NMI	None (Uses Reset Base Address)

Table 7.7 Exception Vectors

Exception	Status _{BEV}	Status _{EXL}	Cause _{IV}	EJTAG ProbEn	Vector
					For Release 2 Implementations, assumes that EBase retains its reset state and that IntCtl _{VS} = 0
Reset, Soft Reset, NMI	x	x	x	x	16#BFC0.0000
EJTAG Debug	x	x	x	0	16#BFC0.0480
EJTAG Debug	x	x	x	1	16#FF20.0200
SRAM Parity Error	0	x	x	x	16#EBase[31:30] 2#1 EBase[28:12] 16#100
SRAM Parity Error	1	x	x	x	16#BFC0.0300
Interrupt	0	0	0	x	16#8000.0180

Table 7.7 Exception Vectors (Continued)

Exception	Status _{BEV}	Status _{EXL}	Cause _{IV}	EJTAG ProbEn	Vector For Release 2 Implementations, assumes that EBase retains its reset state and that IntCtl _{VS} = 0
Interrupt	0	0	1	x	16#8000.0200
Interrupt	1	0	0	x	16#BFC0.0380
Interrupt	1	0	1	x	16#BFC0.0400
All others	0	x	x	x	16#8000.0180
All others	1	x	x	x	16#BFC0.0380
‘x’ denotes don’t care					

7.6 General Exception Processing

With the exception of Reset, Soft 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 (see Table 8.34). The value loaded into the *EPC* register is dependent on whether the processor implements microMIPS, and whether the instruction is in the delay slot of a branch or jump which has delay slots. Table 7.8 shows the value stored in each of the CP0 PC registers, including *EPC*. For implementations of Release 2 of the Architecture if Status_{BEV} = 0, the *CSS* field in the *SRSCtl* register is copied to the *PSS* field, and the *CSS* value is loaded from the appropriate source.

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. For implementations of Release 2 of the Architecture, the *SRSCtl* register is not changed.

Table 7.8 Value Stored in EPC, ErrorEPC, or DEPC on an Exception

microMIPS 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 bits of the address of the instruction, combined with the <i>ISA Mode</i> bit
Yes	Yes	Upper bits of the branch or jump instruction (PC-2 or PC-4 depending on size of the instruction in the microMIPS <i>ISA Mode</i> and PC-4 in the 32-bit <i>ISA Mode</i>), combined with the <i>ISA Mode</i> 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 is started at the exception vector.

Exceptions and Interrupts in the M5100™ Core

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 EPC nor Cause_BD nor SRSCtl are modified */
if Status_EXL = 1 then
    vectorOffset ← 16#180
else
    if InstructionInBranchDelaySlot then
        EPC ← restartPC /* PC of branch/jump */
        Cause_BD ← 1
    else
        EPC ← restartPC /* PC of instruction */
        Cause_BD ← 0
    endif

    /* Compute vector offsets as a function of the type of exception */
    NewShadowSet ← SRSCtl_ESS /* Assume exception, Release 2 only */
    if ExceptionType = TLBRefill then
        vectorOffset ← 16#000
    elseif (ExceptionType = Interrupt) then
        if (Cause_IV = 0) then
            vectorOffset ← 16#180
        else
            if (Status_BEV = 1) or (IntCtl_VS = 0) then
                vectorOffset ← 16#200
            else
                if Config3_VEIC = 1 then
                    VecNum ← Cause_RIPL
                    NewShadowSet ← SRSCtl_EICSS
                else
                    VecNum ← VIntPriorityEncoder()
                    NewShadowSet ← SRSMaIPL×4+3 IPL×4
                endif
                vectorOffset ← 16#200 + (VecNum × (IntCtl_VS || 2#00000))
            endif /* if (Status_BEV = 1) or (IntCtl_VS = 0) then */
        endif /* if (Cause_IV = 0) then */
    endif /* elseif (ExceptionType = Interrupt) then */

    /* Update the shadow set information for an implementation of */
    /* Release 2 of the architecture */
    if ((ArchitectureRevision ≥ 2) and (SRSCtl_HSS > 0) and (Status_BEV = 0) and
        (Status_ERL = 0)) then
        SRSCtl_PSS ← SRSCtl_CSS
        SRSCtl_CSS ← NewShadowSet
    endif
endif /* if Status_EXL = 1 then */

Cause_CE ← FaultingCoprocesorNumber
Cause_ExcCode ← ExceptionType
Status_EXL ← 1
```

```

/* Calculate the vector base address */
if StatusBEV = 1 then
    vectorBase ← 16#BFC0.0200
else
    if ArchitectureRevision ≥ 2 then
        /* The fixed value of EBase31:30 forces the base to be in kseg0 or kseg1 */
        vectorBase ← EBase31:12 || 16#000
    else
        vectorBase ← 16#8000.0000
    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 */

```

7.7 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*, *DINT*, *DIBImpr*, *DDBLImpr*, and *DDBSImpr* bits in the *Debug* register are updated appropriately depending on the debug exception type.
- The *Debug2* register is updated with additional information for complex breakpoints.
- *Halt* and *Doze* bits in the *Debug* register are updated appropriately.
- *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*, *DINT*, *DIBImpr*, *DDBLImpr*, and *DDBSImpr* bits in the *Debug* register.

No other CP0 registers or fields are changed due to the debug exception, thus no additional state is saved.

Operation:

```

if InstructionInBranchDelaySlot then
    DEPC ← PC-4
    DebugDBD ← 1
else
    DEPC ← PC
    DebugDBD ← 0
endif

```

Exceptions and Interrupts in the M5100™ Core

```
DebugD* bits ← DebugExceptionType
DebugHalt ← HaltStatusAtDebugException
DebugDoze ← DozeStatusAtDebugException
DebugDM ← 1
if EJTAGControlRegisterProbTrap = 1 then
    PC ← 0xFF20_0200
else
    PC ← 0xBFC0_0480
endif
```

The same debug exception vector location is used for all debug exceptions. The location is determined by the Prob-Trap bit in the EJTAG Control register (ECR), as shown in [Table 7.9](#).

Table 7.9 Debug Exception Vector Addresses

ProbTrap bit in ECR Register	Debug Exception Vector Address
0	0xBFC0_0480
1	0xFF20_0200 in dmseg

7.8 Exception Descriptions

The following subsections describe each of the exceptions listed in the same sequence as shown in [Table 7.1](#).

7.8.1 Reset/SoftReset Exception

A reset exception occurs when the *SI_ColdReset* signal is asserted to the processor; a soft reset occurs when the *SI_Reset* signal is asserted. These exceptions are not maskable. When one of these exceptions 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/Soft-Reset exception, the state of the processor is not defined, with the following exceptions:

- 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 *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 (0xBFC0_0000)

Operation:

```

Config ← ConfigurationState
StatusRP ← 0
StatusBEV ← 1
StatusTS ← 0
StatusSR ← 0/1 (depending on Reset or SoftReset)
StatusNMI ← 0
StatusERL ← 1
if InstructionInBranchDelaySlot then
    ErrorEPC ← PC - 4
else
    ErrorEPC ← PC
endif
PC ← 0xBFC0_0000

```

7.8.2 Debug Single Step Exception

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 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* points to the SDBBP instruction. However, returning to an instruction (not jump/branch) just before the SDBBP instruction, causes a debug single step exception with the *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

7.8.3 Debug Interrupt Exception

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

7.8.4 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 (0xBFC0_0000)

Operation:

```

StatusBEV ← 1
StatusTS ← 0
StatusSR ← 0
StatusNMI ← 1
StatusERL ← 1
if InstructionInBranchDelaySlot then
    ErrorEPC ← PC - 4
else

```



```

    ErrorEPC ← PC
endif
PC ← 0xBFC0_0000

```

7.8.5 Interrupt Exception

The interrupt exception occurs when one or more of the eight hardware, two software, or timer interrupt requests is enabled by the *Status* register, and the interrupt input is asserted. See 7.3 “Interrupts” on page 121 for more details about the processing of interrupts.

Register ExcCode Value:

Int

Additional State Saved:

Table 7.10 Register States an Interrupt Exception

Register State	Value
<i>CauseIP</i>	indicates the interrupts that are pending.

Entry Vector Used:

See 7.3.2 “Generation of Exception Vector Offsets for Vectored Interrupts” on page 130 for the entry vector used, depending on the interrupt mode the processor is operating in.

7.8.6 Debug Instruction Break Exception

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

7.8.7 Address Error Exception — Instruction Fetch/Data Access

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

Exceptions and Interrupts in the M5100™ Core

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 7.11 CP0 Register States on an Address Exception Error

Register State	Value
BadVAddr	Failing address

Entry Vector Used:

General exception vector (offset 0x180)

7.8.8 SRAM Parity Error Exception

A SRAM error exception occurs when an instruction or data reference detects a data error. This exception is not maskable. To avoid disturbing the error in the cache array the exception vector is to an unmapped, uncached address. This exception is precise.

Cause Register ExcCode Value

N/A

Additional State Saved

Table 7.12 CP0 Register States on a SRAM Parity Error Exception

Register State	Value
<i>CacheErr</i>	Error state
<i>ErrorEPC</i>	Restart PC

Entry Vector Used

Cache error vector (offset 16#100)

7.8.9 Bus Error Exception — Instruction Fetch or Data Access

A bus error exception occurs when an instruction or data access makes a bus request and that request terminates in an error. The bus error exception can occur on either an instruction fetch or a data access. Bus error exceptions that occur on an instruction fetch have a higher priority than bus error exceptions that occur on a data access.

Bus errors taken on any external access on the M5100 core are always precise.

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)

7.8.10 Protection Exception

The protection exception occurs when an access to memory that has been protected by the Memory Protection Unit has been attempted. Or under certain circumstances, attempted write to the EBase register. See the "Security Features of the M14K™ Processor Family" (MD00896) for more information.

Register ExcCode Value:

Prot (Cause Code 29)

Additional State Saved:

MPU Config Register, Triggered Field

MPU StatusN Register, Cause* Fields

Entry Vector Used

General exception vector (offset 0x180)

7.8.11 Debug Software Breakpoint Exception

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

7.8.12 Execution Exception — System Call

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)

7.8.13 Execution Exception — Breakpoint

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)

7.8.14 Execution Exception — Reserved Instruction

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)

7.8.15 Execution Exception — Coprocessor Unusable

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 7.13 Register States on a Coprocessor Unusable Exception

Register State	Value
Cause _{CE}	Unit number of the coprocessor being referenced

Entry Vector Used:

General exception vector (offset 0x180)

7.8.16 Execution Exception — CorExtend Unusable

The CorExtend unusable exception is one of the execution exceptions. All of these exceptions have the same priority. A CorExtend Unusable exception occurs when an attempt is made to execute a CorExtend instruction when *Status_{CEE}* is cleared. It is implementation-dependent whether this functionality is supported. Generally, the functionality will only be supported if a CorExtend block contains local destination registers

Cause Register ExcCode Value:

CEU

Additional State Saved:

None

Entry Vector Used:

General exception vector (offset 0x180)

7.8.17 Execution Exception — DSP Module State Disabled

The DSP Module State Disabled exception is an execution exception. It occurs when an attempt is made to execute a DSP Module 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)

7.8.18 Execution Exception — Coprocessor 2 Exception

The Coprocessor 2 exception is one of the execution exceptions. All of these exceptions have the same priority. A Coprocessor 2 exception occurs when a valid Coprocessor 2 instruction cause a general exception in the Coprocessor 2.

Cause Register ExcCode Value:

C2E

Additional State Saved:

Depending on the Coprocessor 2 implementation, additional state information of the exception can be saved in a Coprocessor 2 control register.

Entry Vector Used:

General exception vector (offset 0x180)

7.8.19 Execution Exception — Implementation-Specific 1 Exception

The Implementation-Specific 1 exception is one of the execution exceptions. All of these exceptions have the same priority. An implementation-specific 1 exception occurs when a valid coprocessor 2 instruction cause an implementation-specific 1 exception in the Coprocessor 2.

Cause Register ExcCode Value:

IS1

Additional State Saved:

Depending on the coprocessor 2 implementation, additional state information of the exception can be saved in a coprocessor 2 control register.

Entry Vector Used:

General exception vector (offset 0x180)

7.8.20 Execution Exception — Integer Overflow

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)

7.8.21 Execution Exception — Trap

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)

7.8.22 Debug Data Break Exception

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

7.8.23 Complex Break Exception

A complex data break exception occurs when the complex hardware breakpoint detects an enabled breakpoint. Complex breaks are taken imprecisely—the instruction that actually caused the exception is allowed to complete and the *DEPC* register and DBD bit in the *Debug* register point to a following instruction.

Debug Register Debug Status Bit Set:

DIBImpr, *DDBLImpr*, and/or *DDBSImpr*

Additional State Saved:

Debug2 fields indicate which type(s) of complex breakpoints were detected.

Entry Vector Used:

Debug exception vector

7.9 Exception Handling and Servicing Flowcharts

The remainder of this chapter contains flowcharts for the following exceptions and guidelines for their handlers:

- General exceptions and their exception handler
- Reset, soft reset and NMI exceptions, and a guideline to their handler
- Debug exceptions

Figure 7.3 General Exception Handler (HW)

Exceptions other than Reset, Soft Reset, NMI, EJTag Debug and cache error, or first-level TLB miss.
 Note: Interrupts can be masked by IE or IMs and Watch is masked if EXL = 1

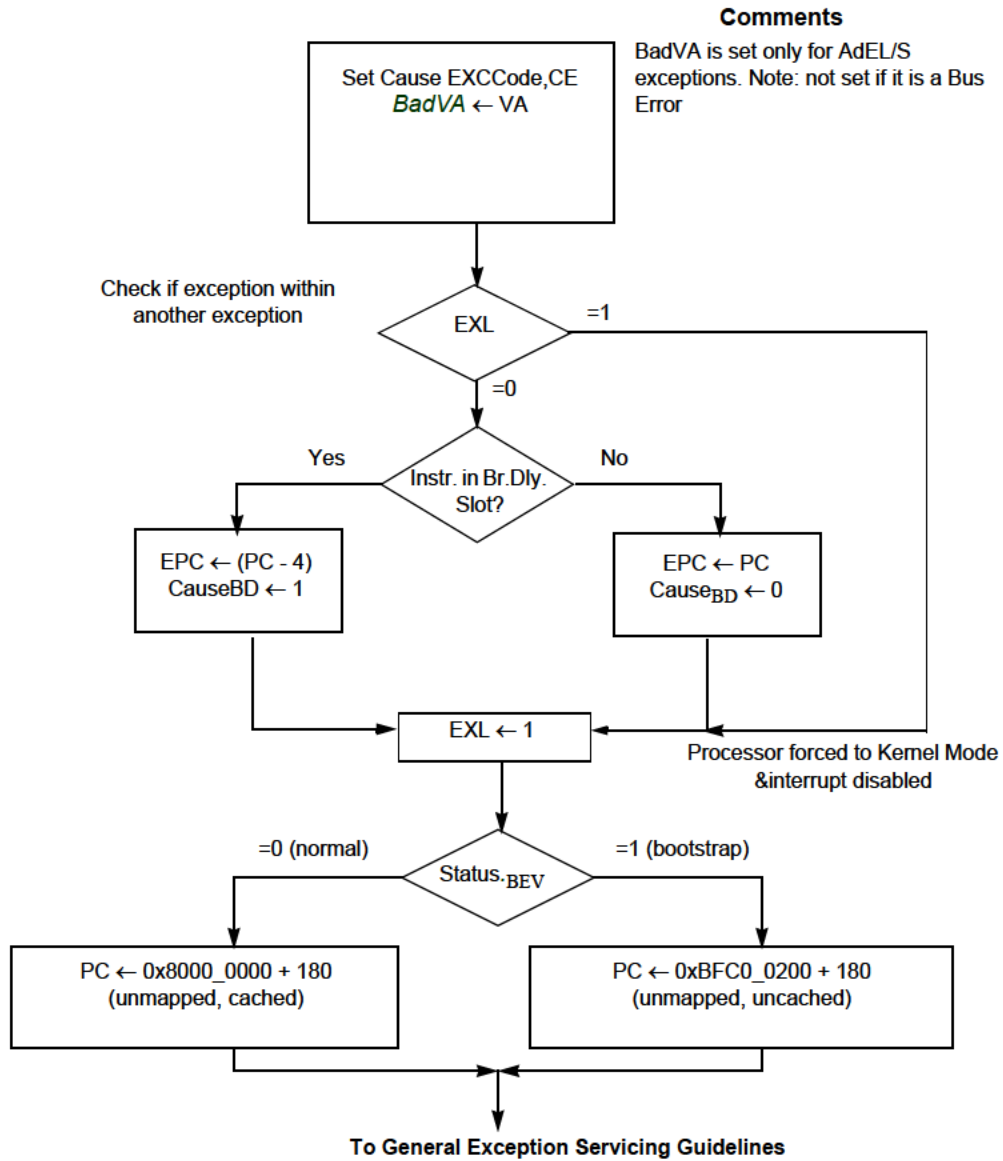


Figure 7.4 General Exception Servicing Guidelines (SW)

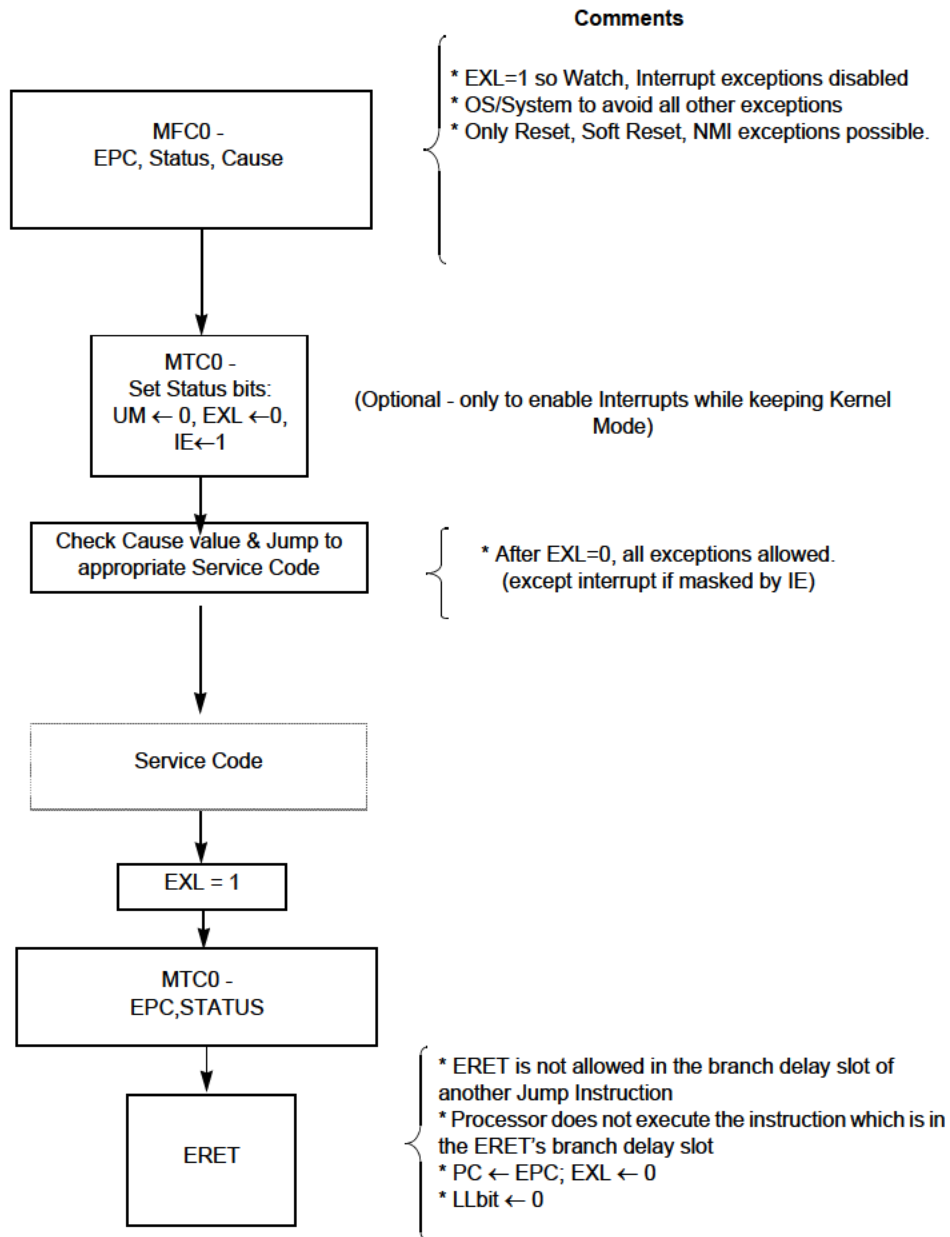
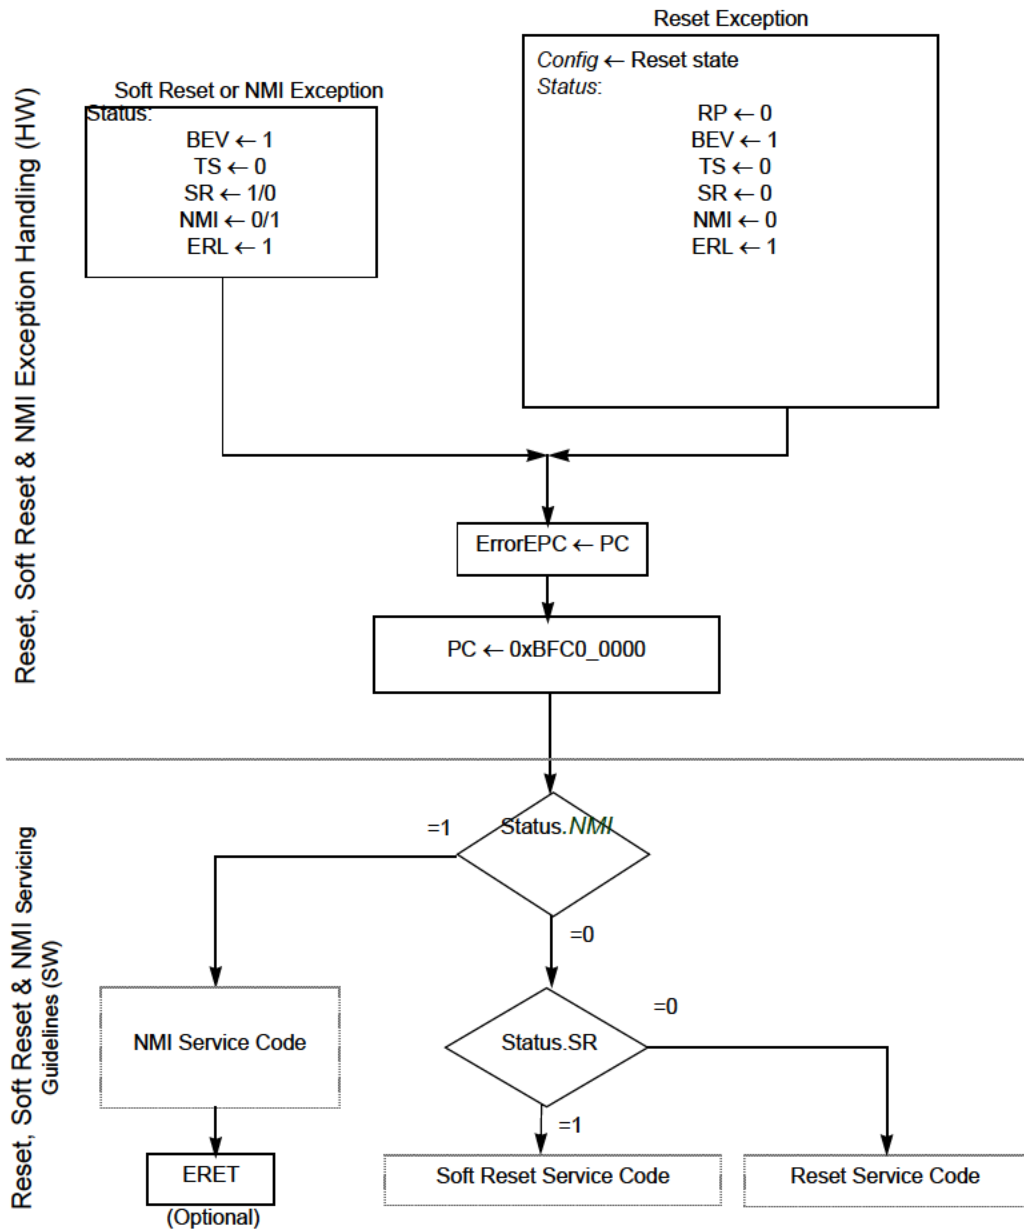


Figure 7.5 Reset, Soft Reset and NMI Exception Handling and Servicing Guidelines



CP0 Registers of the M5100™ Core

The System Control Coprocessor (CP0) provides the register interface to the M5100 processor core for the support of memory management, address translation, exception handling, anti-tamper features, and other privileged operations. Each CP0 register is identified by a *Register Number*, from 0 through 31, and a *Select Number* that is used as the value in the *sel* field of the MFC0 and MTC0 instructions. For instance, the *EBase* register is Register Number 15, Select 1.

After updating a CP0 register, there is a hazard period of zero or more instructions from the update by the MTC0 instruction until the update has taken effect in the core. For a detailed description of CP0 hazards, refer to [Section 2.12 “Hazards”](#).

This chapter contains the following sections:

- [Section 8.1 “CP0 Register Summary”](#)
- [Section 8.2 “CP0 Register Descriptions”](#)

The EJTAG registers are described in [Chapter 11, “EJTAG Debug Support in the M5100™ Core”](#) on page 248.

8.1 CP0 Register Summary

[Table 8.1](#) lists the CP0 registers in numerical order. Individual registers are described in [Section 8.2 “CP0 Register Descriptions”](#).

Table 8.1 CP0 Registers

Register Number	Select Number	Register Name	Function
0	0	<i>Index</i> ³	Index into the TLB array
1	0	<i>Random</i> ³	Randomly generated index into the TLB array
2	0	<i>EntryLo0</i> ³	Low-order portion of the TLB entry for even-numbered virtual pages
3	0	<i>EntryLo1</i> ³	Low-order portion of the TLB entry for odd-numbered virtual pages
4	0 2	<i>Context/</i> <i>UserLocal</i>	Pointer to page table entry in memory User information that can be written by privileged software and read via <i>RDHWR</i> register 29
5	0 1	<i>PageMask</i> <i>PageGrain</i> ³	<i>PageMask</i> controls the variable page sizes in TLB entries. <i>PageGrain</i> enables support of 1KB pages in the TLB.
6	0	<i>Wired3</i>	Controls the number of fixed (“wired”) TLB entries

Table 8.1 CP0 Registers (Continued)

Register Number	Select Number	Register Name	Function
7	0	<i>HWREna</i>	Enables access via the RDHWR instruction to selected hardware registers in non-privileged mode
8	0 1 2	<i>BadVAddr</i> ¹ <i>BadInstr</i> <i>BadInstrP</i>	Reports the address for the most recent address-related exception Reports the instruction that caused the most recent exception Reports the branch instruction if a delay slot caused the most recent exception
9	0	<i>Count</i> ¹	Processor cycle count
10	0 4 5 6	<i>EntryHi</i> ³ <i>GuestCtl1</i> <i>GuestCtl2</i> <i>GuestCtl3</i>	High-order portion of the TLB entry GuestID of virtualization Guest Guest Interrupt Control Guest Shadow Register Set Control
11	0 4	<i>Compare</i> ¹ <i>GuestCtl0Ext</i>	Timer interrupt control Extension to GuestCtl0
12	0 1 2 3 4 5 6 7	<i>Status</i> <i>IntCtl</i> <i>SRSCtl</i> <i>SRSMap1</i> <i>View_IPL</i> <i>SRSMAP2</i> <i>GuestCtl0</i> <i>GTOffset</i>	Processor status and control Interrupt system status and control Shadow Register Sets status and control Shadow set IPL mapping Contiguous view of IM and IPL fields Shadow set IPL mapping Controls Virtualized Guest OS Guest Timer Offset
13	0 4 5	<i>Cause</i> ¹ <i>View_RIPL</i> <i>NestedExc</i>	Cause of last exception
14	0 2	<i>EPC</i> ¹ <i>NestedEPC</i>	Program counter at last exception
15	0 1 2	<i>PRId</i> / <i>EBase</i> <i>CDMMBase</i>	Processor identification and revision; exception base address; Common Device Memory Map Base register
16	0 1 2 3 4 5 7	<i>Config</i> <i>Config1</i> <i>Config2</i> <i>Config3</i> <i>Config4</i> <i>Config5</i> <i>Config7</i>	Configuration registers
17	0	<i>LLAddr</i>	Load linked address
18	0-7	<i>WatchLo</i> ¹	Low-order watchpoint address
19	0-7	<i>WatchHi</i> ¹	High-order watchpoint address
20 - 22		Reserved	Reserved

Table 8.1 CP0 Registers (Continued)

Register Number	Select Number	Register Name	Function
23	0	<i>Debug</i>	EJTAG Debug register
	1	<i>TraceControl</i>	EJTAG Trace Control register
	2	<i>TraceControl2</i>	EJTAG Trace Control register2
	3	<i>UserTraceData1</i>	EJTAG User Trace Data1 register
	4	<i>TraceBPC²</i>	EJTAG Trace Breakpoint Register
	6	<i>Debug2</i>	EJTAG Debug register 2
24	0	<i>DEPC²</i>	Program counter at last debug exception
	3	<i>UserTraceData2</i>	EJTAG User Trace Data2 register
25	0	<i>PerfCtl0</i>	Performance counter 0 control
	1	<i>PerfCnt0</i>	Performance counter 0
	2	<i>PerfCtl1</i>	Performance counter 1 control
	3	<i>PerfCnt1</i>	Performance counter 1
26	0	<i>ErrCtl</i>	Software test enable of way-select and Data RAM arrays for I-Cache and D-Cache
27	0	<i>CacheErr</i>	Records information about cache/SPRAM parity errors
28	0	<i>TagLo</i>	Low-order portion of cache tag interface
	1	<i>DataLo</i>	
29		Reserved	Reserved
30	0	<i>ErrorEPC¹</i>	Program counter at last error
31	0	<i>DeSAVE²</i>	Debug handler scratchpad register
	2	<i>KScratch1</i>	Scratch Register for Kernel Mode
	3	<i>Kscratch2</i>	Scratch Register for Kernel Mode
1. Registers used in exception processing 2. Registers used in debug 3. Registers used in memory management			

8.2 CP0 Register Descriptions

This section contains descriptions of each CP0 register. The registers are listed in numerical order, first by Register Number, then by Select Number.

For each register described below, field descriptions include the read/write properties of the field (shown in [Table 8.2](#)) and the reset state of the field. .

Table 8.2 CP0 Register R/W Field Types

Read/Write 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 UNDEFINED behavior.	

8.2.2 Random Register (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. The width of the Random field is calculated in the same manner as that described for the *Index* register above.

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 that prevents the decrement pseudo-randomly.

The processor initializes the *Random* register to the upper bound on a Reset exception and when the *Wired* register is written.

Figure 8.2 Random Register Format




Table 8.4 Random Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
0	31:5	Must be written as zero; returns zero on reads.	0	0
Random	4:0	TLB Random Index	R	TLB Entries - 1

8.2.3 EntryLo0 and EntryLo1 Registers (CP0 Registers 2 and 3, Select 0)

The pair of *EntryLo* registers act as the interface between the TLB and the TLBR, TLBWI, and TLBWR instructions. For a TLB-based MMU, *EntryLo0* holds the entries for even pages and *EntryLo1* holds the entries for odd pages.

The contents of the *EntryLo0* and *EntryLo1* registers are undefined after an address error, TLB invalid, TLB modified, or TLB refill exception.

Figure 8.3 EntryLo0, EntryLo1 Register Format




Table 8.5 EntryLo0, EntryLo1 Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
RI	31	Read Inhibit. If this bit is set, an attempt to read data from the page causes a TLB Invalid exception, even if the <i>V</i> (Valid) bit is set. The <i>RI</i> bit is enabled only if the <i>RIE</i> bit of the <i>PageGrain</i> register is set. If the <i>RIE</i> bit of <i>PageGrain</i> is not set, the <i>RI</i> bit of <i>EntryLo0/EntryLo1</i> is a reserved 0 bit as per the MIPS32 specification.	R/W	0
XI	30	Execute Inhibit. If this bit is set, an attempt to fetch from the page causes a TLB Invalid exception, even if the <i>V</i> (Valid) bit is set. The <i>XI</i> bit is enabled only if the <i>XIE</i> bit of the <i>PageGrain</i> register is set. If the <i>XIE</i> bit of <i>PageGrain</i> is not set, the <i>XI</i> bit of <i>EntryLo0/EntryLo1</i> is a reserved 0 bit as per the MIPS32 specification.	R/W	0
0	29:26	These 4 bits are normally part of the PFN; however, because the core supports only 32 bits of physical address, the PFN is only 20 bits wide. Therefore, bits 29:26 of this register must be written with zeros.	R/W	0
PFN	25:6	Page Frame Number. Contributes to the definition of the high-order bits of the physical address. If the processor is enabled to support 1KB pages (<i>Config3SP</i> = 1 and <i>PageGrainESP</i> = 1), the PFN field corresponds to bits 29:10 of the physical address (the field is shifted left by 2 bits relative to the Release 1 definition to make room for PA _{11 10}). If the processor is not enabled to support 1KB pages (<i>Config3SP</i> = 0 or <i>PageGrainESP</i> = 0), the PFN field corresponds to bits 31..12 of the physical address.	R/W	Undefined
C	5:3	Coherency attribute of the page. See Table 8.6 .	R/W	Undefined
D	2	“Dirty” or write-enable bit, indicating that the page has been written, and/or is writable. If this bit is a one, then stores to the page are permitted. If this bit is a zero, then stores to the page cause a TLB Modified exception.	R/W	Undefined
V	1	Valid bit, indicating that the TLB entry, and thus the virtual page mapping, are valid. If this bit is a one, then accesses to the page are permitted. If this bit is a zero, then accesses to the page cause a TLB Invalid exception.	R/W	Undefined
G	0	Global bit. On a TLB write, the logical AND of the G bits in both the <i>EntryLo0</i> and <i>EntryLo1</i> register becomes 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 <i>EntryLo0</i> and <i>EntryLo1</i> reflect the state of the TLB G bit.	R/W	Undefined

Table 8.6 lists the encoding of the *C* field of the *EntryLo0* and *EntryLo1* registers and the *K0* field of the *Config* register.

Table 8.6 Cache Coherency Attributes

C[5:3] Value	Cache Coherency Attribute
0	Cacheable, noncoherent, write-through, no write allocate
1	Cacheable, noncoherent, write-through, write allocate
3*, 4, 5, 6	Cacheable, noncoherent, write-back, write allocate
2*, 7	Uncached
* These two values are required by the MIPS32 architecture. Only values 0, 1, 2 and 3 are used in an M5100 core. For example, values 4, 5 and 6 are not used and are mapped to 3. The value 7 is not used and is mapped to 2. Note that these values do have meaning in other MIPS Technologies processor implementations. Refer to the MIPS32 specification for more information.	

8.2.4 Context Register (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.

A TLB exception (TLB Refill, TLB Invalid, or TLB Modified) causes bits $VA_{31:13}$ of the virtual address to be written into the *BadVPN2* field of the *Context* register. The *PTEBase* field is written and used by the operating system.

The *BadVPN2* field of the *Context* register is not defined after an address error exception.

Figure 8.4 Context Register Format

PTEBase	BadVPN2	0
---------	---------	---

Table 8.7 Context Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
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 <i>Context</i> Register as a pointer into the current PTE array in memory.	R/W	Undefined
BadVPN2	22:4	This field is written by hardware on a TLB miss. It contains bits $VA_{31:13}$ of the virtual address that missed.	R	Undefined
0	3:0	Must be written as zero; returns zero on reads.	0	0

8.2.5 UserLocal Register (CP0 Register 4, Select 2)

The *UserLocal* register is a read-write register that is not interpreted by the hardware and conditionally readable via the RDHWR instruction.

Figure 8.5 shows the format of the *UserLocal* register; Table 8.8 describes the *UserLocal* register fields.

Figure 8.5 UserLocal Register Format




Table 8.8 UserLocal Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
UserLocal	31:0	This field contains software information that is not interpreted by hardware.	R/W	Undefined

Programming Notes

Privileged software may write this register with arbitrary information and make it accessible to unprivileged software via register 29 (*ULR*) of the RDHWR instruction. To do so, bit 29 of the *HWREna* register must be set to a 1 to enable unprivileged access to the register. In some operating environments, the *UserLocal* register contains a pointer to a thread-specific storage block that is obtained via the *RDHWR* register.

8.2.6 PageMask Register (CP0 Register 5, Select 0)

The *PageMask* register is a read/write register used for reading from and writing to the TLB. It holds a comparison mask that sets the variable page size for each TLB entry, as shown in Table 8.10. Figure 8.6 shows the format of the *PageMask* register; Table 8.9 describes the *PageMask* register fields.

Figure 8.6 PageMask Register Format




Table 8.9 PageMask Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
Mask	28:13	The <i>Mask</i> field is a bit mask in which a “1” bit indicates that the corresponding bit of the virtual address should not participate in the TLB match.	R/W	Undefined

Table 8.9 PageMask Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
MaskX	12:11	<p>In Release 2 of the Architecture, the <i>MaskX</i> field is an extension to the <i>Mask</i> field to support 1KB pages with definition and action analogous to that of the <i>Mask</i> field, defined above.</p> <ul style="list-style-type: none"> If 1KB pages are enabled ($Config3_{SP} = 1$ and $PageGrain_{ESP} = 1$), these bits are writable and readable, and their values are copied to and from the TLB entry on a TLB write or read, respectively. If 1KB pages are not enabled ($Config3_{SP} = 0$ or $PageGrain_{ESP} = 0$), these bits are not writable, return zero on read, and the effect on the TLB entry on a write is as if they were written with the value $2\#11$. <p>In Release 1 of the Architecture, these bits must be written as zero, return zero on read, and have no effect on the virtual address translation.</p>	R/W	0 (See Description)
0	31:29, 10:0	Ignored on writes; returns zero on reads.	R	0

Table 8.10 Values for the Mask and MaskX¹ Fields of the PageMask Register

Page Size	Bit																	
	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12 ¹	11 ¹
4 KBytes	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
16 KBytes	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1
64 KBytes	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
256 KBytes	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
1 MByte	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1
4 MByte	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
16 MByte	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
64 MByte	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
256 MByte	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

1. $PageMask_{12-11} = PageMask_{MaskX}$ exists only on implementations of Release 2 of the architecture and are treated as if they had the value $2\#11$ if 1K pages are not enabled ($Config3_{SP} = 0$ or $PageGrain_{ESP} = 0$).

It is implementation-dependent how many of the encodings described in Table 8.10 are implemented. All processors must implement the 4KB page size. If a particular page size encoding is not implemented by a processor, a read of the *PageMask* register must return zeros in all bits that correspond to encodings that are not implemented, thereby potentially returning a value different than that written by software.

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 pro-

processor is **UNDEFINED** if software loads the *Mask* field with a value other than one of those listed in [Table 8.10](#) even if the hardware returns a different value on a read. Hardware may depend on this requirement in implementing hardware structures.

8.2.7 PageGrain Register (CP0 Register 5, Select 1)

The *PageGrain* register is a read/write register used for enabling 1KB smaller than 4KB-page support. It is used for reading from and writing to the TLB.

The contents of the *PageGrain* register are not reflected in the contents of the TLB; therefore, the TLB must be flushed before any change to the *PageGrain* register is made. Behavior is **UNDEFINED** if a value other than those listed is used.

Figure 8.7 PageGrain Register Format




Table 8.11 PageGrain Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
RIE	31	Read Inhibit Enable. <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 20%;">Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td><i>RI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is disabled and not writeable by software.</td> </tr> <tr> <td style="text-align: center;">1</td> <td><i>RI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is enabled.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>RI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is disabled and not writeable by software.	1	<i>RI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is enabled.	R/W	0
Encoding	Meaning									
0	<i>RI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is disabled and not writeable by software.									
1	<i>RI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is enabled.									
XIE	30	Execute Inhibit Enable. <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 20%;">Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td><i>XI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is disabled and not writeable by software.</td> </tr> <tr> <td style="text-align: center;">1</td> <td><i>XI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>XI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is disabled and not writeable by software.	1	<i>XI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is enabled	R/W	0
Encoding	Meaning									
0	<i>XI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is disabled and not writeable by software.									
1	<i>XI</i> bit of <i>EntryLo0</i> and <i>EntryLo1</i> registers is enabled									
0	29	Reserved. Must be written as zero; returns zero on read.	0	0						

Table 8.11 PageGrain Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
ESP	28	<p>Enables support for 1KB pages.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>1KB page support is not enabled</td> </tr> <tr> <td>1</td> <td>1KB page support is enabled</td> </tr> </tbody> </table> <p>If this bit is a 1, the following changes occur to coprocessor 0 registers:</p> <ul style="list-style-type: none"> The <i>PFN</i> field of the <i>EntryLo0</i> and <i>EntryLo1</i> registers holds the physical address down to bit 10 (the field is shifted left by 2 bits from the Release 1 definition) The <i>MaskX</i> field of the <i>PageMask</i> register is writable and is concatenated to the right of the <i>Mask</i> field to form the “don’t care” mask for the TLB entry. The <i>VPN2X</i> field of the <i>EntryHi</i> register and bits 12:11 of the virtual address are writable. The virtual address translation algorithm is modified to reflect the smaller page size. <p>If <i>Config3_{SP}</i> = 0, 1KB pages are not implemented, and this bit is ignored on writes and returns zero on reads.</p>	Encoding	Meaning	0	1KB page support is not enabled	1	1KB page support is enabled	R/W	0
Encoding	Meaning									
0	1KB page support is not enabled									
1	1KB page support is enabled									
IEC	27	<p>Enables unique exception codes for the Read-Inhibit and Execute-Inhibit exceptions.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Read-Inhibit and Execute-Inhibit exceptions both use the TLBL exception code.</td> </tr> <tr> <td>1</td> <td>Read-Inhibit exceptions use the TLBRI exception code. Execute-Inhibit exceptions use the TLBXI exception code.</td> </tr> </tbody> </table>	Encoding	Meaning	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
Encoding	Meaning									
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.									
0	26:0	Must be written as zero; returns zero on reads.	0	0						

8.2.8 Wired Register (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 8.8. The width of the *Wired* field is calculated in the same manner as that described for the *Index* register above. Wired entries are fixed, non-replaceable entries that are not overwritten by a TLBWR instruction. Wired entries can be overwritten by a TLBWI instruction.

The *Wired* register is reset to zero by a Reset exception. Writing the *Wired* register causes the *Random* register to reset to its upper bound.

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.

This register is only valid with a TLB-based MMU cores. It is reserved for a FM based MMU core.

Figure 8.8 Wired and Random Entries in the TLB

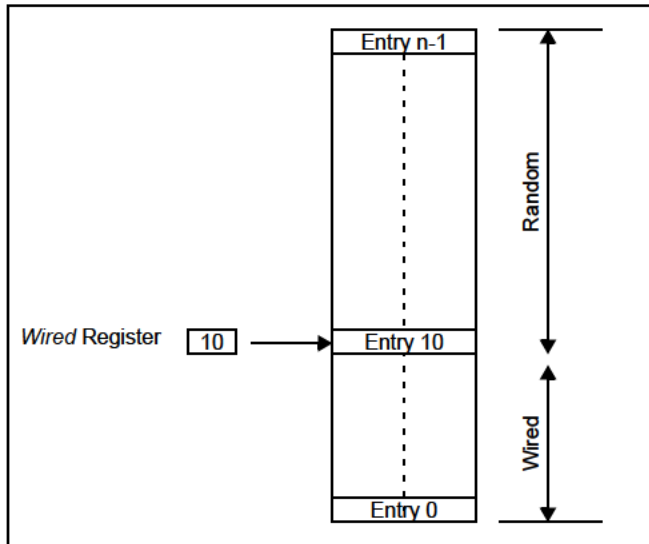



Figure 8.9 Wired Register Format



8.2.9 HWREna Register (CP0 Register 7, Select 0)

The *HWREna* register contains a bit mask that determines which hardware registers are accessible via the RDHWR instruction.

Figure 8.10 shows the format of the *HWREna* Register; Table 8.12 describes the *HWREna* register fields.

Figure 8.10 HWREna Register Format




Table 8.12 HWREna Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
0	31:30	Must be written with zero; returns zero on read	0	0
0	28:4	Must be written with zero; returns zero on read	0	0
ULR	29	User Local Register. This register provides read access to the coprocessor 0 <i>UserLocal</i> register. In some operating environments, the <i>UserLocal</i> register is a pointer to a thread-specific storage block.	R/W	0

Table 8.12 HWREna Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State
Name	Bits			
Mask	3:0	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. See the RDHWR instruction for a list of valid hardware registers.	R/W	0


Privileged software may determine which of the hardware registers are accessible by 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.

8.2.10 BadVAddr Register (CP0 Register 8, Select 0)

The *BadVAddr* register is a read-only register that captures the most recent virtual address that caused the following exception:

- Address error (AdEL or AdES)
- TLB Refill
- TLB Invalid
- TLB Modified

The *BadVAddr* register does not capture address information for bus errors, because they are not addressing errors.

Figure 8.11 BadVAddr Register Format**Table 8.13 BadVAddr Register Field Description**

Fields		Description	Read/Write	Reset State
Name	Bits			
BadVAddr	31:0	Bad virtual address.	R	Undefined

8.2.11 BadInstr Register (CP0 Register 8, Select 1)

The *BadInstr* register is an optional read-only register that captures the most recent instruction that caused one of the following exceptions:

CP0 Registers of the M5100™ Core

- Instruction Validity
Coprorocessor Unusable, Reserved Instruction
- Execution Exception
Integer Overflow, Trap, System Call, Breakpoint, Floating-point, Coprocessor 2 exception
- Addressing
Address Error, TLB Refill, TLB Invalid, TLB Read Inhibit, TLB Execute Inhibit, TLB Modified

The *BadInstr* register is provided to allow acceleration of instruction emulation. The *BadInstr* register is only set by exceptions that are synchronous to an instruction. The *BadInstr* register is not set by Interrupts or by NMI, Machine check, Bus Error, Cache Error, Watch, or EJTAG exceptions.

When a synchronous exception occurs for which there is no valid instruction word (for example TLB Refill - Instruction Fetch), the value stored in *BadInstr* is **UNPREDICTABLE**.

Presence of the *BadInstr* register is indicated by the *Config3_{BI}* bit. The *BadInstr* register is instantiated per-VPE in an MT ASE processor.

Figure 8.12 shows the proposed format of the *BadInstr* register; Table 8.14 describes the *BadInstr* register fields.

Figure 8.12 BadInstr Register Format




Table 8.14 BadInstr Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
BadInstr	31:0	Faulting instruction word. Instruction words smaller than 32 bits are placed in bits 15:0, with bits 31:16 containing zero.	R	Undefined

8.2.12 BadInstrP Register (CP0 Register 8, Select 2)

The *BadInstrP* register is an optional register that is used in conjunction with the *BadInstr* register. The *BadInstrP* register contains the prior branch instruction when the faulting instruction is in a branch delay slot.

The *BadInstrP* register is updated for these exceptions:

- Instruction Validity
Coprorocessor Unusable, Reserved Instruction
- Execution Exception
Integer Overflow, Trap, System Call, Breakpoint, Floating-point, Coprocessor 2 exception

- Addressing

Address Error, TLB Refill, TLB Invalid, TLB Read Inhibit, TLB Execute Inhibit, TLB Modified

The *BadInstrP* register is provided to allow acceleration of instruction emulation. The *BadInstrP* register is only set by exceptions that are synchronous to an instruction. The *BadInstrP* register is not set by Interrupts or by NMI, Machine check, Bus Error, Cache Error, Watch, or EJTAG exceptions. When a synchronous exception occurs, and the faulting instruction is not in a branch delay slot, then the value stored in *BadInstrP* is **UNPREDICTABLE**.

Presence of the *BadInstrP* register is indicated by the *Config3BP* bit. The *BadInstrP* register is instantiated per-VPE in an MT ASE processor.

Figure 8.13 shows the proposed format of the *BadInstrP* register; Table 8.15 describes the *BadInstrP* register fields.

Figure 8.13 BadInstrP Register Format




Table 8.15 BadInstrP Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
BadInstrP	31:0	Prior branch instruction. Instruction words smaller than 32 bits are placed in bits 15:0, with bits 31:16 containing zero.	R	Undefined

8.2.13 Count Register (CP0 Register 9, Select 0)

The *Count* register acts as a timer, incrementing at a constant rate, whether or not an instruction is executed, retired, or any forward progress is made through the pipeline. The counter increments every other clock if the *DC* bit in the *Cause* register is 0.

The *Count* register can be written for functional or diagnostic purposes, including at reset or to synchronize processors.

By writing the *CountDM* bit in the *Debug* register, it is possible to control whether the *Count* register continues incrementing while the processor is in debug mode.

Figure 8.14 Count Register Format




Table 8.16 Count Register Field Description

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

8.2.14 EntryHi Register (CP0 Register 10, Select 0)

The *EntryHi* register contains the virtual address match information used for TLB read, write, and access operations.

A TLB exception (TLB Refill, TLB Invalid, or TLB Modified) causes bits $VA_{31..13}$ of the virtual address to be written into the *VPN2* field of the *EntryHi* register. An implementation of Release 2 of the Architecture which supports 1KB pages also writes $VA_{12..11}$ into the *VPN2X* 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* field is 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 *VPNX2* and *VPN2* fields of the *EntryHi* register are not defined after an address error exception and these 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 write of address-related fields in the *BadVAddr*, *Context* registers.

Figure 8.15 EntryHi Register Format



Table 8.17 EntryHi Register Field Descriptions

VPN2	31:1	$VA_{31..1}$ of the virtual address (virtual page number / 2). This field is written by hardware on a TLB exception or on a TLB read, and is written by software before a TLB write.	R/W	Undefined
VPN2X	12:11	In Release 2 of the Architecture, the <i>VPN2X</i> field is an extension to the <i>VPN2</i> field to support 1KB pages. These bits are not writable by either hardware or software unless <i>Config3_{SP}</i> = 1 and <i>PageGrain_{ESP}</i> = 1. If enabled for write, this field contains $VA_{12..11}$ of the virtual address and is written by hardware on a TLB exception or on a TLB read, and is by software before a TLB write. If writes are not enabled, and in implementations of Release 1 of the Architecture, this field must be written with zero and returns zeros on read.	R/W	0
EHINV	10	TLB HW Invalidate bit. If this bit is set, the TLBWI instruction will invalidate the <i>VPN2</i> field of the selected TLB entry. TLBR instruction will update this field with the <i>VPN2</i> invalid bit of the read TLB entry.	R/W	0
0	9:8	Must be written as zero; returns zero on read.	0	0

Table 8.17 EntryHi Register Field Descriptions (Continued)

ASID	7:0	Address space identifier. This field is written by hardware on a TLBread and by software to establish the current ASID value for TLB write and against which TLB references match each entry's TLB <i>ASID</i> field.	R/W	Undefined
------	-----	---	-----	-----------

8.2.15 GuestCtl1 Register (CP0 Register 10, Select 4)

Compliance Level: *Optional* in the Virtualization Module.

The *GuestCtl1* register defines GuestID control fields for Root (*GuestCtl1RID*) and Guest (*GuestCtl1ID*) that may be used in the context of TLB instructions, instruction and data address translation. The *GuestCtl1RID* field additionally is written by the processor on a TLBR or TLBGR instruction in Root mode, then containing the GuestID read from the TLB entry. A TLBR executed in Guest mode does not cause a write to either *GuestCtl1ID* or *GuestCtl1RID*.

GuestCtl1 is optional, and thus the use of *GuestID* is optional in the context of TLB instructions, and instruction and data address translation. The *GuestCtl1* register only exists in Root Context. GuestID value of 0 is reserved for Root.

The Section titled “Virtualized MMU GuestID Use” in [14] or [15] provides additional detail on GuestID usage as it applies to instruction and data address translation. The Section titled “New CP0 Instructions” of in [14] or [15] describes the TLB instructions and their use of GuestID.

The primary purpose of the GuestID is to provide a unique component of the Guest/Root TLB entry, thus eliminating TLB invalidation overhead on a virtual-machine-level context switch.

Systems that implement a GuestID are required to support a guest identifier field (GID) in each Guest and Root TLB entry. This GuestID field within the TLB is not accessible to the Guest. While operating in guest context, the behavior of guest TLB operations is constrained by the *GuestCtl1ID* field so that only guest TLB entries with a matching GID field are considered.

The actual number of usable bits in the *GuestCtl1ID* and *GuestCtl1RID* fields is implementation-dependent. Software may determine the usable size of these fields by writing all ones and reading the value back. The size of *GuestCtl1ID* and *GuestCtl1RID* must be equal.

The *GuestCtl1* register is instantiated per-VPE in an MT Module processor.

Figure 8.16 shows the format of the Virtualization Module *GuestCtl1* register; Table 8.18 describes the *GuestCtl1* register fields.


Figure 8.16 GuestCtl1 Register Format

Table 8.18 GuestCtl1 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
EID	31:24	External Interrupt Controller Guest ID. Required if an External Interrupt Controller (EIC) is supported. A guest interrupt that is posted by the <i>EIC</i> to the root interrupt bus must cause the Guest ID of the root interrupt bus to be registered in <i>EID</i> when the interrupt is taken. If implemented, the field is read-only and set by hardware. If not implemented, this field must be written as zero and return zero on reads.	R0 or R	0
RID	23:16	Root control GuestID. Used by root TLB operations, and by <i>GuestCtl0_{DRG}</i> =1 mode.	R/W	0
0	15:8	Must be written as zero; returns zero on read.	R0	0
ID	7:0	Guest control GuestID. Identifies resident guest. Applies to guest address translation.	R/W	0

8.2.16 GuestCtl2 Register (CP0 Register 10, Select 5)

Compliance Level: *Optional* in the Virtualization Module.

The *GuestCtl2* register is optional in an implementation. It is only required if support for virtual interrupts in non-EIC mode is included in an implementation. Alternatively, if EIC mode is supported, then *GuestCtl2* is required. Refer to the Section titled “External Interrupts” in [14] or [15] for a description of interrupt handling in EIC and non-EIC modes.

An implementation that supports the virtual interrupt functionality of *GuestCtl2* is not required to support root writes of *Guest.CauseIP*[7:2] or *Guest.CauseRIPL* as described in the Table titled “Guest CP0 Read-only Fields Writable from Root Mode” in [14] or [15].

GuestCtl2 is present in an implementation if *GuestCtl0_{G2}*=1.

Figure 8.17 shows the format of the Virtualization Module *GuestCtl2* register in non-EIC mode. Table 8.19 describes the non-EIC mode *GuestCtl2* register fields.

Figure 8.18 shows the format of the Virtualization Module *GuestCtl2* register in EIC mode. Table 8.20 describes the EIC-mode *GuestCtl2* register fields.

Figure 8.17 GuestCtl2 Register Format for non-EIC mode




Figure 8.18 GuestCtl2 Register Format for EIC mode




Table 8.19 Non-EIC Mode GuestCtl2 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
HC	31:24	<p>Hardware Clear for <i>GuestCtl2_{VIP}</i></p> <p>This set of bits maps one to one to <i>GuestCtl2_{VIP}</i>. <i>HC</i> may be bit-wise Read-only or R/W. If a bit is Read-only, then it may be preset to 0 or 1. Similarly, if a bit is R/W, then it may be reset to 0 or 1. The interpretation of 0 or 1 state follows.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>The deassertion of related external interrupt (IRQ[n]) has no effect on <i>GuestCtl2_{VIP}[n]</i>. Root software must write zero to <i>GuestCtl2_{VIP}[n]</i> to clear the virtual interrupt.</td> </tr> <tr> <td>1</td> <td>The deassertion of related external interrupt (IRQ[n]) causes <i>GuestCtl2_{VIP}[n]</i> to be cleared by hardware.</td> </tr> </tbody> </table> <p>In the case of <i>HC</i>=0, <i>Guest.Status_{IP}[n+2]</i> could continue to be asserted due to an external interrupt when <i>GuestCtl2_{VIP}[n]</i> is cleared by software. The source of the external interrupt must be serviced appropriately.</p> <p>The choice of Read-only vs. R/W is implementation-dependent. Root software can write then read this field to determine the supported configuration.</p>	Encoding	Meaning	0	The deassertion of related external interrupt (IRQ[n]) has no effect on <i>GuestCtl2_{VIP}[n]</i> . Root software must write zero to <i>GuestCtl2_{VIP}[n]</i> to clear the virtual interrupt.	1	The deassertion of related external interrupt (IRQ[n]) causes <i>GuestCtl2_{VIP}[n]</i> to be cleared by hardware.	R, R/W	0 or 1
Encoding	Meaning									
0	The deassertion of related external interrupt (IRQ[n]) has no effect on <i>GuestCtl2_{VIP}[n]</i> . Root software must write zero to <i>GuestCtl2_{VIP}[n]</i> to clear the virtual interrupt.									
1	The deassertion of related external interrupt (IRQ[n]) causes <i>GuestCtl2_{VIP}[n]</i> to be cleared by hardware.									
0	25:18	Must be written as zero; returns zero on read.	R0	0						

Table 8.19 Non-EIC Mode GuestCtl2 Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bits									
VIP	17:10	<p>Virtual Interrupt Pending. The <i>VIP</i> field is used by the root to inject virtual interrupts into the Guest context. <i>VIP</i>[5:0] maps to <i>Guest.Status_{IP}</i>[7:2]. <i>VIP</i> effects <i>Guest.Status_{IP}</i> in the following manner:</p> <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><i>Guest.Status_{IP}</i>[n+2] cannot be asserted due to <i>VIP</i>[n], though it may be asserted by an external interrupt <i>IRQ</i>[n], n = 5..0</td> </tr> <tr> <td>1</td> <td><i>Guest.Status_{IP}</i>[n+2] must at least be asserted due to <i>VIP</i>[n]. It may also be asserted by a concurrent external interrupt. n=5..0</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>Guest.Status_{IP}</i> [n+2] cannot be asserted due to <i>VIP</i> [n], though it may be asserted by an external interrupt <i>IRQ</i> [n], n = 5..0	1	<i>Guest.Status_{IP}</i> [n+2] must at least be asserted due to <i>VIP</i> [n]. It may also be asserted by a concurrent external interrupt. n=5..0	R/W	0
Encoding	Meaning									
0	<i>Guest.Status_{IP}</i> [n+2] cannot be asserted due to <i>VIP</i> [n], though it may be asserted by an external interrupt <i>IRQ</i> [n], n = 5..0									
1	<i>Guest.Status_{IP}</i> [n+2] must at least be asserted due to <i>VIP</i> [n]. It may also be asserted by a concurrent external interrupt. n=5..0									
0	9:5	Must be written as zero; returns zero on read.	R0	0						
Impl	4:0	<p>Implementation. These bits are implementation-dependent and are not defined by the architecture. If not implemented, they must be written as zero and read as zero. If implemented, and if modifying the behavior of the processor, this field must be defined in such a way that correct behavior is preserved if software, with no knowledge of these bits, reads the <i>GuestCtl2</i> register, modifies another field, and writes the updated value back to the <i>GuestCtl2</i> register.</p>	R/W	0						

I

Table 8.20 EIC Mode GuestCtl2 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
GRIPL	31:24	<p>Guest RIPL</p> <p>This field is written only when an interrupt received on the root interrupt bus for a guest is taken. The <i>RIPL</i> (Requested Interrupt Priority Level) sent by <i>EIC</i> on the root interrupt bus is written to this field.</p> <p>Root software can write the field if it needs to modify the <i>EIC</i> value before assigning it to guest. It may also clear this field to prevent a transition to guest mode from causing an interrupt if this field was set earlier with a non-zero value.</p> <p><i>GRIPL</i> is 10 bits only for an implementation that complies with the MCU Module; otherwise it is 8 bits as in the baseline architecture.</p>	R/W	0
GEICSS	21:18	<p>Guest EICSS</p> <p>This field is written only when an interrupt received on the root interrupt bus for a guest is taken. The <i>EICSS</i> (External Interrupt Controller Shadow Set) sent by <i>EIC</i> on the root interrupt bus is written to this field.</p> <p>Root software can write this field if it needs to modify the <i>EIC</i> value before assigning to a guest.</p>	R/W	Undefined
0	17:16	Must be written as zero; returns zero on read.	R0	0
GVEC	15:0	<p>Guest Vector</p> <p>This field is written only when an interrupt is received on the root interrupt bus for a guest. The Vector Offset (or Number) sent by <i>EIC</i> on the root interrupt bus is written to this field.</p> <p><i>GVEC</i> is not loaded into any guest CP0 field, but is used to generate an interrupt vector in guest mode, using the root interrupt bus vector and not the guest interrupt bus vector. This will only occur if the interrupt was first taken in root mode.</p> <p>Root software can write the field if it needs to modify the <i>EIC</i> value before assigning to a guest.</p>	R/W	Undefined

8.2.17 GuestCtl3 Register (CP0 Register 10, Select 6)

Compliance Level: *Optional* in the Virtualization Module.

The *GuestCtl3* register is optional. It is required only if Shadow GPR Sets are supported, and the Shadow Sets used by a guest are virtual and require mapping to physical Shadow Sets. With this mechanism, a pool of Shadow Sets can be physically shared by partitioning the sets among multiple guests and root, under root control.

Virtual mapping of Guest GPR set(s) is supported if Guest *SRSCtl_{HSS}* is writeable by root. Presence of *GuestCtl3* can be detected by root software by writing any non-zero value less than or equal to root *SRSCtl_{HSS}* to Guest *SRSCtl_{HSS}*. If a read returns the value written, then *GuestCtl3* is present.

The *GuestCtl3* register is instantiated per-VPE in an MT Module processor.

Figure 8.19 shows the format of the Virtualization Module *GuestCtl3* register; Table 8.21 describes the *GuestCtl3* register fields.

Figure 8.19 GuestCtl3 Register Format




Table 8.21 GuestCtl3 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:4	This bit must be written as zero, returns zero on read.	R0	0
GLSS	3:0	Guest Lowest Shadow Set number. This determines the lowest physical Shadow Set number assigned by root to guest. Guest is thus assigned physical Shadow Sets <i>GLSS</i> to <i>GLSS</i> plus Guest <i>SRSCtl_{HSS}</i> . If this field is reserved, then all writes must be zero, and reads should return 0.	R0, R/W	0

8.2.18 Compare Register (CP0 Register 11, Select 0)

The *Compare* register acts in conjunction with the *Count* register to implement a timer and timer interrupt function. The timer interrupt is an output of the cores. The *Compare* register maintains a stable value and does not change on its own.

When the value of the *Count* register equals the value of the *Compare* register, the *SI_TimerInt* pin is asserted. This pin will remain asserted until the *Compare* register is written. The *SI_TimerInt* pin can be fed back into the core on one of the interrupt pins to generate an interrupt. However, this is no longer needed as the core will internally route the interrupt to the IP number set by the *IntCtl.IPTI* field.

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

Figure 8.20 Compare Register Format




Table 8.22 Compare Register Field Description

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

8.2.19 GuestCtl0Ext Register (CP0 Register 11, Select 4)

Compliance Level: *Optional* in the Virtualization Module.

GuestCtl0Ext is an optional extension to *GuestCtl0*. If not supported, the register must read as 0.

GuestCtl0_{GOE} should be read by software to determine if *GuestCtl0Ext* is implemented.

The *GuestCtl0Ext* register is instantiated per-VPE in an MT Module processor.

Figure 8.21 shows the format of the Virtualization Module *GuestCtl0Ext* register; Table 8.23 describes the *GuestCtl0Ext* register fields.

Figure 8.21 GuestCtl0Ext Register Format




Table 8.23 GuestCtl0Ext Register Field Descriptions

Fields		Description	Read / Write	Reset State										
Name	Bits													
0	31:8,5	This bit must be written as zero, returns zero on read.	R0	0										
NCC	7:6	Nested Cache Coherency Attributes Determines whether guest CCA is modified by root CCA in 2nd step of guest address translation. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Guest CCA is independent of root CCA</td> </tr> <tr> <td>01</td> <td>Guest CCA is modified by root CCA in manner described in “Guest Nested CCA support” section description of MD00846 or MD00848.</td> </tr> <tr> <td>10</td> <td>Reserved</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	00	Guest CCA is independent of root CCA	01	Guest CCA is modified by root CCA in manner described in “Guest Nested CCA support” section description of MD00846 or MD00848.	10	Reserved	11	Reserved	R	01
Encoding	Meaning													
00	Guest CCA is independent of root CCA													
01	Guest CCA is modified by root CCA in manner described in “Guest Nested CCA support” section description of MD00846 or MD00848.													
10	Reserved													
11	Reserved													
CGI	4	Related to <i>GuestCtl0_{CG}</i> Allows execution of CACHE Index Invalidate operations in guest mode. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Definition of <i>GuestCtl0_{CG}</i> does not change.</td> </tr> <tr> <td>1</td> <td>If <i>GuestCtl0_{CG}</i>=1, then all CACHE Index Invalidate operations may execute in guest mode without causing a GPSI.</td> </tr> </tbody> </table> <p>This field is R0 if such support is not implemented.</p>	Encoding	Meaning	0	Definition of <i>GuestCtl0_{CG}</i> does not change.	1	If <i>GuestCtl0_{CG}</i> =1, then all CACHE Index Invalidate operations may execute in guest mode without causing a GPSI.	R0, R/W	0				
Encoding	Meaning													
0	Definition of <i>GuestCtl0_{CG}</i> does not change.													
1	If <i>GuestCtl0_{CG}</i> =1, then all CACHE Index Invalidate operations may execute in guest mode without causing a GPSI.													

Table 8.23 GuestCtl0Ext Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bits									
FCD	3	<p>Disables Guest Software/Hardware Field Change Exceptions (GSFC/GHFC).</p> <p>This mode is useful for an implementation with root software that is not a full-featured hypervisor. For example, the software may just support memory protection, but may not require protection of CP0 state.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GSFC or GHFC event will cause exception.</td> </tr> <tr> <td>1</td> <td>GSFC or GHFC event will not cause exception.</td> </tr> </tbody> </table>	Encoding	Meaning	0	GSFC or GHFC event will cause exception.	1	GSFC or GHFC event will not cause exception.	R/W	0
Encoding	Meaning									
0	GSFC or GHFC event will cause exception.									
1	GSFC or GHFC event will not cause exception.									
OG	2	<p>Other GPSI Enable. Applies to <i>HWREna</i>, <i>LLAddr</i>, <i>UserTraceData1</i>, <i>UserTraceData2</i>, <i>KScratch1</i> through <i>KScratch6</i>, when implemented. If function is not supported, this field reads as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GPSI not enabled for these registers unless <i>GuestCtl0CP0</i>=0.</td> </tr> <tr> <td>1</td> <td>GPSI enabled for these registers.</td> </tr> </tbody> </table> <p><i>UserTraceData1</i>, <i>UserTraceData2</i> are optional CP0 registers defined in the MIPS iFlowtrace specification.</p>	Encoding	Meaning	0	GPSI not enabled for these registers unless <i>GuestCtl0CP0</i> =0.	1	GPSI enabled for these registers.	R0	0
Encoding	Meaning									
0	GPSI not enabled for these registers unless <i>GuestCtl0CP0</i> =0.									
1	GPSI enabled for these registers.									
BG	1	<p>Bad register GPSI Enable. Applies to <i>BadVAddr</i>, <i>BadInstr</i>, and <i>BadInstrP</i> when implemented. If function is not supported, this field reads as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GPSI not enabled for these registers unless <i>GuestCtl0CP0</i>=0.</td> </tr> <tr> <td>1</td> <td>GPSI enabled for these registers.</td> </tr> </tbody> </table>	Encoding	Meaning	0	GPSI not enabled for these registers unless <i>GuestCtl0CP0</i> =0.	1	GPSI enabled for these registers.	R0	0
Encoding	Meaning									
0	GPSI not enabled for these registers unless <i>GuestCtl0CP0</i> =0.									
1	GPSI enabled for these registers.									
MG	0	<p>MMU GPSI Enable. Applies to <i>Index</i>, <i>Random</i>, <i>EntryLo0</i>, <i>EntryLo1</i>, <i>Context</i>, <i>ContextConfig</i>, <i>UserLocal</i>, <i>PageMask</i>, and <i>EntryHi</i>. If function is not supported, this field reads as 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GPSI not enabled for these registers unless <i>GuestCtl0CP0</i>=0.</td> </tr> <tr> <td>1</td> <td>GPSI enabled for these registers.</td> </tr> </tbody> </table>	Encoding	Meaning	0	GPSI not enabled for these registers unless <i>GuestCtl0CP0</i> =0.	1	GPSI enabled for these registers.	R0	0
Encoding	Meaning									
0	GPSI not enabled for these registers unless <i>GuestCtl0CP0</i> =0.									
1	GPSI enabled for these registers.									

8.2.20 Status Register (CP0 Register 12, Select 0)

The *Status* register is a read/write register that contains the operating mode, interrupt enabling, and the diagnostic states of the processor. Fields of this register combine to create operating modes for the processor. Refer to 6.2 “Modes of Operation” on page 108 for a discussion of operating modes, and 7.3 “Interrupts” on page 121 for a discussion of interrupt modes.

Interrupt Enable: Interrupts are enabled when all of the following conditions are true:

- $IE = 1$
- $EXL = 0$
- $ERL = 0$
- $DM = 0$

If these conditions are met, then the settings of the *IM* and *IE* bits enable the interrupts.

Operating Modes: If the *DM* bit in the *Debug* register is 1, then the processor is in debug mode; otherwise the processor is in either kernel or user mode. The following CPU *Status* register bit settings determine user or kernel mode:

- User mode: $UM = 1$, $EXL = 0$, and $ERL = 0$
- Kernel mode: $UM = 0$, or $EXL = 1$, or $ERL = 1$

Coprocessor Accessibility: The *Status* register *CU* bits control coprocessor accessibility. If any coprocessor is unable, then an instruction that accesses it generates an exception.

Figure 8.22 shows the format of the *Status* register; Table 8.24 describes the *Status* register fields.

Figure 8.22 Status Register Format




Table 8.24 Status Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
CU3	31	Controls access to coprocessor 3. COP3 is not supported. This bit cannot be written and will read as 0.	R	0
CU2	30	Controls access to coprocessor 2. This bit can only be written if coprocessor is attached to the COP2 interface. (C2 bit in Config1 is set). This bit will read as 0 if no coprocessor is present.	R/W	0
CU1	29	Controls access to coprocessor 1. This bit can only be written if the FPU is configured. This bit will read as 0 if the FPU is not present.	R/W	0

Table 8.24 Status Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
CU0	28	<p>Controls access to coprocessor 0:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Access not allowed</td> </tr> <tr> <td>1</td> <td>Access allowed</td> </tr> </tbody> </table> <p>Coprocessor 0 is always usable when the processor is running in kernel mode, independent of the state of the <i>CU0</i> bit.</p>	Encoding	Meaning	0	Access not allowed	1	Access allowed	R/W	Undefined
Encoding	Meaning									
0	Access not allowed									
1	Access allowed									
RP	27	<p>Enables reduced power mode. The state of the <i>RP</i> bit is available on the external core interface as the <i>SI_RP</i> signal.</p>	R/W	0 for Cold Reset only.						
FR	26	<p>This bit is used to control the floating-point register mode for 64-bit floating-point units:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Floating-point registers can contain any 32-bit data type. 64-bit data types are stored in even-odd pairs of registers</td> </tr> <tr> <td>1</td> <td>Floating-point registers can contain any data type</td> </tr> </tbody> </table> <p>This bit must be ignored on write and read as zero under the following conditions:</p> <ul style="list-style-type: none"> • No floating-point unit is implementation • 64-bit floating-point unit is not implemented. 	Encoding	Meaning	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 data type	R/W	0
Encoding	Meaning									
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 data type									
RE	25	<p>Used to enable reverse-endian memory references while the processor is running in user mode:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>User mode uses configured endianness</td> </tr> <tr> <td>1</td> <td>User mode uses reversed endianness</td> </tr> </tbody> </table> <p>Neither Debug Mode nor Kernel Mode nor Supervisor Mode references are affected by the state of this bit.</p>	Encoding	Meaning	0	User mode uses configured endianness	1	User mode uses reversed endianness	R/W	Undefined
Encoding	Meaning									
0	User mode uses configured endianness									
1	User mode uses reversed endianness									
MX	24	<p>MIPS DSP Extension. Enables access to DSP Module resources:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Access not allowed</td> </tr> <tr> <td>1</td> <td>Access allowed</td> </tr> </tbody> </table> <p>An attempt to execute any DSP Module instruction before this bit has been set to 1 will cause a DSP State Disabled exception. The state of this bit is reflected in <i>Config3_DSPP</i></p>	Encoding	Meaning	0	Access not allowed	1	Access allowed	R/W	0
Encoding	Meaning									
0	Access not allowed									
1	Access allowed									
R	23	<p>Reserved. This field is ignored on writes and reads as 0.</p>	R	0						

Table 8.24 Status Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
BEV	22	Controls the location of exception vectors: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Bootstrap</td> </tr> </tbody> </table>	Encoding	Meaning	0	Normal	1	Bootstrap	R/W	1
Encoding	Meaning									
0	Normal									
1	Bootstrap									
TS	21	TLB shutdown. Indicates that the TLB has detected a match on multiple entries. This bit is set if a TLBWI or TLBWR instruction is issued that would cause a TLB shutdown condition if allowed to complete. A machine check exception is also issued. This bit is only used in the TLB-based MMU processors and is reserved in the FM-based MMU processors. Software can only write a 0 to this bit to clear it and cannot force a 0-1 transition.	R/W	0						
SR	20	Indicates that the entry through the reset exception vector was due to a Soft Reset: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Not Soft Reset (NMI or Reset)</td> </tr> <tr> <td>1</td> <td>Soft Reset</td> </tr> </tbody> </table> Software can only write a 0 to this bit to clear it and cannot force a 0-1 transition.	Encoding	Meaning	0	Not Soft Reset (NMI or Reset)	1	Soft Reset	R/W	1 for Soft Reset; 0 otherwise
Encoding	Meaning									
0	Not Soft Reset (NMI or Reset)									
1	Soft Reset									
NMI	19	Indicates that the entry through the reset exception vector was due to an NMI: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Not NMI (Soft Reset or Reset)</td> </tr> <tr> <td>1</td> <td>NMI</td> </tr> </tbody> </table> Software can only write a 0 to this bit to clear it and cannot force a 0-1 transition.	Encoding	Meaning	0	Not NMI (Soft Reset or Reset)	1	NMI	R/W	1 for NMI; 0 otherwise
Encoding	Meaning									
0	Not NMI (Soft Reset or Reset)									
1	NMI									
CEE	17	CorExtend Enable: Implementation-dependent. If CorExtend block indicates that this bit should be used, any attempt to execute a CorExtend instruction with this bit cleared will result in a CorExtend Unusable exception. This bit is reserved if CorExtend is not present.	R/W	Undefined						
IM9:IM2	18, 16:10	Interrupt Mask: Controls the enabling of each of the hardware interrupts. Refer to 7.3 “Interrupts” on page 121 for a complete discussion of enabled interrupts. <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> In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled ($Config3_{VEIC} = 1$), these bits have a different meaning and are interpreted as the <i>IPL</i> field, described below.	Encoding	Meaning	0	Interrupt request disabled	1	Interrupt request enabled	R/W	Undefined for IM7:IM2 0 for IM9:IM8
Encoding	Meaning									
0	Interrupt request disabled									
1	Interrupt request enabled									

Table 8.24 Status Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
IPL	18, 16:10	<p>Interrupt Priority Level.</p> <p>In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled ($Config3_{VEIC} = 1$), this field is the encoded (0:255) 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 ($Config3_{VEIC} = 0$), these bits have a different meaning and are interpreted as the $IM7..IM2$ bits, described above.</p>	R/W	Undefined for IPL15:IPL10 0 for IPL18:IPL17						
IM1:IM0	9:8	<p>Interrupt Mask: Controls the enabling of each of the software interrupts. Refer to Section 7.3 “Interrupts” for a complete discussion of enabled interrupts.</p> <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>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, 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									
R	7:5	Reserved. This field is ignored on writes and reads as 0.	R	0						
UM	4	<p>This bit denotes the base operating mode of the processor. See Section 6.2 “Modes of Operation” for a full discussion of operating modes. The encoding of this bit is:</p> <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>Base mode is Kernel Mode</td> </tr> <tr> <td>1</td> <td>Base mode is User Mode</td> </tr> </tbody> </table> <p>Note that the processor can also be in kernel mode if <i>ERL</i> or <i>EXL</i> is set, regardless of the state of the <i>UM</i> bit.</p>	Encoding	Meaning	0	Base mode is Kernel Mode	1	Base mode is User Mode	R/W	Undefined
Encoding	Meaning									
0	Base mode is Kernel Mode									
1	Base mode is User Mode									
R	3	This bit is reserved. This bit is ignored on writes and reads as zero.	R	0						

Table 8.24 Status Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
ERL	2	<p>Error Level; Set by the processor when a Reset, Soft Reset, NMI or Cache Error exception are taken.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal level</td> </tr> <tr> <td>1</td> <td>Error level</td> </tr> </tbody> </table> <p>When <i>ERL</i> is set:</p> <ul style="list-style-type: none"> • The processor is running in kernel mode • Interrupts are disabled • The ERET instruction will use the return address held in <i>ErrorEPC</i> instead of <i>EPC</i> • The lower 2^{29} bytes of kuseg are treated as an unmapped and uncached region. See Chapter 6, “Memory Management of the M5100™ Core” on page 107. This allows main memory to be accessed in the presence of cache errors. The operation of the processor is UNDEFINED if the <i>ERL</i> bit is set while the processor is executing instructions from kuseg. 	Encoding	Meaning	0	Normal level	1	Error level	R/W	1
Encoding	Meaning									
0	Normal level									
1	Error level									
EXL	1	<p>Exception Level; Set by the processor when any exception other than Reset, Soft Reset, or NMI exceptions is taken.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal level</td> </tr> <tr> <td>1</td> <td>Exception level</td> </tr> </tbody> </table> <p>When <i>EXL</i> is set:</p> <ul style="list-style-type: none"> • The processor is running in Kernel Mode • Interrupts are disabled. • <i>EPC</i>, <i>Cause_{BD}</i> and <i>SRSCtl</i> (implementations of Release 2 of the Architecture only) will not be updated if another exception is taken 	Encoding	Meaning	0	Normal level	1	Exception level	R/W	Undefined
Encoding	Meaning									
0	Normal level									
1	Exception level									
IE	0	<p>Interrupt Enable: Acts as the master enable for software and hardware interrupts:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupts are disabled</td> </tr> <tr> <td>1</td> <td>Interrupts are enabled</td> </tr> </tbody> </table> <p>In Release 2 of the Architecture, this bit may be modified separately via the DI and EI instructions.</p>	Encoding	Meaning	0	Interrupts are disabled	1	Interrupts are enabled	R/W	Undefined
Encoding	Meaning									
0	Interrupts are disabled									
1	Interrupts are enabled									

8.2.21 IntCtl Register (CP0 Register 12, Select 1)

The *IntCtl* register controls the expanded interrupt capability added in Release 2 of the Architecture, including vectored interrupts and support for an external interrupt controller. This register does not exist in implementations of Release 1 of the Architecture.

Figure 8.23 shows the format of the *IntCtl* register; Table 8.25 describes the *IntCtl* register fields.

Figure 8.23 IntCtl Register Format




Table 8.25 IntCtl Register Field Descriptions

Fields		Description	Read/Write	Reset State																					
Name	Bits																								
IPTI	31:29	<p>For Interrupt Compatibility and Vectored Interrupt modes, this field specifies the IP number to which the Timer Interrupt request is merged, and allows software to determine whether to consider <i>Cause_{TI}</i> for a potential interrupt.</p> <table border="1" style="width: 100%; border-collapse: collapse; margin: 10px 0;"> <thead> <tr> <th style="text-align: center;">Encoding</th> <th style="text-align: center;">IP bit</th> <th style="text-align: center;">Hardware Interrupt Source</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">2</td> <td style="text-align: center;">2</td> <td style="text-align: center;">HW0</td> </tr> <tr> <td style="text-align: center;">3</td> <td style="text-align: center;">3</td> <td style="text-align: center;">HW1</td> </tr> <tr> <td style="text-align: center;">4</td> <td style="text-align: center;">4</td> <td style="text-align: center;">HW2</td> </tr> <tr> <td style="text-align: center;">5</td> <td style="text-align: center;">5</td> <td style="text-align: center;">HW3</td> </tr> <tr> <td style="text-align: center;">6</td> <td style="text-align: center;">6</td> <td style="text-align: center;">HW4</td> </tr> <tr> <td style="text-align: center;">7</td> <td style="text-align: center;">7</td> <td style="text-align: center;">HW5</td> </tr> </tbody> </table> <p>The value of this bit is set by the static input, <i>SI_IPTI[2:0]</i>. This allows external logic to communicate the specific <i>SI_Int</i> hardware interrupt pin to which the timer interrupt is attached internally.</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>	Encoding	IP bit	Hardware Interrupt Source	2	2	HW0	3	3	HW1	4	4	HW2	5	5	HW3	6	6	HW4	7	7	HW5	R	Externally Set
Encoding	IP bit	Hardware Interrupt Source																							
2	2	HW0																							
3	3	HW1																							
4	4	HW2																							
5	5	HW3																							
6	6	HW4																							
7	7	HW5																							

Table 8.25 IntCtl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State																					
Name	Bits																								
IPPCI	28:26	<p>For Interrupt Compatibility and Vectored Interrupt modes, this field specifies the IP number to which the Performance Counter Interrupt request is merged, and allows software to determine whether to consider <i>Cause_{PCI}</i> for a potential interrupt.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>IP bit</th> <th>Hardware Interrupt Source</th> </tr> </thead> <tbody> <tr> <td>2</td> <td>2</td> <td>HW0</td> </tr> <tr> <td>3</td> <td>3</td> <td>HW1</td> </tr> <tr> <td>4</td> <td>4</td> <td>HW2</td> </tr> <tr> <td>5</td> <td>5</td> <td>HW3</td> </tr> <tr> <td>6</td> <td>6</td> <td>HW4</td> </tr> <tr> <td>7</td> <td>7</td> <td>HW5</td> </tr> </tbody> </table> <p>The value of this bit is set by the static input, <i>SI_IPPCI[2:0]</i>. This allows external logic to communicate the specific <i>SI_Int</i> hardware interrupt pin to which the performance counter interrupt is attached internally. 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>	Encoding	IP bit	Hardware Interrupt Source	2	2	HW0	3	3	HW1	4	4	HW2	5	5	HW3	6	6	HW4	7	7	HW5	R	0
Encoding	IP bit	Hardware Interrupt Source																							
2	2	HW0																							
3	3	HW1																							
4	4	HW2																							
5	5	HW3																							
6	6	HW4																							
7	7	HW5																							

Table 8.25 IntCtl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State																					
Name	Bits																								
IPFDC	25:23	<p>For Interrupt Compatibility and Vectored Interrupt 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 <i>Cause_{FDC}</i> for a potential interrupt.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>IP bit</th> <th>Hardware Interrupt Source</th> </tr> </thead> <tbody> <tr> <td>2</td> <td>2</td> <td>HW0</td> </tr> <tr> <td>3</td> <td>3</td> <td>HW1</td> </tr> <tr> <td>4</td> <td>4</td> <td>HW2</td> </tr> <tr> <td>5</td> <td>5</td> <td>HW3</td> </tr> <tr> <td>6</td> <td>6</td> <td>HW4</td> </tr> <tr> <td>7</td> <td>7</td> <td>HW5</td> </tr> </tbody> </table> <p>The value of this field is UNPREDICTABLE if External Interrupt Controller Mode is both implemented and enabled. The external interrupt controller is expected to provide this information for that interrupt mode. If EJTAG FDC is not implemented, this field returns zero on read.</p> <p>The value of this bit is set by the static input, SI_IPFDCI[2:0]. This allows external logic to communicate the specific SI_Int hardware interrupt pin to which the Fast Debug Channel interrupt is attached internally.</p>	Encoding	IP bit	Hardware Interrupt Source	2	2	HW0	3	3	HW1	4	4	HW2	5	5	HW3	6	6	HW4	7	7	HW5	R	Preset or Externally Set
Encoding	IP bit	Hardware Interrupt Source																							
2	2	HW0																							
3	3	HW1																							
4	4	HW2																							
5	5	HW3																							
6	6	HW4																							
7	7	HW5																							
PF	22	<p>Enables Vector Prefetching Feature.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Vector Prefetching disabled.</td> </tr> <tr> <td>1</td> <td>Vector Prefetching enabled.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Vector Prefetching disabled.	1	Vector Prefetching enabled.	RW	0															
Encoding	Meaning																								
0	Vector Prefetching disabled.																								
1	Vector Prefetching enabled.																								
ICE	21	<p>For IRET instruction. Enables Interrupt Chaining.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt Chaining disabled</td> </tr> <tr> <td>1</td> <td>Interrupt Chaining enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	Interrupt Chaining disabled	1	Interrupt Chaining enabled	RW	0															
Encoding	Meaning																								
0	Interrupt Chaining disabled																								
1	Interrupt Chaining enabled																								

Table 8.25 IntCtl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State									
Name	Bits												
StkDec	20:16	<p>For Auto-Prologue feature. This is the number of 4-byte words that is decremented from the value of GPR29</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Decrement Amount in Words</th> <th>Decrement Amount in Bytes</th> </tr> </thead> <tbody> <tr> <td>0-3</td> <td>3</td> <td>12</td> </tr> <tr> <td>Others</td> <td>As encoded, e.g. 0x5 means 5 words</td> <td>4 * encoded value e.g. 0x5 means 20 bytes</td> </tr> </tbody> </table>	Encoding	Decrement Amount in Words	Decrement Amount in Bytes	0-3	3	12	Others	As encoded, e.g. 0x5 means 5 words	4 * encoded value e.g. 0x5 means 20 bytes	RW	0x3
Encoding	Decrement Amount in Words	Decrement Amount in Bytes											
0-3	3	12											
Others	As encoded, e.g. 0x5 means 5 words	4 * encoded value e.g. 0x5 means 20 bytes											
ClrEXL	15	<p>For Auto-Prologue feature and IRET instruction. If set, during Auto-Prologue and IRET interrupt chaining, the <i>KSU/ERL/EXL</i> fields are cleared.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Fields are not cleared by these operations.</td> </tr> <tr> <td>1</td> <td>Fields are cleared by these operations.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Fields are not cleared by these operations.	1	Fields are cleared by these operations.	RW	0			
Encoding	Meaning												
0	Fields are not cleared by these operations.												
1	Fields are cleared by these operations.												
APE	14	<p>Enables Auto-Prologue feature.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Auto-Prologue disabled</td> </tr> <tr> <td>1</td> <td>Auto-Prologue enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	Auto-Prologue disabled	1	Auto-Prologue enabled	RW	0			
Encoding	Meaning												
0	Auto-Prologue disabled												
1	Auto-Prologue enabled												
UseKStk	13	<p>Chooses which Stack to use during Interrupt Automated Prologue.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Copy \$29 of the Previous SRS to the Current SRS at the beginning of IAP. This is used for Bare-Iron environments with only one stack.</td> </tr> <tr> <td>1</td> <td>Use \$29 of the Current SRS at the beginning of IAP. This is used for environments where there are separate User-mode and Kernel mode stacks. In this case, \$29 of the SRS used during IAP must be pre-initialized by software to hold the Kernel mode stack pointer.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Copy \$29 of the Previous SRS to the Current SRS at the beginning of IAP. This is used for Bare-Iron environments with only one stack.	1	Use \$29 of the Current SRS at the beginning of IAP. This is used for environments where there are separate User-mode and Kernel mode stacks. In this case, \$29 of the SRS used during IAP must be pre-initialized by software to hold the Kernel mode stack pointer.	RW	0			
Encoding	Meaning												
0	Copy \$29 of the Previous SRS to the Current SRS at the beginning of IAP. This is used for Bare-Iron environments with only one stack.												
1	Use \$29 of the Current SRS at the beginning of IAP. This is used for environments where there are separate User-mode and Kernel mode stacks. In this case, \$29 of the SRS used during IAP must be pre-initialized by software to hold the Kernel mode stack pointer.												
0	12:10	Must be written as zero; returns zero on read.	0	0									

Table 8.25 IntCtl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State																					
Name	Bits																								
VS	9:5	Vector Spacing. If vectored interrupts are implemented (as denoted by <i>Config3_{VInt}</i> or <i>Config3_{VEIC}</i>), this field specifies the spacing between vectored interrupts. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Spacing Between Vectors (hex)</th> <th>Spacing Between Vectors (decimal)</th> </tr> </thead> <tbody> <tr> <td>16#00</td> <td>16#000</td> <td>0</td> </tr> <tr> <td>16#01</td> <td>16#020</td> <td>32</td> </tr> <tr> <td>16#02</td> <td>16#040</td> <td>64</td> </tr> <tr> <td>16#04</td> <td>16#080</td> <td>128</td> </tr> <tr> <td>16#08</td> <td>16#100</td> <td>256</td> </tr> <tr> <td>16#10</td> <td>16#200</td> <td>512</td> </tr> </tbody> </table> All other values are reserved. The operation of the processor is UNDEFINED if a reserved value is written to this field.	Encoding	Spacing Between Vectors (hex)	Spacing Between Vectors (decimal)	16#00	16#000	0	16#01	16#020	32	16#02	16#040	64	16#04	16#080	128	16#08	16#100	256	16#10	16#200	512	R/W	0
Encoding	Spacing Between Vectors (hex)	Spacing Between Vectors (decimal)																							
16#00	16#000	0																							
16#01	16#020	32																							
16#02	16#040	64																							
16#04	16#080	128																							
16#08	16#100	256																							
16#10	16#200	512																							
0	4:0	Must be written as zero; returns zero on read.	0	0																					

8.2.22 SRSCtl Register (CP0 Register 12, Select 2)

The *SRSCtl* register controls the operation of GPR shadow sets in the processor. This register does not exist in implementations of the architecture prior to Release 2.

Figure 8.24 shows the format of the *SRSCtl* register; Table 8.26 describes the *SRSCtl* register fields.

Figure 8.24 SRSCtl Register Format




Table 8.26 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 8.26 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 M5100 processor are:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>One shadow set (normal GPR set) is present.</td> </tr> <tr> <td>1</td> <td>Two shadow sets are present.</td> </tr> <tr> <td>3</td> <td>Four shadow sets are present.</td> </tr> <tr> <td>7</td> <td>Eight shadow sets are present</td> </tr> <tr> <td>15</td> <td>Sixteen shadow sets are present</td> </tr> <tr> <td>2, 4-6, 8-14</td> <td>Reserved</td> </tr> </tbody> </table> <p>The value in this field also represents the highest value that can be written to the <i>ESS</i>, <i>EICSS</i>, <i>PSS</i>, and <i>CSS</i> fields of this register, or to any of the fields of the <i>SRSMap</i> register. The operation of the processor is UNDEFINED if a value larger than the one in this field is written to any of these other fields.</p>	Encoding	Meaning	0	One shadow set (normal GPR set) is present.	1	Two shadow sets are present.	3	Four shadow sets are present.	7	Eight shadow sets are present	15	Sixteen shadow sets are present	2, 4-6, 8-14	Reserved	R	Preset
Encoding	Meaning																	
0	One shadow set (normal GPR set) is present.																	
1	Two shadow sets are present.																	
3	Four shadow sets are present.																	
7	Eight shadow sets are present																	
15	Sixteen shadow sets are present																	
2, 4-6, 8-14	Reserved																	
0	25:22	Must be written as zeros; returns zero on read.	0	0														
EICSS	21:18	<p>EIC interrupt mode shadow set. If <i>Config3_{VEIC}</i> 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 <i>SRSMap</i> register to select the current shadow set for the interrupt.</p> <p>See Section 7.3.1 “Interrupt Modes” for a discussion of EIC interrupt mode. If <i>Config3_{VEIC}</i> is 0, this field must be written as zero, and 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 UNDEFINED if software writes a value into this field that is greater than the value in the <i>HSS</i> field.</p>	R/W	0														
0	11:10	Must be written as zeros; returns zero on read.	0	0														

Table 8.26 SRSCtl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State
Name	Bits			
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 <i>CSS</i> field when an exception or interrupt occurs. An ERET instruction copies this value back into the <i>CSS</i> field if <i>Status_{BEV}</i> = 0.</p> <p>This field is not updated on any exception which sets <i>Status_{ERL}</i> to 1 (i.e., Reset, Soft Reset, NMI, cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with <i>Status_{EXL}</i> = 1, or <i>Status_{BEV}</i> = 1. This field is not updated on an exception that occurs while <i>Status_{ERL}</i> = 1.</p> <p>The operation of the processor is UNDEFINED if software writes a value into this field that is greater than the value in the <i>HSS</i> field.</p>	R/W	0
0	5:4	Must be written as zeros; returns zero on read.	0	0
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 <i>PSS</i> field on an ERET. Table 8.27 describes the various sources from which the <i>CSS</i> field is updated on an exception or interrupt.</p> <p>This field is not updated on any exception which sets <i>Status_{ERL}</i> to 1 (i.e., Reset, Soft Reset, NMI, cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with <i>Status_{EXL}</i> = 1, or <i>Status_{BEV}</i> = 1. Neither is it updated on an ERET with <i>Status_{ERL}</i> = 1 or <i>Status_{BEV}</i> = 1. This field is not updated on an exception that occurs while <i>Status_{ERL}</i> = 1.</p> <p>The value of <i>CSS</i> can be changed directly by software only by writing the <i>PSS</i> field and executing an ERET instruction.</p>	R	0

8.2.23 SRSSMap Register (CP0 Register 12, Select 3)

Table 8.27 Sources for new SRSSCtl_{CSS} on an Exception or Interrupt

Exception Type	Condition	SRSSCtl _{CSS} Source	Comment
Exception	All	SRSSCtl _{ESS}	
Non-Vectored Interrupt	$Cause_{IV} = 0$	SRSSCtl _{ESS}	Treat as exception
Vectored Interrupt	$Cause_{IV} = 1$ and $Config3_{VEIC} = 0$ and $Config3_{VInt} = 1$	SRSSMap _{VECTNUM}	Source is internal map register. (for VECTNUM see Table 7.3)
Vectored EIC Interrupt	$Cause_{IV} = 1$ and $Config3_{VEIC} = 1$	SRSSCtl _{EICSS}	Source is external interrupt controller.

The *SRSSMap* register contains 8, 4-bit fields that provide the mapping from a 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 ($Cause_{IV} = 0$ or $IntCtl_{VS} = 0$). In such cases, the shadow set number comes from *SRSSCtl_{ESS}*.

If *SRSSCtl_{HSS}* 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 *SRSSCtl_{HSS}*.

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 8.25](#) shows the format of the *SRSSMap* register; [Table 8.28](#) describes the *SRSSMap* register fields.

Figure 8.25 SRSSMap Register Format




Table 8.28 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

8.2.24 View_IPL Register (CP0 Register 12, Select 4)

Figure 8-26 View_IPL Register Format




Table 8.29 View_IPL Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
IM	9:0	Interrupt Mask. If EIC interrupt mode is not enabled, controls which interrupts are enabled.	R/W	Undefined for IM7:IM2 0 for IM9:IM8
IPL	9:2	Interrupt Priority Level. If EIC interrupt mode is enabled, this field is the encoded value of the current <i>IPL</i> .	R/W	Undefined
0	31:10,1:0	Must be written as zero; returns zero on read.	0	0

This register gives read and write access to the *IM* or *IPL* field that is also available in the *Status* Register. The use of this register allows the Interrupt Mask or the Priority Level to be read/written without extracting/inserting that bit field from/to the *Status* Register.

The *IPL* field might be located in non-contiguous bits within the *Status* Register. All of the *IPL* bits are presented as a contiguous field within this register.

8.2.25 SRSSMap2 Register (CP0 Register 12, Select 5)

The *SRSSMap2* register contains 2 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-vectorized interrupt ($Cause_{IV} = 0$ or $IntCtl_{VS} = 0$). In such cases, the shadow set number comes from *SRSCtl_{ESS}*.

If *SRSCtl_{HSS}* 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_{HSS}*.

The *SRSSMap2* register contains the shadow register set numbers for vector numbers 9:8. 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 8-27 shows the format of the *SRSSMap2* register; Table 8.30 describes the *SRSSMap2* register fields.

Figure 8-27 SRSSMap Register Format



Table 8.30 SRSSMap Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:8	Must be written as zero; returns zero on read.	R	0
SSV9	7:4	Shadow register set number for Vector Number 9	R/W	0
SSV8	3:0	Shadow register set number for Vector Number 8	R/W	0

8.2.26 GuestCtl0 Register (CP0 Register 12, Select 6)

The *GuestCtl0* register contains control bits that indicate whether the base mode of the processor is guest mode or root mode, plus additional bits controlling guest mode access to privileged resources. The *GuestCtl0* register is accessible only in root mode.

The *GuestCtl0* register is instantiated per-VPE in an MT Module processor. This register is added by the Virtualization Module.

Note on behavior of *GuestCtl0_{DRG/RAD}*: These R/W fields define additional functions for the Guest and Root TLBs. Both must be interpreted together. An implementation does not have to support all valid combinations. Root software can test supported combinations by writing then reading legal values. Legal values for $(RAD, DRG) = \{00, 01, 11\}$.

Figure 8.28 shows the format of the Virtualization Module *GuestCtl0* register; Table 8.31 describes the *GuestCtl0* register fields.

Figure 8.28 GuestCtl0 Register Format




Table 8.31 GuestCtl0 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
GM	31	<p>Guest Mode</p> <p>The processor is in guest mode when $GM=1$, $Root.Status_{EXL}=0$ and $Root.Status_{ERL}=0$ and $Root.Debug_{DM}=0$.</p>	R/W	0						
RI	30	<p>Guest Reserved Instruction Redirect.</p> <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>Reserved Instruction exceptions during guest-mode execution are taken in guest mode.</td> </tr> <tr> <td>1</td> <td>Reserved Instruction exceptions during guest-mode execution result in a Guest Reserved Instruction Redirect exception, taken in root mode.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Reserved Instruction exceptions during guest-mode execution are taken in guest mode.	1	Reserved Instruction exceptions during guest-mode execution result in a Guest Reserved Instruction Redirect exception, taken in root mode.	R/W	0
Encoding	Meaning									
0	Reserved Instruction exceptions during guest-mode execution are taken in guest mode.									
1	Reserved Instruction exceptions during guest-mode execution result in a Guest Reserved Instruction Redirect exception, taken in root mode.									
MC	29	<p>Guest Mode-Change exception enable. The purpose of this enable is to provide Root software control over certain mode-changing events within guest context that may be frequent in guest context by causing Field Change exceptions.</p> <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> <p>During guest mode execution a hardware initiated change to $Guest.Status_{EXL}$ will not trigger a Guest Hardware Field Change Exception.</p> <p>During guest mode execution, a software initiated change to $Guest.Status_{UM/KSU}$ will not trigger a Guest Software Field Change Exception.</p> </td> </tr> <tr> <td>1</td> <td> <p>During guest mode execution a hardware initiated change to $Guest.Status_{EXL}$ will trigger a Guest Hardware Field Change Exception.</p> <p>During guest mode execution, a software initiated change to $Guest.Status_{UM/KSU}$ will trigger a Guest Software Field Change Exception.</p> </td> </tr> </tbody> </table>	Encoding	Meaning	0	<p>During guest mode execution a hardware initiated change to $Guest.Status_{EXL}$ will not trigger a Guest Hardware Field Change Exception.</p> <p>During guest mode execution, a software initiated change to $Guest.Status_{UM/KSU}$ will not trigger a Guest Software Field Change Exception.</p>	1	<p>During guest mode execution a hardware initiated change to $Guest.Status_{EXL}$ will trigger a Guest Hardware Field Change Exception.</p> <p>During guest mode execution, a software initiated change to $Guest.Status_{UM/KSU}$ will trigger a Guest Software Field Change Exception.</p>	R/W	0
Encoding	Meaning									
0	<p>During guest mode execution a hardware initiated change to $Guest.Status_{EXL}$ will not trigger a Guest Hardware Field Change Exception.</p> <p>During guest mode execution, a software initiated change to $Guest.Status_{UM/KSU}$ will not trigger a Guest Software Field Change Exception.</p>									
1	<p>During guest mode execution a hardware initiated change to $Guest.Status_{EXL}$ will trigger a Guest Hardware Field Change Exception.</p> <p>During guest mode execution, a software initiated change to $Guest.Status_{UM/KSU}$ will trigger a Guest Software Field Change Exception.</p>									

Table 8.31 GuestCtl0 Register Field Descriptions

Fields		Description	Read / Write	Reset State										
Name	Bits													
CP0	28	<p>Guest access to coprocessor 0.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Guest-kernel use of a GuestPrivileged Sensitive instruction will trigger a Guest Privileged Sensitive Instruction exception.</td> </tr> <tr> <td>1</td> <td>Guest-kernel use of a GuestPrivileged Sensitive instructions is permitted, subject to all other exception conditions.</td> </tr> </tbody> </table> <p>The CP0 bit has no other effect on the operation of Coprocessor 0 in guest mode.</p>	Encoding	Meaning	0	Guest-kernel use of a GuestPrivileged Sensitive instruction will trigger a Guest Privileged Sensitive Instruction exception.	1	Guest-kernel use of a GuestPrivileged Sensitive instructions is permitted, subject to all other exception conditions.	R/W	0				
Encoding	Meaning													
0	Guest-kernel use of a GuestPrivileged Sensitive instruction will trigger a Guest Privileged Sensitive Instruction exception.													
1	Guest-kernel use of a GuestPrivileged Sensitive instructions is permitted, subject to all other exception conditions.													
AT	27:26	<p>Guest Address Translation control.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Reserved.</td> </tr> <tr> <td>1</td> <td>Guest MMU under Root control. Guest and Root MMU both implemented and active in hardware. This mode is optional.</td> </tr> <tr> <td>2</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Guest MMU under Guest control. Guest and Root MMU both implemented and active in hardware. This mode is required.</td> </tr> </tbody> </table> <p>Guest TLB resources are:</p> <ul style="list-style-type: none"> • TLB related Instructions - TLBWR, TLBWI, TLBR, TLBP, TLBINV, TLBINVF. • Supporting Registers - <i>Index</i>, <i>Random</i>, <i>EntryLo0</i>, <i>EntryLo1</i>, <i>EntryHi</i>, <i>Context</i>, <i>XContext</i>, <i>ContextConfig</i>, <i>PageMask</i>, <i>PageGrain</i>, <i>SegCtl0</i>, <i>SegCtl1</i>, <i>SegCtl2</i>, <i>PWBase</i>, <i>PWField</i>, <i>PWSize</i>, <i>PWCtl</i>. <p>If the Guest TLB resources (excluding <i>Index</i>, <i>Random</i>, <i>EntryLo0</i>, <i>EntryLo1</i>, <i>Context</i>, <i>XContext</i>, <i>ContextConfig</i>, <i>PageMask</i> and <i>EntryHi</i>) are under Root control ($GuestCtl0_{AT}=1$), Guest use of these instructions or access to any of these registers will trigger a Guest Privileged Sensitive Instruction exception, allowing Root to control Guest address translation directly.</p> <p>In default mode ($GuestCtl0_{AT}=3$), the Guest TLB resources are active under Guest control.</p>	Encoding	Meaning	0	Reserved.	1	Guest MMU under Root control. Guest and Root MMU both implemented and active in hardware. This mode is optional.	2	Reserved	3	Guest MMU under Guest control. Guest and Root MMU both implemented and active in hardware. This mode is required.	R or R/W if more than default mode implemented.	Implementation defined
Encoding	Meaning													
0	Reserved.													
1	Guest MMU under Root control. Guest and Root MMU both implemented and active in hardware. This mode is optional.													
2	Reserved													
3	Guest MMU under Guest control. Guest and Root MMU both implemented and active in hardware. This mode is required.													

Table 8.31 GuestCtl0 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
GT	25	<p>Guest Timer register access.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Guest-kernel access to <i>Count</i> or <i>Compare</i> registers, or a read of the <i>CC</i> field in the <i>HWREna</i> register with the RDHWR instruction will trigger a Guest Privileged Sensitive Instruction exception.</td> </tr> <tr> <td>1</td> <td>Guest-kernel read access from <i>Count</i> and guest-kernel read or write access to <i>Compare</i> is permitted, and guest reads from <i>CC</i> using RDHWR instruction are permitted in any mode.</td> </tr> </tbody> </table> <p>The <i>GT</i> bit has no other effect on the operation of timers in guest mode.</p>	Encoding	Meaning	0	Guest-kernel access to <i>Count</i> or <i>Compare</i> registers, or a read of the <i>CC</i> field in the <i>HWREna</i> register with the RDHWR instruction will trigger a Guest Privileged Sensitive Instruction exception.	1	Guest-kernel read access from <i>Count</i> and guest-kernel read or write access to <i>Compare</i> is permitted, and guest reads from <i>CC</i> using RDHWR instruction are permitted in any mode.	R/W	0
Encoding	Meaning									
0	Guest-kernel access to <i>Count</i> or <i>Compare</i> registers, or a read of the <i>CC</i> field in the <i>HWREna</i> register with the RDHWR instruction will trigger a Guest Privileged Sensitive Instruction exception.									
1	Guest-kernel read access from <i>Count</i> and guest-kernel read or write access to <i>Compare</i> is permitted, and guest reads from <i>CC</i> using RDHWR instruction are permitted in any mode.									
CG	24	<p>Cache Instruction Guest-mode enable. When read-only, the GPSI exception will always occur.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Use of the CACHE instruction will cause a Guest Privileged Sensitive Instruction exception to occur.</td> </tr> <tr> <td>1</td> <td>Use of the CACHE instruction is permitted only when used with an Effective Address Operand type of 'Address'. All other operand types will cause a a Guest Privileged Sensitive Instruction exception.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Use of the CACHE instruction will cause a Guest Privileged Sensitive Instruction exception to occur.	1	Use of the CACHE instruction is permitted only when used with an Effective Address Operand type of 'Address'. All other operand types will cause a a Guest Privileged Sensitive Instruction exception.	R0, R/W	0
Encoding	Meaning									
0	Use of the CACHE instruction will cause a Guest Privileged Sensitive Instruction exception to occur.									
1	Use of the CACHE instruction is permitted only when used with an Effective Address Operand type of 'Address'. All other operand types will cause a a Guest Privileged Sensitive Instruction exception.									
CF	23	<p>Config register access.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Guest-kernel write access to <i>Config0-7</i> will trigger a Guest Privileged Sensitive Instruction exception.</td> </tr> <tr> <td>1</td> <td>Guest-kernel access to <i>Config0-7</i> is permitted.</td> </tr> </tbody> </table> <p>The CF bit has no other effect on the operation of <i>Config</i> register fields in guest mode.</p>	Encoding	Meaning	0	Guest-kernel write access to <i>Config0-7</i> will trigger a Guest Privileged Sensitive Instruction exception.	1	Guest-kernel access to <i>Config0-7</i> is permitted.	R/W	0
Encoding	Meaning									
0	Guest-kernel write access to <i>Config0-7</i> will trigger a Guest Privileged Sensitive Instruction exception.									
1	Guest-kernel access to <i>Config0-7</i> is permitted.									
G1	22	<p><i>GuestCtl1</i> register implementation. Set by hardware.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>GuestCTL1</i> register is not implemented.</td> </tr> <tr> <td>1</td> <td><i>GuestCTL1</i> register is implemented.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>GuestCTL1</i> register is not implemented.	1	<i>GuestCTL1</i> register is implemented.	R	preset
Encoding	Meaning									
0	<i>GuestCTL1</i> register is not implemented.									
1	<i>GuestCTL1</i> register is implemented.									

Table 8.31 GuestCtl0 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
Impl	21:20	<p>Implementation defined.</p> <p>These bits are not defined by the MIPS32 architecture and are implementation-dependent.</p> <p>If these bits are not implemented, they must be ignored on writes and must read as zero.</p> <p>If they are implemented and they modify the behavior of the processor, they must be defined in such a way that preserves correct behavior if software, with no knowledge of these bits, reads the <i>GuestCtl0</i> register, modifies another field, and writes the updated value back to the <i>GuestCtl0</i> register.</p>	R/W	0						
G0E	19	<p><i>GuestCtl0Ext</i> register implemented. Set by hardware.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>GuestCtl0Ext</i> register is not implemented.</td> </tr> <tr> <td>1</td> <td><i>GuestCtl0Ext</i> register is implemented.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>GuestCtl0Ext</i> register is not implemented.	1	<i>GuestCtl0Ext</i> register is implemented.	R	preset
Encoding	Meaning									
0	<i>GuestCtl0Ext</i> register is not implemented.									
1	<i>GuestCtl0Ext</i> register is implemented.									
PT	18	<p>Defines the existence of the Pending Interrupt Pass-through feature.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>GuestCtl0_PIP</i> is not supported. <i>GuestCtl0_PIP</i> is a reserved field. All external interrupts are processed via Root intervention.</td> </tr> <tr> <td>1</td> <td><i>GuestCtl0_PIP</i> supported. Interrupts may be assigned to Root or Guest.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>GuestCtl0_PIP</i> is not supported. <i>GuestCtl0_PIP</i> is a reserved field. All external interrupts are processed via Root intervention.	1	<i>GuestCtl0_PIP</i> supported. Interrupts may be assigned to Root or Guest.	R	preset
Encoding	Meaning									
0	<i>GuestCtl0_PIP</i> is not supported. <i>GuestCtl0_PIP</i> is a reserved field. All external interrupts are processed via Root intervention.									
1	<i>GuestCtl0_PIP</i> supported. Interrupts may be assigned to Root or Guest.									
PIP	17:10	<p>Pending Interrupt Pass-through (PIP).</p> <p>In non-EIC mode, controls how external interrupts are passed through to the guest CP0 context. Interpreted as a bit mask and applies 1:1 to <i>Guest.Cause_IP[9:2]</i>.</p> <p><i>GuestCtl1_PIP</i> may be extended by <i>GuestCtl1_ASE</i>.</p> <p>Existence of the PIP feature is defined by the <i>GuestCtl0_PT</i> field.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Corresponding interrupt request is not visible in guest context.</td> </tr> <tr> <td>1</td> <td>Corresponding interrupt request is visible in guest context.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Corresponding interrupt request is not visible in guest context.	1	Corresponding interrupt request is visible in guest context.	R/W R0 if unimplemented	0
Encoding	Meaning									
0	Corresponding interrupt request is not visible in guest context.									
1	Corresponding interrupt request is visible in guest context.									

Table 8.31 GuestCtl0 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
RAD	9	<p>Root ASID De-alias mode. Determines the means by which a Virtualized MMU implementation uses Root ASID to de-alias different contexts.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GuestID is used to de-alias both Guest and Root TLB entries.</td> </tr> <tr> <td>1</td> <td>Root ASID is used to de-alias Root TLB entries when Guest TLB contains only one context at any given time.</td> </tr> </tbody> </table>	Encoding	Meaning	0	GuestID is used to de-alias both Guest and Root TLB entries.	1	Root ASID is used to de-alias Root TLB entries when Guest TLB contains only one context at any given time.	R	0
Encoding	Meaning									
0	GuestID is used to de-alias both Guest and Root TLB entries.									
1	Root ASID is used to de-alias Root TLB entries when Guest TLB contains only one context at any given time.									
DRG	8	<p>Direct Root to Guest access. Determines whether an implementation provides root software the means to access guest entries directly in the Root TLB for access to guest memory. If $GuestCtl0_{DRG}=1$ and $GuestCtl0_{RID}$ is non-zero, then all root accesses are mapped. H/W must set G=1 as if the access were for guest.</p> <p><i>DRG</i> is R0 if only <i>DRG</i>=0 supported; otherwise it must be R/W.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Root software cannot access guest entries directly.</td> </tr> <tr> <td>1</td> <td>Root software can access guest entries directly.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Root software cannot access guest entries directly.	1	Root software can access guest entries directly.	R	0
Encoding	Meaning									
0	Root software cannot access guest entries directly.									
1	Root software can access guest entries directly.									
G2	7	<p><i>GuestCtl2</i> register implemented. Set by hardware.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Unimplemented</td> </tr> <tr> <td>1</td> <td>Implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	Unimplemented	1	Implemented	R	preset
Encoding	Meaning									
0	Unimplemented									
1	Implemented									

Table 8.31 GuestCtl0 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
GExc-Code	6:2	Hypervisor exception cause code. Described in Table 8.32 . This field is UNDEFINED on a root exception.	R	Undefined						
SFC2	1	<p>Guest Software Field Change exception enable for <i>Guest.Status_{CU2}</i>. The purpose of this enable is to provide Root software control over guest COP2 enable related Field Change exception. Guest software may utilize <i>Status_{CU2}</i> for COP2 specific context switching.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GSFC exception taken if <i>Status_{CU2}</i> is modified by guest.</td> </tr> <tr> <td>1</td> <td>GSFC exception not taken if <i>Status_{CU2}</i> modified by guest.</td> </tr> </tbody> </table>	Encoding	Meaning	0	GSFC exception taken if <i>Status_{CU2}</i> is modified by guest.	1	GSFC exception not taken if <i>Status_{CU2}</i> modified by guest.	R/W if implemented, 0 otherwise	0
Encoding	Meaning									
0	GSFC exception taken if <i>Status_{CU2}</i> is modified by guest.									
1	GSFC exception not taken if <i>Status_{CU2}</i> modified by guest.									
SFC1	0	<p>Guest Software Field Change exception enable for <i>Guest.Status_{CU1}</i>. The purpose of this enable is to provide Root software control over guest COP1 enable related Field Change exception. Guest software may utilize <i>Status_{CU1}</i> for COP1 specific context switching.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>GSFC exception taken if <i>Status_{CU1}</i> is modified by guest.</td> </tr> <tr> <td>1</td> <td>GSFC exception not taken if <i>Status_{CU1}</i> modified by guest.</td> </tr> </tbody> </table>	Encoding	Meaning	0	GSFC exception taken if <i>Status_{CU1}</i> is modified by guest.	1	GSFC exception not taken if <i>Status_{CU1}</i> modified by guest.	R/W if implemented, 0 otherwise.	0
Encoding	Meaning									
0	GSFC exception taken if <i>Status_{CU1}</i> is modified by guest.									
1	GSFC exception not taken if <i>Status_{CU1}</i> modified by guest.									

[Table 8.32](#) describes the cause codes use for GExcCode.

Table 8.32 GuestCtl0 GExcCode values

Exception code value		Mnemonic	Description
Decimal	Hexadecimal		
0	0x00	GPSI	Guest Privileged Sensitive instruction. Taken when execution of a Guest Privileged Sensitive Instruction was attempted from guest-kernel mode, but the instruction was not enabled for guest-kernel mode.
1	0x01	GSFC	Guest Software Field Change event
2	0x02	HC	Hypercall
3	0x03	GRR	Guest Reserved Instruction Redirect. A Reserved Instruction exception would be taken in guest mode. When <i>GuestCtl0_{RI}</i> =1, this root-mode exception is raised before the guest-mode exception can be taken.
4 - 8	0x4 - 0x8	IMP	N.A.
9	0x09	GHFC	Guest Hardware Field Change event

Table 8.32 GuestCtl0 GExcCode values

Exception code value		Mnemonic	Description
Decimal	Hexadecimal		
10	0x0A	GPA	Guest mode initiated Root TLB exception has Guest Physical Address available. Set when a Guest mode initiated TLB translation results in a Root TLB related exception occurring in Root mode and the Guest Physical Address is available.
11 - 31	0xB - 0x1f	-	Reserved

8.2.27 GTOffset Register (CP0 Register 12, Select 7)

Compliance Level: *Required* by the Virtualization Module.

Timekeeping within the guest context is controlled by root mode. The guest time value is generated by adding the two’s complement offset in the *Root.GTOffset* register to the root timer in value *Root.Count*.

The guest time value is used to generate timer interrupts within the guest context, by comparison with the *Guest.Compare* register. The guest time value can be read from the *Guest.Count* register. Guest writes to the *Guest.Count* register always result in a Guest Privileged Sensitive Instruction exception.

The number of bits supported in *GTOffset* is implementation-dependent but must be non-zero. It is recommended that a minimum of 16 bits be implemented. Root software can check the number of implemented bits by writing all ones and then reading. Unimplemented bits will return zero.

The *GTOffset* register is instantiated per-VPE in an MT Module processor. This register is added by the Virtualization Module.

See [Section 4.6.8 “Guest Timer”](#) of MD00846 “Virtualization Module for the MIPS32 Architecture”.

[Figure 8.29](#) shows the Virtualization Module format of the *GTOffset* register; [Table 8.33](#) describes the *GTOffset* register fields.

Figure 8.29 GTOffset Register Format



Table 8.33 GTOffset Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
GTOffset	31:0	Two’s complement offset from <i>Root.Count</i>	R/W	0

8.2.28 Cause Register (CP0 Register 13, Select 0)

The *Cause* register primarily describes the cause of the most recent exception. In addition, fields also control software interrupt requests and the vector through which interrupts are dispatched. With the exception of the *IP1..0*, *DC*,

Table 8.34 Cause Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
PCI	26	<p>Performance Counter Interrupt. In an implementation of Release 2 of the Architecture, this bit denotes whether a performance counter interrupt is pending (analogous to the IP bits for other interrupt types):</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No timer interrupt is pending</td> </tr> <tr> <td>1</td> <td>Timer interrupt is pending</td> </tr> </tbody> </table> <p>The state of the <i>PCI</i> bit is available on the external M5100 interface as the <i>SI_PCInt</i> signal.</p>	Encoding	Meaning	0	No timer interrupt is pending	1	Timer interrupt is pending	R	0
Encoding	Meaning									
0	No timer interrupt is pending									
1	Timer interrupt is pending									
IC	25	<p>Indicates if Interrupt Chaining occurred on the last IRET instruction.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt Chaining did not happen on last IRET</td> </tr> <tr> <td>1</td> <td>Interrupt Chaining occurred during last IRET</td> </tr> </tbody> </table>	Encoding	Meaning	0	Interrupt Chaining did not happen on last IRET	1	Interrupt Chaining occurred during last IRET	R	Undefined
Encoding	Meaning									
0	Interrupt Chaining did not happen on last IRET									
1	Interrupt Chaining occurred during last IRET									
AP	24	<p>Indicates whether an exception occurred during Interrupt Auto-Prologue.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Exception did not occur during Auto-Prologue operation.</td> </tr> <tr> <td>1</td> <td>Exception occurred during Auto-Prologue operation.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Exception did not occur during Auto-Prologue operation.	1	Exception occurred during Auto-Prologue operation.	R	Undefined
Encoding	Meaning									
0	Exception did not occur during Auto-Prologue operation.									
1	Exception occurred during Auto-Prologue operation.									
IV	23	<p>Indicates whether an interrupt exception uses the general exception vector or a special interrupt vector:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Use the general exception vector (16#180)</td> </tr> <tr> <td>1</td> <td>Use the special interrupt vector (16#200)</td> </tr> </tbody> </table> <p>In implementations of Release 2 of the architecture, if the <i>Cause_{IV}</i> is 1 and <i>Status_{BEV}</i> is 0, the special interrupt vector represents the base of the vectored interrupt table.</p>	Encoding	Meaning	0	Use the general exception vector (16#180)	1	Use the special interrupt vector (16#200)	R/W	Undefined
Encoding	Meaning									
0	Use the general exception vector (16#180)									
1	Use the special interrupt vector (16#200)									

Table 8.34 Cause Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
WP	22	<p>Indicates that a watch exception was deferred because <i>Status_{EXL}</i> or <i>Status_{ERL}</i> had a value of 1 at the time the watch exception was detected. This bit indicates that the watch exception was deferred, and it causes the exception to be initiated when <i>Status_{EXL}</i> and <i>Status_{ERL}</i> are both zero. As such, software must clear this bit as part of the watch exception handler to prevent a watch exception loop.</p> <p>Software should not write a 1 to this bit when its value is 0, thereby causing a 0-to-1 transition. If such a transition is caused by software, it is UNPREDICTABLE whether hardware ignores the write, accepts the write with no side effects, or accepts the write and initiates a watch exception when <i>Status_{EXL}</i> and <i>Status_{ERL}</i> are both zero.</p> <p>Because watch registers are not implemented on the M5100 core, this bit is ignored on writes and reads as zero.</p>	R	0						
FDCI	21	<p>Fast Debug Channel Interrupt. This bit denotes whether a FDC Interrupt is pending (analogous to the <i>IP</i> bits for other interrupt types):</p> <table border="1" data-bbox="505 999 1029 1192"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No Fast Debug Channel interrupt is pending</td> </tr> <tr> <td>1</td> <td>Fast Debug Channel interrupt is pending</td> </tr> </tbody> </table>	Encoding	Meaning	0	No Fast Debug Channel interrupt is pending	1	Fast Debug Channel interrupt is pending	R	Undefined
Encoding	Meaning									
0	No Fast Debug Channel interrupt is pending									
1	Fast Debug Channel interrupt is pending									

Table 8.34 Cause Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State																											
Name	Bits																														
IP9:IP2	17:10	<p>Indicates an interrupt is pending:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Name</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>17</td> <td>IP9</td> <td>Hardware Interrupt 7</td> </tr> <tr> <td>16</td> <td>IP8</td> <td>Hardware Interrupt 6</td> </tr> <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>In implementations of Release 1 of the Architecture, timer and performance counter interrupts are combined in an implementation-dependent way with hardware interrupt 5.</p> <p>In implementations of Release 2 of the Architecture in which EIC interrupt mode is not enabled (<i>Config3_{VEIC}</i> = 0), timer and performance counter interrupts are combined in an implementation-dependent way with any hardware interrupt. If EIC interrupt mode is enabled (<i>Config3_{VEIC}</i> = 1), these bits have a different meaning, and are interpreted as the <i>RIPL</i> field, described below.</p>	Bit	Name	Meaning	17	IP9	Hardware Interrupt 7	16	IP8	Hardware Interrupt 6	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	<p>Undefined for IP7:IP2</p> <p>0 for IP9:IP8</p>
Bit	Name	Meaning																													
17	IP9	Hardware Interrupt 7																													
16	IP8	Hardware Interrupt 6																													
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																													
RIPL	17:10	<p>Requested Interrupt Priority Level.</p> <p>In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled (<i>Config3_{VEIC}</i> = 1), this field is the encoded (0..255) value of the requested interrupt. A value of zero indicates that no interrupt is requested.</p> <p>If EIC interrupt mode is not enabled (<i>Config3_{VEIC}</i> = 0), these bits have a different meaning and are interpreted as the <i>IP7..IP2</i> bits, described above.</p>	R	<p>Undefined for bits 15:10</p> <p>0 for bits 17:16</p>																											
IP1:IP0	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 is available on the external core interface as the <i>SI_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																													
ExcCode	6:2	Exception code - see Table 8.35	R	Undefined																											
0	20:18, 7, 1:0	Must be written as zero; returns zero on read.	0	0																											

Table 8.35 Cause Register ExcCode Field

Exception Code Value		Mnemonic	Description
Decimal	Hexadecimal		
0	16#00	Int	Interrupt
1-3	16#01-16#03	-	Reserved
4	16#04	AdEL	Address error exception (load or instruction fetch)
5	16#05	AdES	Address error exception (store)
6	16#06	IBE	Bus error exception (instruction fetch)
7	16#07	DBE	Bus error exception (data reference: load or store)
8	16#08	Sys	Syscall exception
9	16#09	Bp	Breakpoint exception
10	16#0a	RI	Reserved instruction exception
11	16#0b	CpU	Coprocessor Unusable exception
12	16#0c	Ov	Arithmetic Overflow exception
13	16#0d	Tr	Trap exception
14	16#0e	-	Reserved
15	16#0f	FPE	Floating-point exception
16	16#10	IS1	Implementation-Specific Exception 1 (COP2)
17	16#11	CEU	CorExtend Unusable
18	16#12	C2E	Coprocessor 2 exceptions
19-25	16#13-16#19	-	Reserved
26	16#1a	DSPDis	DSP Module State Disabled exception
27	16#1b	GE	Hypervisor Exception (Guest Exit). Hypervisor-intervention exception occurred during guest code execution. <i>GuestCtl0</i> _{ExcCode} contains additional cause information.
28-29	16#1c-16#1d	-	Reserved
30	16#1e	Parity Error	Parity error. In normal mode, a parity error exception has a dedicated vector and the <i>Cause</i> register is not updated. If a parity error occurs while in Debug Mode, this code is written to the <i>DebugDExcCode</i> field to indicate that re-entry to Debug Mode was caused by a parity error.
31	16#1f	-	Reserved

8.2.29 View_RIPL Register (CP0 Register 13, Select 4)

This register gives read access to the *IP* or *RIPL* field that is also available in the *Cause* Register. The use of this register allows the Interrupt Pending or the Requested Priority Level to be read without extracting that bit field from the *Cause* Register.

Figure 8-31 View_RIPL Register Format




Table 8.36 View_RIPL Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:10	Must be written as zero; returns zero on read.	0	0
IP9:IP2	9:2	HW Interrupt Pending. If EIC interrupt mode is not enabled, indicates which HW interrupts are pending.	R	Undefined for IP7:IP2 0 for IP9:IP8
RIPL	9:2	Interrupt Priority Level. If EIC interrupt mode is enabled, this field indicates the Requested Priority Level of the pending interrupt.	R	Undefined
IP1:IP0	1:0	SW Interrupt Pending. If EIC interrupt mode is not enabled, controls which SW interrupts are pending.	R/W	Undefined

8.2.30 NestedExc (CP0 Register 13, Select 5)

The *Nested Exception (NestedExc)* register is an optional read-only register containing the values of *Status_{EXL}* and *Status_{ERL}* prior to acceptance of the current exception.

This register is part of the Nested Fault feature. The existence of the register can be determined by reading the *Config5_{NFExists}* bit.

Figure 8-32 shows the format of the *NestedExc* register; Table 8.37 describes the *NestedExc* register fields.

Figure 8-32 NestedExc Register Format



Table 8.37 NestedExc Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:3	Reserved, read as 0.	R0	0
ERL	2	Value of <i>Status_{ERL}</i> prior to acceptance of current exception. Updated by all exceptions that would set either <i>Status_{EXL}</i> or <i>Status_{ERL}</i> . Not updated by Debug exceptions.	R	Undefined
EXL	1	Value of <i>Status_{EXL}</i> prior to acceptance of current exception. Updated by exceptions which would update EPC if <i>Status_{EXL}</i> is not set (MCheck, Interrupt, Address Error, all TLB exceptions, Bus Error, CopUnusable, Reserved Instruction, Overflow, Trap, Syscall, FPU, etc.) . For these exception types, this register field is updated regardless of the value of <i>Status_{EXL}</i> . Not updated by exception types which update <i>ErrorEPC</i> - (Reset, Soft Reset, NMI, Cache Error). Not updated by Debug exceptions.	R	Undefined
0	0	Reserved, read as 0.	R0	0

8.2.31 Exception Program Counter (CP0 Register 14, Select 0)

The *Exception Program Counter (EPC)* is a read/write register that contains the address at which processing resumes after an exception has been serviced. All bits of the *EPC* register are significant and must be writable.

For synchronous (precise) exceptions, the *EPC* contains one of the following:

- The virtual address of the instruction that was the direct cause of the exception
- 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.

On new exceptions, 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 microMIPS, a read of the *EPC* register (via MFC0) returns the following value in the destination GPR:

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

That is, the upper 31 bits of the exception PC are combined with the lower bit of the *ISAMode* field 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:

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

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.

Figure 8.33 EPC Register Format




Table 8.38 EPC Register Field Description

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

8.2.32 NestedEPC (CP0 Register 14, Select 2)

The *Nested Exception Program Counter (NestedEPC)* is an optional read/write register with the same behavior as the *EPC* register, except that:

- The *NestedEPC* register ignores the value of *Status_{EXL}* and is therefore updated on the occurrence of any exception, including nested exceptions.
- The *NestedEPC* register is not used by the *ERET/DERET/IRET* instructions. To return to the address stored in *NestedEPC*, software must copy the value of the *NestedEPC* register to the *EPC* register.

This register is part of the Nested Fault feature. The existence of the register can be determined by reading the *Config5NFExists* bit.

Figure 8-32 shows the format of the *NestedEPC* register; Table 8.37 describes the *NestedEPC* register fields.

Figure 8-34 NestedEPC Register Format




Table 8.39 NestedEPC Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
NestedEPC	0	<p>Nested Exception Program Counter</p> <p>Updated by exceptions which would update EPC if <i>Status_{EXL}</i> is not set (MCheck, Interrupt, Address Error, all TLB exceptions, Bus Error, CopUnusable, Reserved Instruction, Overflow, Trap, Syscall, FPU, etc.) . For these exception types, this register field is updated regardless of the value of <i>Status_{EXL}</i>.</p> <p>Not updated by exception types which update <i>ErrorEPC</i> i.e., Reset, Soft Reset, NMI, and Cache Error. Not updated by Debug exceptions.</p>	R/W	Undefined

8.2.33 Processor Identification (CP0 Register 15, Select 0)

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

Figure 8.35 PRId Register Format




Table 8.40 PRId Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
Company Opt	31:24	Company Option. Whatever name is specified by the SoC builder who synthesizes the M5100 core; refer to your SoC manual. This field should be preset by the config GUI with a number between 0x00 and 0x7F; higher values (0x80-0xFF) are reserved by MIPS Technologies.	R	Preset
Company ID	23:16	Company Identifier. Identifies the company that designed or manufactured the processor. In the M5100 this field contains a value of 1 to indicate MIPS Technologies, Inc.	R	1
Processor ID	15:8	Processor Identifier. Identifies the type of processor. This field allows software to distinguish between the various types of MIPS Technologies processors.	R	0xA6

Table 8.40 PRId Register Field Descriptions (Continued)

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

8.2.34 EBase Register (CP0 Register 15, Select 1)

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

The *EBase* register provides the ability for software to identify the specific processor within a multiprocessor system, and allows the exception vectors for each processor to be different, especially in systems composed of heterogeneous processors. Bits 31:12 of the *EBase* register are concatenated with zeros to form the base of the exception vectors when *Status_{BEV}* is 0. The exception vector base address comes from the fixed defaults (see [Section 7.5 “Exception Vector Locations”](#)) when *Status_{BEV}* 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 16#8000.0000, providing backward compatibility with Release 1 implementations.

Bits 31:30 of the *EBase* Register are fixed with the value 2#10 to force the exception base address to be in the kseg0 or kseg1 unmapped virtual address segments.

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

Combining bits 31:12 with the *Exception Base* field allows the base address of the exception vectors to be placed at any 4KByte page boundary. If vectored interrupts are used, a vector offset greater than 4KBytes can be generated. In

this case, bit 12 of the *Exception Base* field must be zero. The operation of the processor is **UNDEFINED** if software writes bit 12 of the *Exception Base* field with a 1 and enables the use of a vectored interrupt whose offset is greater than 4KBytes from the exception base address.

Figure 8.36 shows the format of the *EBase* Register; Table 8.41 describes the *EBase* register fields.

Figure 8.36 EBase Register Format




Table 8.41 EBase Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
1	31	This bit is ignored on writes and returns one on reads.	R	1
0	30	This bit is ignored on writes and returns zero on reads.	R	0
Exception Base	29:12	In conjunction with bits 31:30, this field specifies the base address of the exception vectors when <i>Status_{BEV}</i> is zero.	R/W	0
0	11:10	Must be written as zero; returns zero on reads.	0	0
CPUNum	9:0	This field specifies the number of the CPU in a multiprocessor system and can be used by software to distinguish a particular processor from the others. The value in this field is set by the <i>SI_CPUNum[9:0]</i> static input pins to the core. In a single processor system, this value should be set to zero.	R	Externally Set

8.2.35 CDMMBase Register (CP0 Register 15, Select 2)

The 36-bit physical base address for the Common Device Memory Map facility is defined by this register. This register only exists if *Config3_{CDMM}* is set to one.

Figure 8.37 shows the format of the *CDMMBase* register, and Table 8.42 describes the register fields.

Figure 8.37 CDMMBase Register Format




Table 8.42 CDMMBase Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
CDMM_UPPER_ADDR	31:11	Bits 35:15 of the base physical address of the memory mapped registers. The number of implemented physical address bits is implementation-specific. For the unimplemented address bits, writes are ignored and reads return zero.	R/W	Undefined

Table 8.42 CDMMBase Register Field Descriptions

Fields		Description	Read/Write	Reset State												
Name	Bits															
EN	10	Enables the CDMM region. If this bit is cleared, memory requests to this address region access regular system memory. If this bit is set, memory requests to this region access the CDMM logic	R/W	0												
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>CDMM Region is disabled.</td> </tr> <tr> <td>1</td> <td>CDMM Region is enabled.</td> </tr> </tbody> </table>	Encoding	Meaning	0	CDMM Region is disabled.	1	CDMM Region is enabled.								
Encoding	Meaning															
0	CDMM Region is disabled.															
1	CDMM Region is enabled.															
CI	9	If set to 1, this indicates that the first 64-byte Device Register Block of the CDMM is reserved for additional registers that manage CDMM region behavior and are not IO device registers.	R	0												
CDMMSize	8:0	This field represents the number of 64-byte Device Register Blocks instantiated in the core.	R	Preset												
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>1 DRB</td> </tr> <tr> <td>1</td> <td>2 DRBs</td> </tr> <tr> <td>2</td> <td>3 DRBs</td> </tr> <tr> <td>...</td> <td>...</td> </tr> <tr> <td>511</td> <td>512 DRBs</td> </tr> </tbody> </table>	Encoding	Meaning	0	1 DRB	1	2 DRBs	2	3 DRBs	511	512 DRBs		
Encoding	Meaning															
0	1 DRB															
1	2 DRBs															
2	3 DRBs															
...	...															
511	512 DRBs															

8.2.36 Config Register (CP0 Register 16, Select 0)

The *Config* register specifies various configuration and capabilities information. Most of the fields in the *Config* register are initialized by hardware during the Reset exception process, or are constant. Figure 8.38 shows the format of the Config Register Format - Select 0, and Table 8.43 describes the register fields.

Figure 8.38 Config Register Format — Select 0




Table 8.43 Config Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
M	31	This bit is hardwired to ‘1’ to indicate the presence of the Config1 register.	R	1
K23	30:28	This field controls the cacheability of the kseg2 and kseg3 address segments in FM implementations. Refer to Table 8.44 for the field encoding.	FM: R/W	FM: 010

Table 8.43 Config Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
KU	27:25	This field controls the cacheability of the kuseg and useg address segments in FM implementations. Refer to Table 8.44 for the field encoding.	FM: R/W	FM: 010
0	24:23	Must be written as 0. Returns zero on reads.	0	0
UDI	22	This bit indicates that CorExtend User Defined Instructions have been implemented. 0 = No User Defined Instructions are implemented 1 = User Defined Instructions are implemented	R	Preset
SB	21	Indicates whether SimpleBE bus mode is enabled. Set via <i>Sl_SimpleBE[0]</i> input pin: 0 = No reserved byte enables on SRAM interface 1 = Only simple byte enables allowed on SRAM interface	R	Externally Set
MDU	20	This bit indicates the type of Multiply/Divide Unit present. 0 = Fast, high-performance MDU 1 = Iterative, area-efficient MDU	R	Preset
0	19:17	Must be written as 0. Returns zero on reads.	0	0
DS	16	Dual SRAM interface. 0: Unified instruction/data SRAM interface 1: Dual instruction/data SRAM interfaces	R	Preset
BE	15	Indicates the endian mode in which the processor is running. Set via <i>Sl_Endian</i> input pin. 0: Little endian 1: Big endian	R	Externally Set
AT	14:13	Architecture type implemented by the processor. This field is always 00 to indicate the MIPS32 architecture.	R	00
AR	12:10	Architecture revision level. This field is always 001 to indicate MIPS32 Release 2. 0: Release 1 1: Release 2 2-7: Reserved	R	001
MT	9:7	MMU Type: 3: Fixed Mapping 0-2, 4-7: Reserved	R	3
0	6:3	Must be written as zeros; returns zeros on reads.	0	0
K0	2:0	Kseg0 coherency algorithm. Refer to Table 8.44 for the field encoding.	R/W	010

Table 8.44 Cache Coherency Attributes

C(2:0) Value	Cache Coherency Attribute
2	Uncached.
3	Cached (Core treats as uncached, but passes attribute to the system for use with any external caching mechanisms)

8.2.37 Config1 Register (CP0 Register 16, Select 1)

The *Config1* register is an adjunct to the *Config* register and encodes additional information about capabilities present on the core. All fields in the *Config1* register are read-only.

Figure 8.39 Config1 Register Format — Select 1




Table 8.45 Config1 Register Field Descriptions — Select 1


Fields		Description	Read/Write	Reset State
Name	Bit(s)			
M	31	This bit is hardwired to '1' to indicate the presence of the <i>Config2</i> register.	R	1
MMU Size	30:25	This field contains the number of entries in the TLB minus one.	R	0
IS	24:22	This field contains the number of instruction cache sets per way. Because the M5100 core does not include caches, this field is always read as 0.	R	0
IL	21:19	This field contains the instruction cache line size. Because the M5100 core does not include caches, this field is always read as 0.	R	0
IA	18:16	This field contains the level of instruction cache associativity. Because the M5100 core does not include caches, this field is always read as 0.	R	0
DS	15:13	This field contains the number of data cache sets per way. Because the M5100 core does not include caches, this field is always read as 0.	R	0
DL	12:10	This field contains the data cache line size. Because the M5100 core does not include caches, this field is always read as 0.	R	0
DA	9:7	This field contains the type of set associativity for the data cache. Because the M5100 core does not include caches, this field is always read as 0.	R	0
C2	6	Coprocessor 2 present. 0: No coprocessor is attached to the COP2 interface 1: A coprocessor is attached to the COP2 interface If the Cop2 interface logic is not implemented, this bit will read 0.	R	Preset
MD	5	MDMX implemented. This bit always reads as 0 because MDMX is not supported.	R	0
PC	4	Performance Counter registers implemented. 0: No Performance Counter registers are implemented 1: Performance Counter registers are implemented	R	Preset

Table 8.45 Config1 Register Field Descriptions — Select 1 (Continued)

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
WR	3	Watch registers implemented. 0: No Watch registers are implemented 1: One or more Watch registers are implemented This bit is always read as 0, because the M5100 core does not contain Watch registers.	R	0
CA	2	Code compression (MIPS16e) implemented. 0: MIPS16e is not implemented 1: MIPS16e is implemented	R	0
EP	1	EJTAG present: This bit is always set to indicate that the core implements EJTAG.	R	1
FP	0	FPU implemented. 0: No FPU 1: FPU is implemented	R	Preset

8.2.38 Config2 Register (CP0 Register 16, Select 2)

The *Config2* register is an adjunct to the *Config* register and is reserved to encode additional capabilities information. *Config2* is allocated for showing the configuration of level 2/3 caches. These fields are reset to 0 because L2/L3 caches are not supported by the M5100 core. All fields in the *Config2* register are read-only.

Figure 8.40 Config2 Register Format — Select 2**Table 8.46 Config2 Register Field Descriptions — Select 1**

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
M	31	This bit is hardwired to '1' to indicate the presence of the <i>Config3</i> register.	R	1
0	30:0	These bits are reserved.	R	0

8.2.39 Config3 Register (CP0 Register 16, Select 3)

The *Config3* register encodes additional capabilities. All fields in the *Config3* register are read-only.

Figure 8-41 shows the format of the *Config3* register; Table 8.47 describes the *Config3* register fields.

Figure 8-41 Config3 Register Format




Table 8.47 Config3 Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bits									
M	31	This bit is reserved to indicate that a <i>Config4</i> register is present.	R	1						
0	30:28,24,9,2	Must be written as zeros; returns zeros on read.	0	0						
BP	27	BadInstrP register implemented. This bit indicates whether the faulting prior branch instruction word register is present. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>BadInstrP</i> is not implemented.</td> </tr> <tr> <td>1</td> <td><i>BadInstrP</i> is implemented.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>BadInstrP</i> is not implemented.	1	<i>BadInstrP</i> is implemented.	R	1
Encoding	Meaning									
0	<i>BadInstrP</i> is not implemented.									
1	<i>BadInstrP</i> is implemented.									
BI	26	BadInstr register implemented. This bit indicates whether the faulting instruction word register is present. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>BadInstrP</i> is not implemented.</td> </tr> <tr> <td>1</td> <td><i>BadInstrP</i> is implemented.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>BadInstrP</i> is not implemented.	1	<i>BadInstrP</i> is implemented.	R	1
Encoding	Meaning									
0	<i>BadInstrP</i> is not implemented.									
1	<i>BadInstrP</i> is implemented.									
VZ	23	Virtualization Module implemented. This bit indicates whether the Virtualization Module is present. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Virtualization Module is not implemented.</td> </tr> <tr> <td>1</td> <td>Virtualization Module is implemented.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Virtualization Module is not implemented.	1	Virtualization Module is implemented.	R	Preset (Always 0 in Guest context)
Encoding	Meaning									
0	Virtualization Module is not implemented.									
1	Virtualization Module is implemented.									

Table 8.47 Config3 Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State								
Name	Bits											
IPLW	22:21	<p>Width of the <i>Status</i>_{IPL} and <i>Cause</i>_{RIPL} fields:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>IPL</i> and <i>RIPL</i> fields are 6-bits in width.</td> </tr> <tr> <td>1</td> <td><i>IPL</i> and <i>RIPL</i> fields are 8-bits in width.</td> </tr> <tr> <td>Others</td> <td>Reserved.</td> </tr> </tbody> </table> <p>If the <i>IPL</i> field is 8-bits in width, bits 18 and 16 of <i>Status</i> are used as the most significant bit and second most significant bit, respectively, of that field.</p> <p>If the <i>RIPL</i> field is 8-bits in width, bits 17 and 16 of <i>Cause</i> are used as the most significant bit and second most significant bit, respectively, of that field.</p>	Encoding	Meaning	0	<i>IPL</i> and <i>RIPL</i> fields are 6-bits in width.	1	<i>IPL</i> and <i>RIPL</i> fields are 8-bits in width.	Others	Reserved.	R	1
Encoding	Meaning											
0	<i>IPL</i> and <i>RIPL</i> fields are 6-bits in width.											
1	<i>IPL</i> and <i>RIPL</i> fields are 8-bits in width.											
Others	Reserved.											
MMAR	20:18	<p>microMIPS Architecture revision level:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Release 1</td> </tr> <tr> <td>1-7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	0	Release 1	1-7	Reserved	R	0		
Encoding	Meaning											
0	Release 1											
1-7	Reserved											
MCU	17	<p>MIPS MCU ASE implemented.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>MCU ASE is not implemented.</td> </tr> <tr> <td>1</td> <td>MCU ASE is implemented.</td> </tr> </tbody> </table>	Encoding	Meaning	0	MCU ASE is not implemented.	1	MCU ASE is implemented.	R	1		
Encoding	Meaning											
0	MCU ASE is not implemented.											
1	MCU ASE is implemented.											
ISAOnExc	16	<p>Reflects the Instruction Set Architecture used when vectoring to an exception. Affects exceptions whose vectors are offsets from EBASE.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>MIPS32 ISA is used on entrance to an exception vector.</td> </tr> <tr> <td>1</td> <td>microMIPS is used on entrance to an exception vector.</td> </tr> </tbody> </table>	Encoding	Meaning	0	MIPS32 ISA is used on entrance to an exception vector.	1	microMIPS is used on entrance to an exception vector.	RW	Preset, driven by signal external to CPU core		
Encoding	Meaning											
0	MIPS32 ISA is used on entrance to an exception vector.											
1	microMIPS is used on entrance to an exception vector.											

Table 8.47 Config3 Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State										
Name	Bits													
ISA	15:14	Indicates Instruction Set Availability. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Only MIPS32 is implemented.</td> </tr> <tr> <td>1</td> <td>Only microMIPS is implemented.</td> </tr> <tr> <td>2</td> <td>Both MIPS32 and microMIPS are implemented. MIPS32 ISA used when coming out of reset.</td> </tr> <tr> <td>3</td> <td>Both MIPS32 and microMIPS are implemented. microMIPS is used when coming out of reset.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Only MIPS32 is implemented.	1	Only microMIPS is implemented.	2	Both MIPS32 and microMIPS are implemented. MIPS32 ISA used when coming out of reset.	3	Both MIPS32 and microMIPS are implemented. microMIPS is used when coming out of reset.	R	Preset, driven by signal external to CPU core
Encoding	Meaning													
0	Only MIPS32 is implemented.													
1	Only microMIPS is implemented.													
2	Both MIPS32 and microMIPS are implemented. MIPS32 ISA used when coming out of reset.													
3	Both MIPS32 and microMIPS are implemented. microMIPS is used when coming out of reset.													
ULRI	13	<i>UserLocal</i> register implemented. This bit indicates whether the <i>UserLocal</i> coprocessor 0 register is implemented. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>UserLocal</i> register is not implemented</td> </tr> <tr> <td>1</td> <td><i>UserLocal</i> register is implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>UserLocal</i> register is not implemented	1	<i>UserLocal</i> register is implemented	R	1				
Encoding	Meaning													
0	<i>UserLocal</i> register is not implemented													
1	<i>UserLocal</i> register is implemented													
RXI	12	Indicates whether the <i>RIE</i> and <i>XIE</i> bits exist within the <i>PageGrain</i> register.. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>The <i>RIE</i> and <i>XIE</i> bits are not implemented within the <i>PageGrain</i> register.</td> </tr> <tr> <td>1</td> <td>The <i>RIE</i> and <i>XIE</i> bits are implemented within the <i>PageGrain</i> register</td> </tr> </tbody> </table>	Encoding	Meaning	0	The <i>RIE</i> and <i>XIE</i> bits are not implemented within the <i>PageGrain</i> register.	1	The <i>RIE</i> and <i>XIE</i> bits are implemented within the <i>PageGrain</i> register	R	0				
Encoding	Meaning													
0	The <i>RIE</i> and <i>XIE</i> bits are not implemented within the <i>PageGrain</i> register.													
1	The <i>RIE</i> and <i>XIE</i> bits are implemented within the <i>PageGrain</i> register													
DSP2P	11	Reads 1 to indicate that Revision 2 of the MIPS DSP Module is implemented	R	Preset										
DSPP	10	Reads 1 to indicate that the MIPS DSP Module extension is implemented.	R	Preset										
ITL	8	Indicates that iFlowtrace hardware is present.	R	Preset										

Table 8.47 Config3 Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
LPA	7	<p>Denotes the presence of support for large physical addresses on MIPS64 processors. Not used by MIPS32 processors and returns zero on read.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Large physical address support is not implemented</td> </tr> <tr> <td>1</td> <td>Large physical address support is implemented</td> </tr> </tbody> </table> <p>For implementations of Release 1 of the Architecture, this bit returns zero on read.</p>	Encoding	Meaning	0	Large physical address support is not implemented	1	Large physical address support is implemented	R	0
Encoding	Meaning									
0	Large physical address support is not implemented									
1	Large physical address support is implemented									
VEIC	6	<p>Indicates support for an external interrupt controller.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Support for EIC interrupt mode is not implemented</td> </tr> <tr> <td>1</td> <td>Support for EIC interrupt mode is implemented</td> </tr> </tbody> </table> <p>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 or not.</p>	Encoding	Meaning	0	Support for EIC interrupt mode is not implemented	1	Support for EIC interrupt mode is implemented	R	Externally Set
Encoding	Meaning									
0	Support for EIC interrupt mode is not implemented									
1	Support for EIC interrupt mode is implemented									
VInt	5	<p>Indicates implementation of Vectored interrupts.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Vector interrupts are not implemented</td> </tr> <tr> <td>1</td> <td>Vectored interrupts are implemented</td> </tr> </tbody> </table> <p>On the M5100 core, this bit is always a 1, because vectored interrupts are implemented.</p>	Encoding	Meaning	0	Vector interrupts are not implemented	1	Vectored interrupts are implemented	R	1
Encoding	Meaning									
0	Vector interrupts are not implemented									
1	Vectored interrupts are implemented									
SP	4	<p>When set, indicates that Small (1KByte) page support is implemented.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Small page support is not implemented</td> </tr> <tr> <td>1</td> <td>Small page support is implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	Small page support is not implemented	1	Small page support is implemented	R	0
Encoding	Meaning									
0	Small page support is not implemented									
1	Small page support is implemented									
CDMM	3	<p>Common Device Memory Map implemented. This bit indicates whether the CDMM is implemented.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>CDMM is not implemented</td> </tr> <tr> <td>1</td> <td>CDMM is implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	CDMM is not implemented	1	CDMM is implemented	R	Preset
Encoding	Meaning									
0	CDMM is not implemented									
1	CDMM is implemented									

Table 8.47 Config3 Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
SM	1	SmartMIPS™ ASE implemented. This bit indicates whether the SmartMIPS ASE is implemented. Because SmartMIPS is not present on the M5100 core, this bit will always be 0. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>SmartMIPS ASE is not implemented</td> </tr> <tr> <td>1</td> <td>SmartMIPS ASE is implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	SmartMIPS ASE is not implemented	1	SmartMIPS ASE is implemented	R	0
Encoding	Meaning									
0	SmartMIPS ASE is not implemented									
1	SmartMIPS ASE is implemented									
TL	0	Trace Logic implemented. This bit indicates whether PC or data trace is implemented. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Trace logic is not implemented</td> </tr> <tr> <td>1</td> <td>Trace logic is implemented</td> </tr> </tbody> </table>	Encoding	Meaning	0	Trace logic is not implemented	1	Trace logic is implemented	R	Preset
Encoding	Meaning									
0	Trace logic is not implemented									
1	Trace logic is implemented									

8.2.40 Config4 Register (CP0 Register 16, Select 4)

The *Config4* register encodes additional capabilities. This register is required if any optional feature described by this register is implemented and is otherwise optional.

Figure 8-42 shows the format of the *Config4* register; Table 8.48 describes the *Config4* register fields.

Figure 8-42 Config4 Register Format




Table 8.48 Config4 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
M	31	This bit is reserved to indicate that a <i>Config5</i> register is present.	R	1

Table 8.48 Config4 Register Field Descriptions

Fields		Description	Read / Write	Reset State										
Name	Bits													
IE	30:29	<p>TLB Invalidate instruction support/configuration.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>TLBINV, TLBINVF, <i>EntryHi</i>_{EHINV} not supported by hardware</td> </tr> <tr> <td>01</td> <td>TLBINV, TLBINVF not supported. <i>EntryHi</i>_{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions.</td> </tr> <tr> <td>10</td> <td>TLBINV, TLBINVF supported. <i>EntryHi</i>_{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions. TLBINV* instructions operate on one TLB entry</td> </tr> <tr> <td>11</td> <td>TLBINV, TLBINVF supported. <i>EntryHi</i>_{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions. TLBINV* instructions operate on entire TLB.</td> </tr> </tbody> </table>	Encoding	Meaning	00	TLBINV, TLBINVF, <i>EntryHi</i> _{EHINV} not supported by hardware	01	TLBINV, TLBINVF not supported. <i>EntryHi</i> _{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions.	10	TLBINV, TLBINVF supported. <i>EntryHi</i> _{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions. TLBINV* instructions operate on one TLB entry	11	TLBINV, TLBINVF supported. <i>EntryHi</i> _{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions. TLBINV* instructions operate on entire TLB.	R	0x1
Encoding	Meaning													
00	TLBINV, TLBINVF, <i>EntryHi</i> _{EHINV} not supported by hardware													
01	TLBINV, TLBINVF not supported. <i>EntryHi</i> _{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions.													
10	TLBINV, TLBINVF supported. <i>EntryHi</i> _{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions. TLBINV* instructions operate on one TLB entry													
11	TLBINV, TLBINVF supported. <i>EntryHi</i> _{EHINV} supported. Refer to Volume II of MIPS32® Architecture for Programmers for the full description of these instructions. TLBINV* instructions operate on entire TLB.													
KScr Exist	23:16	<p>Indicates how many scratch registers are available to kernel-mode software within COPO Register 31.</p> <p>Each bit represents a select for Coprocessor0 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.</p> <p>Scratch registers meant for other purposes are not represented in this field. For example, if EJTAG is implemented, bit 16 is preset to zero even though DESAVE register is implemented as Select 0. Select 1 is reserved for future debug purposes and should not be used as a kernel scratch register, so bit 17 is preset to zero.</p>	R	0x0C										
0	28:24,15:0	Must be written as zeros; returns zeros on read.	0	0										

8.2.41 Config5 Register (CP0 Register 16, Select 5)

The *Config5* register encodes additional capabilities. This register is required if any optional feature described by this register is implemented and is otherwise optional.

Figure 8-43 shows the format of the *Config5* register; Table 8.49 describes the *Config5* register fields.

Figure 8-43 Config5 Register Format




Table 8.49 Config5 Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
M	31	This bit is reserved. With the current architectural definition, this bit should always read as a 0.	R	0						
0	30:3,1	Must be written as zeros; returns zeros on read.	0	0						
UFR	2	Release 5 feature. This feature allows user mode access to <i>Status_{FR}</i> with <i>CTC1</i> and <i>CFC1</i> . <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>User mode FR instructions not allowed</td> </tr> <tr> <td style="text-align: center;">1</td> <td>User mode FR instructions allowed</td> </tr> </tbody> </table>	Encoding	Meaning	0	User mode FR instructions not allowed	1	User mode FR instructions allowed	R/W	0
Encoding	Meaning									
0	User mode FR instructions not allowed									
1	User mode FR instructions allowed									
NFExists	0	Indicates that the Nested Fault feature is present. The Nested Fault feature allows recognition of faulting behavior within an exception handler.	R	1						

8.2.42 Config7 Register (CP0 Register 16, Select 7)

The *Config7* register contains implementation specific configuration information. A number of these bits are writeable to disable certain performance enhancing features within the M5100 core.

Figure 8.44 Config7 Register Format




Table 8.50 Config7 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
WII	31	Wait IE Ignore. Indicates that this processor will allow an interrupt to unblock a WAIT instruction, even if <i>IE</i> is preventing the interrupt from being taken. This avoids problems using the WAIT instruction for ‘bottom half’ interrupt servicing. In WII mode when <i>IE</i> =0, waking up from Sleep mode will not enter an Interrupt Service Routine.	R	1
0	30:19,17:0	These bits are unused and should be written as 0.	R	0
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	0

8.2.43 SecurityCtrl Register (CP0 Register 22, Select 0)

This register controls the security features in the core: random slip generation . In security critical code, these options may be enabled to reduce the possibility of an attacker being able to reveal code or data being run on the processor. The *SecurityCtrl* register also controls whether *SecurityHwPRNG* or *SecurityScrambling* is used to write the upper or lower half of a 64-bit value.

The core is able to generate slips at random intervals. This will make the timing and power profile of an algorithm indeterministic, thereby garbling information that is essential for side channel attacks. Random slips are enabled by setting the *RSI* field to a nonzero value. The average interval between the slips is controlled by the value in the *RSI* field. Thus, the average number of clock cycles between each slip is 2^n , where $n=RSI$. The average frequency can thereby be varied from 1/128 to 1/2. A random slip lasts one clock cycle. However, a new slip may be issued immediately after a preceding slip. Random slips are superimposed on other slips, which means that if a random slip is issued at the same time as any other slip or stall,, then the random slip will not impose any further delay.

The *SEL* field functions as an address selector for data written to *SecurityHwPRNG*. For *SecurityHwPRNG*, only *SEL*[0] is used. If it is set, the upper half of the 64-bit value seed value will be written; otherwise the lower half will be written.

Figure 8.45 SecurityCtrl Register Format




Table 8.51 SecurityCtrl Register Field Descriptions

Fields		Description	Read/Wri te	Reset State
Name	Bit(s)			
0	31:3	Must be written as zeros; returns zeros on reads.	0	0
SEL	23:16	Address selector for data written to <i>SecurityHwPRNG</i> .	R/W	0
RSI	2:0	Random slip interval. This field determines the average interval between random slips. A slip is generated whenever the n lower bits of the random generator value equals zero where n is the number given by <i>RSI</i> . The average frequency of slips per clock cycle is therefore $1/2^n$ when $n > 0$. Setting RSI to zero will disable random slips.	R/W	0

8.2.44 SecuritySwPRNG Register (CP0 Register 22, Select 1)

The software pseudo random number generator provides 32-bit numbers that can be used by software. The generator is accessed through the *SecuritySwPRNG* register. A value written to the register will be used as a seed value by the generator. Each time the register is read, the next number in a pseudo random number sequence will be returned. The generator may use up to 4 clock cycles to compute each number. The generator will start computing the next number as soon as the previous has been read; however, if two numbers are read within 4 clock cycles, then the second number will be incomplete.

Figure 8.46 SecuritySwPRNG Register Format




Table 8.52 SecuritySwPRNG Register Field Descriptions

Fields		Description	Read/Wri te	Reset State
Name	Bit(s)			
Random	31:0	Each read returns the next number in a pseudo random number sequence. A value written to the register will be used as a seed value by the PRNG.	R/W	Undefined

8.2.45 SecurityHwPRNG Register (CP0 Register 22, Select 2)

A value written to this register will function as a seed value for the pseudo random number generator, which is used for random slip generation. As the PRNG uses a 64-bit seed value, the *HL* field in the *SecurityCtrl* register is used to select whether the upper or lower 32 bits should be written. Before writing *SecurityHwPRNG*, the PRNG should be stopped by clearing the *RSI* fields in *SecurityCtrl*, to ensure that the correct 64-bit seed value is written.

Figure 8.47 SecurityHwPRNG Register Format




Table 8.53 SecurityHwPRNG Register Field Descriptions

Fields		Description	Read/Wri te	Reset State
Name	Bit(s)			
Seed	31:0	Seed value. A value written to this field will be used as either the upper or lower 32 bits of the 64-bit seed value used by the hardware pseudo random generator. The <i>SEL[0]</i> field in register <i>SecurityCtrl</i> determines which part of the seed value is written. Reads from this register return an undefined value which does not depend on previous writes to the register.	W	Undefined

8.2.46 Debug Register (CP0 Register 23, Select 0)

The *Debug* register is used to control the debug exception and provide information about the cause of the debug exception and also when re-entering at the debug exception vector due to a normal exception in debug mode. The read-only information bits are updated every time the debug exception is taken, or when a normal exception is taken when already in debug mode.


Only the *DM* bit and the *EJTAGver* field are valid when read from non-debug mode; the values of all other bits and fields are UNPREDICTABLE. Operation of the processor is UNDEFINED if the *Debug* register is written in non-debug mode.

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*, *DIBImpr*, *DDBLImpr*, *DDBSImpr* 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 8.48 Debug Register Format



18	17	15	14	10	9	8	7	6	5	4	3	2	1	0
DDB LImpr	Ver	DExcCode			NoSSt	SSt	R	DIBI mpr	DINT	DIB	DDBS	DDBL	DBp	DSS

Table 8.54 Debug Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
DBD	31	Indicates whether the last debug exception or exception in debug mode occurred in a branch delay slot: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Not in delay slot</td> </tr> <tr> <td>1</td> <td>In delay slot</td> </tr> </tbody> </table>	Encoding	Meaning	0	Not in delay slot	1	In delay slot	R	Undefined
Encoding	Meaning									
0	Not in delay slot									
1	In delay slot									
DM	30	Indicates that the processor is operating in debug mode: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Processor is operating in non-debug mode</td> </tr> <tr> <td>1</td> <td>Processor is operating in debug mode</td> </tr> </tbody> </table>	Encoding	Meaning	0	Processor is operating in non-debug mode	1	Processor is operating in debug mode	R	0
Encoding	Meaning									
0	Processor is operating in non-debug mode									
1	Processor is operating in debug mode									
NoDCR	29	Indicates whether the dseg memory segment is present and the Debug Control Register is accessible: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>dseg is present</td> </tr> <tr> <td>1</td> <td>No dseg present</td> </tr> </tbody> </table>	Encoding	Meaning	0	dseg is present	1	No dseg present	R	0
Encoding	Meaning									
0	dseg is present									
1	No dseg present									
LSNM	28	Controls access of load/store between dseg and main memory: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Load/stores in dseg address range goes to dseg</td> </tr> <tr> <td>1</td> <td>Load/stores in dseg address range goes to main memory</td> </tr> </tbody> </table>	Encoding	Meaning	0	Load/stores in dseg address range goes to dseg	1	Load/stores in dseg address range goes to main memory	R/W	0
Encoding	Meaning									
0	Load/stores in dseg address range goes to dseg									
1	Load/stores in dseg address range goes to main memory									
Doze	27	Indicates that the processor was in any kind of low power mode when a debug exception occurred: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Processor not in low-power mode when debug exception occurred</td> </tr> <tr> <td>1</td> <td>Processor in low-power mode when debug exception occurred</td> </tr> </tbody> </table>	Encoding	Meaning	0	Processor not in low-power mode when debug exception occurred	1	Processor in low-power mode when debug exception occurred	R	Undefined
Encoding	Meaning									
0	Processor not in low-power mode when debug exception occurred									
1	Processor in low-power mode when debug exception occurred									

Table 8.54 Debug Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
Halt	26	Indicates that the internal system bus clock was stopped when the debug exception occurred: <table border="1" data-bbox="527 415 1053 548"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Internal system bus clock stopped</td> </tr> <tr> <td>1</td> <td>Internal system bus clock running</td> </tr> </tbody> </table>	Encoding	Meaning	0	Internal system bus clock stopped	1	Internal system bus clock running	R	Undefined
Encoding	Meaning									
0	Internal system bus clock stopped									
1	Internal system bus clock running									
CountDM	25	Indicates the Count register behavior in debug mode: <table border="1" data-bbox="527 611 1053 772"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Count register stopped in debug mode</td> </tr> <tr> <td>1</td> <td>Count register is running in debug mode</td> </tr> </tbody> </table>	Encoding	Meaning	0	Count register stopped in debug mode	1	Count register is running in debug mode	R/W	1
Encoding	Meaning									
0	Count register stopped in debug mode									
1	Count register is running in debug mode									
IBusEP	24	Instruction fetch Bus Error exception Pending. Set when an instruction fetch bus error event occurs, or if a 1 is written to the bit by software. Cleared when a Bus Error exception on an instruction fetch is taken by the processor, and by reset. If <i>IBusEP</i> is set when <i>IEXI</i> is cleared, a Bus Error exception on an instruction fetch is taken by the processor, and <i>IBusEP</i> is cleared.	R/W1	0						
MCheckP	23	Indicates that an imprecise Machine Check exception is pending. All Machine Check exceptions are precise on the M5100 processor, so this bit will always read as 0.	R	0						
CacheEP	22	Indicates that an imprecise Cache Error is pending. Cache Errors cannot be taken by the M5100 core, so this bit will always read as 0	R	0						
DBusEP	21	Data access Bus Error exception Pending. Covers imprecise bus errors on data access, similar to the behavior of <i>IBusEP</i> for imprecise bus errors on an instruction fetch.	R/W1	0						
IEXI	20	Imprecise Error eXception Inhibit controls exceptions taken due to imprecise error indications. Set when the processor takes a debug exception or exception in debug mode. Cleared by execution of the DERET instruction; otherwise modifiable by debug mode software. When <i>IEXI</i> 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	Indicates that an imprecise Debug Data Break Store exception was taken. Imprecise data breaks only occur on complex breakpoints.	R	Undefined						
DDBLImpr	18	Indicates that an imprecise Debug Data Break Load exception was taken. Imprecise data breaks only occur on complex breakpoints.	R	Undefined						
Ver	17:15	EJTAG version.	R	101						

Table 8.54 Debug Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
DExcCode	14:10	Indicates the cause of the latest exception in debug mode. The field is encoded as the <i>ExcCode</i> field in the <i>Cause</i> register for those normal exceptions that may occur in debug mode. Value is undefined after a debug exception.	R	Undefined						
NoSST	9	Indicates whether the single-step feature controllable by the <i>SSt</i> bit is available in this implementation: <table border="1" data-bbox="527 573 1053 709"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Single-step feature available</td> </tr> <tr> <td>1</td> <td>No single-step feature available</td> </tr> </tbody> </table>	Encoding	Meaning	0	Single-step feature available	1	No single-step feature available	R	0
Encoding	Meaning									
0	Single-step feature available									
1	No single-step feature available									
SSt	8	Controls if debug single step exception is enabled: <table border="1" data-bbox="527 768 1053 932"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug single-step exception enabled</td> </tr> <tr> <td>1</td> <td>Debug single step exception enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	No debug single-step exception enabled	1	Debug single step exception enabled	R/W	0
Encoding	Meaning									
0	No debug single-step exception enabled									
1	Debug single step exception enabled									
R	7	Reserved. Must be written as zeros; returns zeros on reads.	R	0						
DIBImpr	6	Indicates that an Imprecise debug instruction break exception occurred (due to a complex breakpoint). Cleared on exception in debug mode.	R	Undefined						
DINT	5	Indicates that a debug interrupt exception occurred. Cleared on exception in debug mode. <table border="1" data-bbox="527 1192 1053 1329"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug interrupt exception</td> </tr> <tr> <td>1</td> <td>Debug interrupt exception</td> </tr> </tbody> </table>	Encoding	Meaning	0	No debug interrupt exception	1	Debug interrupt exception	R	Undefined
Encoding	Meaning									
0	No debug interrupt exception									
1	Debug interrupt exception									
DIB	4	Indicates that a debug instruction break exception occurred. Cleared on exception in debug mode. <table border="1" data-bbox="527 1415 1053 1551"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug instruction exception</td> </tr> <tr> <td>1</td> <td>Debug instruction exception</td> </tr> </tbody> </table>	Encoding	Meaning	0	No debug instruction exception	1	Debug instruction exception	R	Undefined
Encoding	Meaning									
0	No debug instruction exception									
1	Debug instruction exception									
DDBS	3	Indicates that a debug data break exception occurred on a store. Cleared on exception in debug mode. <table border="1" data-bbox="527 1638 1053 1774"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug data exception on a store</td> </tr> <tr> <td>1</td> <td>Debug instruction exception on a store</td> </tr> </tbody> </table>	Encoding	Meaning	0	No debug data exception on a store	1	Debug instruction exception on a store	R	Undefined
Encoding	Meaning									
0	No debug data exception on a store									
1	Debug instruction exception on a store									

Table 8.54 Debug Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State		
Name	Bit(s)					
DDBL	2	Indicates that a debug data break exception occurred on a load. Cleared on exception in debug mode.	R	Undefined		
					Encoding	Meaning
					0	No debug data exception on a load
	1	Debug instruction exception on a load				
DBp	1	Indicates that a debug software breakpoint exception occurred. Cleared on exception in debug mode.	R	Undefined		
					Encoding	Meaning
					0	No debug software breakpoint exception
	1	Debug software breakpoint exception				
DSS	0	Indicates that a debug single-step exception occurred. Cleared on exception in debug mode.	R	Undefined		
					Encoding	Meaning
					0	No debug single-step exception
	1	Debug single-step exception				

8.2.47 Trace Control Register (CP0 Register 23, Select 1)

The *TraceControl* register configuration is shown below.

This register is only implemented if the EJTAG PDtrace capability is present.

Figure 8.49 TraceControl Register Format




Table 8.55 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 this software control register.	R/W	0

Table 8.55 TraceControl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State
Name	Bits			
UT	30	This bit is used to indicate the type of user-triggered trace record. A value of zero implies a usertype 1, and a value of one implies a user type 2. The actual triggering of a user trace record occurs on a write to the <i>UserTraceData</i> register.	R/W	Undefined
0	29:28	Reserved for future use; Must be written as zero; returns zero on read.	0	0
TB	27	Trace All Branch. When set to one, this tells the processor to trace the PC value for all taken branches, 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 core trace logic that slow but complete tracing is desired. When set to one, the core tracing logic does not allow a FIFO overflow or 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	When set to one, this enables tracing in Debug Mode. For trace to be enabled in Debug mode, the On bit must be one. When set to zero, trace is disabled in Debug Mode, irrespective of other bits.	R/W	Undefined
E	24	When set to one, this enables tracing in Exception Mode. For trace to be enabled in Exception mode, the On bit must be one. When set to zero, trace is disabled in Exception Mode, irrespective of other bits.	R/W	Undefined
K	23	When set to one, this enables tracing in Kernel Mode. For trace to be enabled in Kernel mode, the On bit must be one. When set to zero, trace is disabled in Kernel Mode, irrespective of other bits.	R/W	Undefined
0	22	This bit is reserved. Must be written as zero; returns zero on read.	0	0
U	21	When set to one, this enables tracing in User Mode. For trace to be enabled in User mode, the On bit must be one. When set to zero, trace is disabled in User Mode, irrespective of other bits.	R/W	Undefined
ASID_M	20:13	In an FM-based MMU core in which ASID is not supported, this field is ignored on writes and returns zero on reads.	R	0

Table 8.55 TraceControl Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State																		
Name	Bits																					
ASID	12:5	In an FM-based MMU core in which ASID is not supported, this field is ignored on writes and returns zero on reads.	R	0																		
G	4	In an FM-based MMU core in which ASID is not supported, this field is ignored on writes and returns 1 on reads. This causes all match equations to work correctly in the absence of an ASID.	R	1																		
Mode	3:1	<p>These three bits control the trace mode function.</p> <table border="1"> <thead> <tr> <th>Mode</th> <th>Trace Mode</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>Trace PC</td> </tr> <tr> <td>001</td> <td>Trace PC and load address</td> </tr> <tr> <td>010</td> <td>Trace PC and store address</td> </tr> <tr> <td>011</td> <td>Trace PC and both load/store addresses</td> </tr> <tr> <td>100</td> <td>Trace PC and load data</td> </tr> <tr> <td>101</td> <td>Trace PC and load address and data</td> </tr> <tr> <td>110</td> <td>Trace PC and store address and data</td> </tr> <tr> <td>111</td> <td>Trace PC and both load/store address and data</td> </tr> </tbody> </table> <p>The <i>TraceControl2</i>_{ValidModes} field determines which of these encodings are supported by the processor. The operation of the processor is UNPREDICTABLE if this field is set to a value which is not supported by the processor.</p>	Mode	Trace Mode	000	Trace PC	001	Trace PC and load address	010	Trace PC and store address	011	Trace PC and both load/store addresses	100	Trace PC and load data	101	Trace PC and load address and data	110	Trace PC and store address and data	111	Trace PC and both load/store address and data	R/W	Undefined
Mode	Trace Mode																					
000	Trace PC																					
001	Trace PC and load address																					
010	Trace PC and store address																					
011	Trace PC and both load/store addresses																					
100	Trace PC and load data																					
101	Trace PC and load address and data																					
110	Trace PC and store address and data																					
111	Trace PC and both load/store address and data																					
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																		

8.2.48 Trace Control2 Register (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) (see [Section 11.8.6 “ITCB Register Interface for Software Configurability”](#)). 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 EJTAG PDTrace capability is present.

Figure 8.50 TraceControl2 Register Format

31	7	6	5	4	3	2	0
0				Valid-Modes	TBI	TBU	SyP

Table 8.56 TraceControl2 Register Field Descriptions

Fields		Description	Read/Write	Reset State										
Name	Bits													
0	31:5	Reserved for future use; Must be written as zero; returns zero on read.	0	0										
ValidModes	6:5	<p>This field specifies the type of tracing that is supported by the processor.</p> <table border="1"> <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													
TBI	4	<p>This bit indicates how many trace buffers are implemented by the TCB.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Only one trace buffer is implemented, and the Debug single-step exception bit of this register indicates which trace buffer is implemented</td> </tr> <tr> <td>1</td> <td>Both on-chip and off-chip trace buffers are implemented by the TCB, and the <i>TBU</i> bit of this register indicates to which trace buffer the trace is currently written.</td> </tr> </tbody> </table>	Encoding	Meaning	0	Only one trace buffer is implemented, and the Debug single-step exception 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 <i>TBU</i> bit of this register indicates to which trace buffer the trace is currently written.	R	Per implementation				
Encoding	Meaning													
0	Only one trace buffer is implemented, and the Debug single-step exception 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 <i>TBU</i> bit of this register indicates to which trace buffer the trace is currently written.													
TBU	3	<p>This bit denotes which trace buffer is currently being written by the trace and is used to select the appropriate interpretation of the <i>TraceControl2_{SyP}</i> field.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Trace data is being sent to an on-chip trace buffer</td> </tr> <tr> <td>1</td> <td>Trace Data is being sent to an off-chip trace buffer</td> </tr> </tbody> </table>	Encoding	Meaning	0	Trace data is being sent to an on-chip trace buffer	1	Trace Data is being sent to an off-chip trace buffer	R	Undefined				
Encoding	Meaning													
0	Trace data is being sent to an on-chip trace buffer													
1	Trace Data is being sent to an off-chip trace buffer													

Table 8.56 TraceControl2 Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State																											
Name	Bits																														
SyP	2:0	<p>Used to indicate the synchronization period. The period (in cycles) between which the periodic synchronization information is to be sent is defined as shown below, for both when the trace buffer is on-chip and off-chip.</p> <table border="1"> <thead> <tr> <th>SyP</th> <th>On-chip</th> <th>Off-chip</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>2^2</td> <td>2^7</td> </tr> <tr> <td>001</td> <td>2^3</td> <td>2^8</td> </tr> <tr> <td>010</td> <td>2^4</td> <td>2^9</td> </tr> <tr> <td>011</td> <td>2^5</td> <td>2^{10}</td> </tr> <tr> <td>100</td> <td>2^6</td> <td>2^{11}</td> </tr> <tr> <td>101</td> <td>2^7</td> <td>2^{12}</td> </tr> <tr> <td>110</td> <td>2^8</td> <td>2^{13}</td> </tr> <tr> <td>111</td> <td>2^9</td> <td>2^{14}</td> </tr> </tbody> </table> <p>The “On-chip” column value is used when the trace data is being written to an on-chip trace buffer (e.g., <i>TraceControl2_{TBU}</i> = 0). Conversely, the “Off-chip” column is used when the trace data is being written to an off-chip trace buffer (e.g., <i>TraceControl2_{TBU}</i> = 1).</p>	SyP	On-chip	Off-chip	000	2^2	2^7	001	2^3	2^8	010	2^4	2^9	011	2^5	2^{10}	100	2^6	2^{11}	101	2^7	2^{12}	110	2^8	2^{13}	111	2^9	2^{14}	R	Undefined
SyP	On-chip	Off-chip																													
000	2^2	2^7																													
001	2^3	2^8																													
010	2^4	2^9																													
011	2^5	2^{10}																													
100	2^6	2^{11}																													
101	2^7	2^{12}																													
110	2^8	2^{13}																													
111	2^9	2^{14}																													

8.2.49 User Trace Data1 Register (CP0 Register 23, Select 3)/User Trace Data2 Register (CP0 Register 24, Select 3)

A software write to any bits in the *UserTraceData1* or *UserTraceData2* registers will trigger a trace record to be written indicating a type 1 or type 2 user format respectively. The trace output data is **UNPREDICTABLE** if these registers are written in consecutive cycles.

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

Figure 8.51 User Trace Data1/User Trace Data2 Register Format




Table 8.57 UserTraceData1/UserTraceData2 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

8.2.50 TraceBPC Register (CP0 Register 23, Select 4)

This register is used to start and stop tracing using an EJTAG Hardware breakpoint. The Hardware breakpoint can then be set as a trigger source and optionally also as a Debug exception breakpoint.

This register is only implemented if hardware breakpoints and the EJTAG PDTrace capability are both present.

Figure 8.52 Trace BPC Register Format




Table 8.58 TraceBPC Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bits									
DE	31	Used to specify whether the trigger signal from EJTAG data breakpoint should trigger tracing functions or not: <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr> <th style="width: 30%;">Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>Disables trigger signals from data breakpoints</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Enables trigger signals from data breakpoints</td> </tr> </tbody> </table>	Encoding	Meaning	0	Disables trigger signals from data breakpoints	1	Enables trigger signals from data breakpoints	R/W	0
Encoding	Meaning									
0	Disables trigger signals from data breakpoints									
1	Enables trigger signals from data breakpoints									
0	30:18	Reserved	0	0						
DBPOn	17:16	Each of the 2 bits corresponds to the 2 possible EJTAG hardware data breakpoints that may be implemented. For example, bit 16 corresponds to the first data breakpoint. If 2 data breakpoints are present in the EJTAG implementation, then they correspond to bits 16 and 17. The rest are always ignored by the tracing logic because they will never be triggered. A value of one for each bit implies that a trigger from the corresponding data breakpoint should start tracing. And a value of zero implies that tracing should be turned off with the trigger signal.	R/W	0						
IE	15	Used to specify whether the trigger signal from EJTAG instruction breakpoint should trigger tracing functions or not: <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr> <th style="width: 30%;">Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td>Disables trigger signals from instruction breakpoints</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Enables trigger signals from instruction breakpoints</td> </tr> </tbody> </table>	Encoding	Meaning	0	Disables trigger signals from instruction breakpoints	1	Enables trigger signals from instruction breakpoints	R/W	0
Encoding	Meaning									
0	Disables trigger signals from instruction breakpoints									
1	Enables trigger signals from instruction breakpoints									
0	14:6	Reserved	0	0						

Table 8.58 TraceBPC Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State
Name	Bits			
IBPOn	5:0	Each of the 6 bits corresponds to the 6 possible EJTAG hardware instruction breakpoints that may be implemented. Bit 0 corresponds to the first instruction breakpoint, and so on. If only 2 instruction breakpoints are present in the EJTAG implementation, then only bits 0 and 1 are used. The rest are always ignored by the tracing logic because they will never be triggered. A value of one for each bit implies that a trigger from the corresponding instruction breakpoint should start tracing. And a value of zero implies that tracing should be turned off with the trigger signal.	R/W	0

8.2.51 Debug2 Register (CP0 Register 23, Select 6)

This register holds additional information about Complex Breakpoint exceptions.

This register is only implemented if complex hardware breakpoints are present.

Figure 8.53 Debug2 Register Format




Table 8.59 Debug2 Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
0	31:4	Reserved	0	0
Prm	3	Primed - indicates whether a complex breakpoint with an active priming condition was seen on the last debug exception.	R	Undefined
DQ	2	Data Qualified - indicates whether a complex breakpoint with an active data qualifier was seen on the last debug exception.	R	Undefined
Tup	1	Tuple - indicates whether a tuple breakpoint was seen on the last debug exception.	R	Undefined
PaCo	0	Pass Counter - indicates whether a complex breakpoint with an active pass counter was seen on the last debug exception	R	Undefined

8.2.52 Debug Exception Program Counter Register (CP0 Register 24, Select 0)

The Debug Exception Program Counter (*DEPC*) register is a read/write register that contains the address at which processing resumes after a debug exception or debug mode exception has been serviced.

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, complex break), the *DEPC* contains the virtual address of the instruction where execution should resume after the debug handler code is executed.

In processors that implement microMIPS, 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 *ISAMode* 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 8.54 DEPC Register Format




Table 8.60 DEPC Register Formats

Fields		Description	Read/Write	Reset
Name	Bit(s)			
DEPC	31:0	The <i>DEPC</i> 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 DERET instruction causes a jump to the address in the <i>DEPC</i> .	R/W	Undefined

8.2.53 Performance Counter Register (CP0 Register 25, select 0-3)

The M5100 processor defines two performance counters and two associated control registers, 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 8.61.

Table 8.61 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 ANDed with an interrupt enable bit, *IE*, of their respective control register to determine if a performance counter interrupt should be signalled. The two values are then ORed together to create the *SI_PCI* output. Traditionally, this signal is combined with one of the *SI_Int* pins to signal an interrupt to the M5100. However, this is no longer needed as the core will internally route the interrupt to the IP number set by the *IntCtl.IPPCI* field. Counting is not affected by the interrupt indication. This output is cleared when the counter wraps to zero, and may be cleared in software by writing a value with bit 31 = 0 to the *Performance Counter Count* registers.

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 8.55 Performance Counter Control Register




Table 8.62 Performance Counter Control Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
M	31	If this bit is one, another pair of <i>Performance Control</i> and <i>Counter</i> registers is implemented at an MTC0 or MFC0 select field value of 'n+2' and 'n+3'.	R	Preset

Table 8.62 Performance Counter Control Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State										
Name	Bits													
EC	24:23	<p>Event Class. Read/write access in Root mode only; read-only 0 in all other contexts. An implementation may detect the existence of this feature by writing a non-zero value to the field and then reading it back. If the value read is 0, then EC is not supported.</p> <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>Root events counted. [default] Active in Root context.</td> </tr> <tr> <td>1</td> <td>Root intervention events counted, Active in Root context.</td> </tr> <tr> <td>2</td> <td>Guest events counted. Active in Guest context.</td> </tr> <tr> <td>3</td> <td>Guest events plus Root intervention events counted. Active in Guest context. Root will only assign encoding if it wants to give Guest visibility into Root intervention events.</td> </tr> </tbody> </table> <p>Root events are those that occur when $GuestCtl0_{GM}=0$. Root intervention events are those that occur when $GuestCtl0_{GM}=1$ and $!(Root.Status_{EXL}=0$ and $Root.Status_{ERL}=0$ and $Root.Debug_{DM}=0$) Guest events are those that occur when $GuestCtl0_{GM}=1$ and $Root.Status_{EXL}=0$ and $Root.Status_{ERL}=0$ and $Root.Debug_{DM}=0$</p>	Encoding	Meaning	0	Root events counted. [default] Active in Root context.	1	Root intervention events counted, Active in Root context.	2	Guest events counted. Active in Guest context.	3	Guest events plus Root intervention events counted. Active in Guest context. Root will only assign encoding if it wants to give Guest visibility into Root intervention events.	R/W in Root mode. R0 in all others	0
Encoding	Meaning													
0	Root events counted. [default] Active in Root context.													
1	Root intervention events counted, Active in Root context.													
2	Guest events counted. Active in Guest context.													
3	Guest events plus Root intervention events counted. Active in Guest context. Root will only assign encoding if it wants to give Guest visibility into Root intervention events.													
EventExt	14:11	Event specific to Virtualization Module if supported. Possible events are listed in Table 8.63 .	R/W	Undefined										
Event	10:5	Counter event enabled for this counter. Possible events are listed in Table 8.63 .	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										
K	1	Count in Kernel Mode. When this bit is set, count the event in Kernel Mode when EXL and ERL both are 0.	R/W	Undefined										
EXL	0	Count when EXL . When this bit is set, count the event when $EXL = 1$ and $ERL = 0$.	R/W	Undefined										
0	30:25, 22:12, 2	Must be written as zeroes; returns zeroes when read.	0	0										

Table 8.63 describes the events countable with the two performance counters. The mode column indicates whether the event counting is influenced by the mode bits (*U,K,EXL*). The operation of a counter is **UNPREDICTABLE** for events which are specified as Reserved.

Performance counters never count in debug mode or when *ERL* = 1.

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.

Table 8.63 Performance Counter Events Sorted by Event Number

Event Num	Counter 0	Mode	Counter 1	Mode
0	Cycles	No	Cycles	No
1	Instructions completed	Yes	Instructions completed	Yes
2	branch instructions	Yes	Reserved	NA
3	JR r31 (return) instructions	Yes	Reserved	NA
4	JR (not r31) instructions	Yes	Reserved	NA
5	Reserved	NA	Reserved	NA
6	Reserved	NA	Reserved	NA
7	Reserved	NA	Reserved	NA
8	Reserved	NA	Reserved	NA
9	Reserved	NA	Reserved	NA
10	Reserved	NA	Reserved	NA
11	Reserved	NA	Reserved	NA
12	Reserved	NA	Reserved	NA
13	Reserved	NA	Reserved	NA
14	integer instructions completed	Yes	Reserved	NA
15	loads completed	Yes	Stores completed	Yes
16	J/JAL completed	Yes	microMIPS instructions completed	Yes
17	no-ops completed	Yes	Integer multiply/divide completed	Yes
18	Stall cycles	No	Reserved	NA
19	SC instructions completed	Yes	SC instructions failed	Yes
20	Prefetch instructions completed	Yes	Reserved	NA
21	Reserved	NA	Reserved	NA
22	Reserved	NA	Reserved	NA
23	Exceptions taken	Yes	Reserved	NA
24	Reserved	NA	Reserved	NA
25	Reserved	NA	ALU stall cycles	No
26	Reserved	NA	Reserved	NA
27	Reserved	NA	Reserved	NA
28	Reserved	NA	Implementation-specific CP2 event	Yes
29	Reserved	NA	Reserved	NA
30	Implementation-specific CorExtend event	Yes	Reserved	NA

Table 8.63 Performance Counter Events Sorted by Event Number (Continued)

Event Num	Counter 0	Mode	Counter 1	Mode
31	Reserved	NA	Reserved	NA
32	Reserved	NA	Reserved	NA
33	Reserved	NA	Reserved	NA
34	Reserved	NA	Reserved	NA
35	Reserved	NA	CP2 To/From Instructions completed	Yes
36	Reserved	NA	Reserved	NA
37	Reserved	NA	Reserved	NA
38	Reserved	NA	Reserved	NA
39	Reserved	NA	Reserved	NA
40	Uncached stall cycles	Yes	Reserved	NA
41	MDU stall cycles	Yes	Reserved	NA
42	CP2 stall cycles	Yes	CorExtend stall cycles	Yes
43	Reserved	NA	Reserved	NA
44	Reserved	NA	Reserved	NA
45	Load to Use stall cycles	Yes	Reserved	NA
46	Other interlock stall cycles	Yes	Reserved	NA
47	Reserved	NA	Reserved	NA
48	Reserved	NA	Reserved	NA
49	EJTAG Instruction Triggerpoints	Yes	EJTAG Data Triggerpoints	Yes
50	Reserved	NA	Reserved	NA
51	Reserved	NA	Reserved	NA
52	Reserved	NA	Reserved	NA
53	Reserved	NA	Reserved	NA
54	Reserved	NA	Reserved	NA
55	Reserved	NA	Reserved	NA
56:62, 64:127	Reserved	NA	Reserved	NA
63	User-Defined-0	NA	User-defined-1	NA
128	Number of root exceptions taken in guest mode	Yes	Number of guest exits (guest mode to root mode transitions)	Yes
129	Number of GSFC exceptions	Yes	Number of GHFC exceptions	Yes
130	Number of GPSI exceptions	Yes	Number of GRIR exceptions	Yes
131	Number of Hypercall exceptions	Yes	Number of guest related root TLB exceptions that are taken with GuestCtl0.GExcCode=GVA	Yes
132	Number of root TLB exceptions caused by I-side guest translation requests	Yes	Number of root TLB exceptions caused by D-side guest translation requests	Yes
133	Number of root writes that set Guest.Cause.TI to 1	Yes	Number of root writes to Guest.PerfCnt that cause Guest.Cause.PCI to be set to 1	Yes

Table 8.63 Performance Counter Events Sorted by Event Number (Continued)

Event Num	Counter 0	Mode	Counter 1	Mode
134	Reserved	NA	Number of guest accesses to PerfCtr/Ctl registers that cause GPSI when virtually shared	Yes
135	Number of interrupts that cause guest exit in EIC mode	Yes	Number of interrupts that cause guest exit in non-EIC mode	Yes
136-1023	Reserved	NA	Reserved	NA

Table 8.64 Performance Counter Event Descriptions Sorted by Event Type


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 M5100 is built with that clock gater present, none of the counters will increment while the clock is stopped, e.g., due to a WAIT instruction.
Instruction Completion: 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.
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.
microMIPS	1	16	All microMIPS instructions.
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
Cp2 To/From instns	1	35	Includes move to/from, control to/from, and cop2 loads and stores.
Instruction execution events			
SC instructions 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.
Exceptions Taken	0	23	Any type of exception taken.

Table 8.64 Performance Counter Event Descriptions Sorted by Event Type (Continued)

Event Name	Counter	Event Number	Description
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.
General Stalls			
ALU stall cycles	1	25	Counts the number of cycles in which the ALU pipeline cannot advance.
Stall cycles	0	18	Counts the total number of cycles in which no instructions are issued by SRAM to ALU (the RF stage does not advance). This includes both of the previous two events. However, this is different from the sum of them, because cycles when both stalls are active will only be counted once.
Specific stalls - 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.			
Uncached stall cycles	0	40	Cycles in which the processor is stalled on an uncached fetch, load, or store.
MDU stall cycles	0	41	Counts all cycles in which the integer pipeline waits on MDU return data.
Cp2 stall cycles	0	42	Counts all cycles in which the integer pipeline waits on CP2 return data.
CorExtend stall cycles	1	42	Counts all cycles in which the integer pipeline waits on CorExtend return data.
Load to Use stall cycles	0	45	Counts all cycles in which the integer pipeline waits on Load return dependent data.
Other interlocks stall cycles	0	46	Counts all cycles in which the integer pipeline waits on return data from MFC0 and RDHWR instructions.
Implementation-specific events - Modules that can be replaced by the customer will have an event signal associated with them.			
Cp2	1	28	Set to 1 if COP2 is implemented.
CorExtend	0	30	Set to 1 if CorExtend is implemented.
User defined event #0	0	63	User defined event. This is an external controlled event via the pin PM_USER_0. At every rising clock edge where PM_USER_0 is asserted, the performance counter will increment by 1.
User defined event #1	1	63	User defined event. This is an external controlled event via the pin PM_USER_1. At every rising clock edge where PM_USER_1 is asserted, the performance counter will increment by 1.
Virtualization events			
Root exceptions taken in guest mode	0	128	Number of root exceptions taken in guest mode
Guest exits	1	128	Number of guest exits (guest mode to root mode transitions)
GSFC exceptions	0	129	Number of GSFC exceptions
GHFC exceptions	1	129	Number of GHFC exceptions
GPSI exceptions	0	130	Number of GPSI exceptions
GRIR exceptions	1	130	Number of GRIR exceptions

Table 8.64 Performance Counter Event Descriptions Sorted by Event Type (Continued)

Event Name	Counter	Event Number	Description
Hypercall exceptions	0	131	Number of Hypercall exceptions
Guest related root TLB exceptions that are taken with GuestCtl0.GExcCode=GVA	1	131	Number of guest related root TLB exceptions that are taken with GuestCtl0.GExcCode=GVA
Root TLB exceptions caused by I-side guest translation requests	0	132	Number of root TLB exceptions caused by I-side guest translation requests
Root TLB exceptions caused by D-side guest translation requests	1	132	Number of root TLB exceptions caused by D-side guest translation requests
Root writes that set Guest.Cause.TI to 1	0	133	Number of root writes that set Guest.Cause.TI to 1
Root writes to Guest.PerfCnt that cause Guest.Cause.PCI to be set to 1	1	133	Number of root writes to Guest.PerfCnt that cause Guest.Cause.PCI to be set to 1
Guest accesses to PerfCtr/Ctl registers that cause GPSI when virtually shared	1	134	Number of guest accesses to PerfCtr/Ctl registers that cause GPSI when virtually shared
Interrupts that cause guest exit in EIC mode	0	135	Number of interrupts that cause guest exit in EIC mode
Interrupts that cause guest exit in non-EIC mode	1	135	Number of interrupts that cause guest exit in non-EIC mode

Figure 8.56 Performance Counter Count Register**Table 8.65 Performance Counter Count Register Field Descriptions**

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

8.2.54 ErrCtl Register (CP0 Register 26, Select 0)

The *ErrCtl* register controls parity protection of data and instruction SRAM. Parity protection can be enabled or disabled using the *PE* bit.


Figure 8.57 ErrCtl Register Format

Table 8.66 Errctl Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
PE	31	<p>Parity Enable. This bit enables or disables the parity protection for both the instruction SRAM and the data SRAM.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Parity disabled</td> </tr> <tr> <td>1</td> <td>Parity enabled</td> </tr> </tbody> </table> <p>This field is only write-able if the parity option was implemented when the M5100 was built. If parity is not supported, this field is always read as 0. Software can test for parity support by attempting to write a 1 to this field, then read back the value.</p>	Encoding	Meaning	0	Parity disabled	1	Parity enabled	R or R/W	0
Encoding	Meaning									
0	Parity disabled									
1	Parity enabled									
R	30:0	Must be written as zero; returns zero on reads.	0	0						

8.2.55 CacheErr Register (CP0 Register 27, Select 0)

The *CacheErr* register provides an interface with the cache error-detection logic. When a SRAM Parity Error exception is signaled, the fields of this register are set accordingly.

Figure 8.58 CacheErr Register (Primary Caches)




Table 8.67 CacheErr Register Field Descriptions (Primary Caches)

Fields		Description	Read / Write	Reset State						
Name	Bits									
ER	31	<p>Error Reference. Indicates the type of reference that encountered an error.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Instruction</td> </tr> <tr> <td>1</td> <td>Data</td> </tr> </tbody> </table>	Encoding	Meaning	0	Instruction	1	Data	R	Undefined
Encoding	Meaning									
0	Instruction									
1	Data									
EB	30	<p>Error Both. Indicates that a data SRAM parity error occurred in addition to an instruction SRAM parity error.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No additional data SRAM parity error</td> </tr> <tr> <td>1</td> <td>Additional data SRam parity error</td> </tr> </tbody> </table> <p>In the case of an additional data SRAM parity error, the remainder of the bits in this register are set according to the instruction SRAM parity error.</p>	Encoding	Meaning	0	No additional data SRAM parity error	1	Additional data SRam parity error	R	Undefined
Encoding	Meaning									
0	No additional data SRAM parity error									
1	Additional data SRam parity error									
Addr	29:0	Error address. Specifies on which address the error was detected.	R	Undefined						

8.2.56 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.

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

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

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

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

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

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

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

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the error exception PC, and the lower bit of the error 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 8.59 ErrorEPC Register Format




Table 8.68 ErrorEPC Register Field Description

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

8.2.57 DeSave Register (CP0 Register 31, Select 0)

The *Debug Exception Save (DeSave)* register is a read/write register that functions as a simple memory location. This register is used by the debug exception handler to save one of the GPRs that is then used to save the rest of the context

to a pre-determined memory area (such as in the EJTAG Probe). This register allows the safe debugging of exception handlers and other types of code in which the existence of a valid stack for context saving cannot be assumed.

Figure 8.60 DeSave Register Format



Table 8.69 DeSave Register Field Description

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

8.2.58 KScratchn Registers (CP0 Register 31, Selects 2 to 3)

The *KScratchn* registers are optional read/write registers available for scratchpad storage by kernel-mode software. These registers are 32 bits in width for 32-bit processors and 64 bits for 64-bit processors.

The existence of these registers is indicated by the *KScrExist* field in the *Config4* register. The *KScrExist* field specifies which of the selects are populated with a kernel scratch register.

Debug-mode software should not use these registers, but should instead use the *DeSave* register. If EJTAG is implemented, select 0 should not be used for a *KScratch* register. Select 1 is being reserved for future debug use and should not be used for a *KScratch* register.

Figure 8-61 KScratchn Register Format




Table 8.70 KScratchn Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
Data	31:0	Scratch pad data saved by kernel software.	R/W	Undefined

Hardware and Software Initialization of the M5100™ Core

The M5100 processor 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 9.1 “Hardware-Initialized Processor State”](#)
- [Section 9.2 “Software Initialized Processor State”](#)

9.1 Hardware-Initialized Processor State

The M5100 processor core, like most other MIPS processors, 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. *SL_ColdReset* is asserted after power-up to bring the device into a known state. Soft reset can be forced by asserting the *SL_Reset* pin. This distinction is made for compatibility with other MIPS processors. In practice, both resets are handled identically with the exception of the setting of *Status_{SR}*.

9.1.1 Coprocessor 0 State

Much of the hardware initialization occurs in Coprocessor 0.

- *Status_{BEV}* - cleared to 1 on Reset/SoftReset
- *Status_{TS}* - cleared to 0 on Reset/SoftReset
- *Status_{SR}* - cleared to 0 on Reset, set to 1 on SoftReset
- *Status_{NMI}* - cleared to 0 on Reset/SoftReset
- *Status_{ERL}* - set to 1 on Reset/SoftReset
- *Status_{RP}* - cleared to 0 on Reset/SoftReset
- *Config* fields related to static inputs - set to input value by Reset/SoftReset
- *Config_{K0}* - set to 010 (uncached) on Reset/SoftReset
- *Debug_{DM}* - cleared to 0 on Reset/SoftReset (unless EJTAGBOOT option is used to boot into DebugMode, see [Chapter 11, “EJTAG Debug Support in the M5100™ Core”](#) on page 248 for details)
- *Debug_{LSNM}* - cleared to 0 on Reset/SoftReset

Hardware and Software Initialization of the M5100™ Core

- *Debug_{IBusEP}* - cleared to 0 on Reset/SoftReset
- *Debug_{DBusEP}* - cleared to 0 on Reset/SoftReset
- *Debug_{IEXI}* - cleared to 0 on Reset/SoftReset
- *Debug_{SSi}* - cleared to 0 on Reset/SoftReset

9.1.2 Bus State Machines

All pending bus transactions are aborted and the state machines in the SRAM interface unit are reset when a Reset or SoftReset exception is taken.

9.1.3 Static Configuration Inputs

All static configuration inputs should only be changed during Reset.

9.1.4 Fetch Address

Upon Reset/SoftReset, 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.

9.2 Software Initialized Processor State

Software is required to initialize the following parts of the device.

9.2.1 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 in hardware. However, when simulating the operation of the core, unknown values can cause problems. Thus, initializing the register file in the boot code may avoid simulation problems.

9.2.2 Coprocessor 0 State

Miscellaneous COP0 states need to be initialized prior to leaving 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), SW0/1 (Software Interrupts) should be cleared.
- *Config*: Typically, the K0, KU and K23 fields should be set to the desired Cache Coherency Algorithm (CCA) value prior to accessing the corresponding memory regions. But in the M5100 core, all CCA values are treated identically, so the hardware reset value of these fields need not be modified.
- *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. The write to compare will also clear any pending Timer Interrupts (Thus, *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 writeable, 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.

Power Management of the M5100™ Core

The M5100 processor core offers a number of power management features, including low-power design, active power management and power-down modes of operation. The core is a static design that supports a WAIT instruction designed to signal the rest of the device that execution and clocking should be halted, reducing system power consumption during idle periods.

The core provides two mechanisms for system level low-power support discussed in the following sections.

- [Section 10.1 “Register-Controlled Power Management”](#)
- [Section 10.2 “Instruction-Controlled Power Management”](#)

10.1 Register-Controlled Power Management

The *RP* bit in the CP0 *Status* register enables a standard software mechanism for placing the system into a low power state. The state of the RP bit is available externally via the *SI_RP* output signal. Three additional pins, *SI_EXL*, *SI_ERL*, and *EJ_DebugM* support the power management function by allowing the user to change the power state if an exception or error occurs while the core is in a low power state.

Setting the *RP* bit of the CP0 *Status* register causes the core to assert the *SI_RP* signal. The external agent can then decide whether to reduce the clock frequency and place the core into power down mode.

If an interrupt is taken while the device is in power down mode, that interrupt may need to be serviced depending on the needs of the application. The interrupt causes an exception which in turn causes the *EXL* bit to be set. The setting of the *EXL* bit causes the assertion of the *SI_EXL* signal on the external bus, indicating to the external agent that an interrupt has occurred. At this time the external agent can choose to either speed up the clocks and service the interrupt or let it be serviced at the lower clock speed.

The setting of the *ERL* bit causes the assertion of the *SI_ERL* signal on the external bus, indicating to the external agent that an error has occurred. At this time the external agent can choose to either speed up the clocks and service the error or let it be serviced at the lower clock speed.

Similarly, the *EJ_DebugM* signal indicates that the processor is in debug mode. Debug mode is entered when the processor takes a debug exception. If fast handling of this is desired, the external agent can speed up the clocks.

The core provides four power down signals that are part of the system interface. Three of the pins change state as the corresponding bits in the CP0 *Status* register are set or cleared. The fourth pin indicates that the processor is in debug mode:

- The *SI_RP* signal represents the state of the *RP* bit (27) in the CP0 *Status* register.
- The *SI_EXL* signal represents the state of the *EXL* bit (1) in the CP0 *Status* register.
- The *SI_ERL* signal represents the state of the *ERL* bit (2) in the CP0 *Status* register.

- The *EJ_DebugM* signal indicates that the processor has entered debug mode.

10.2 Instruction-Controlled Power Management

The second mechanism for invoking power down mode is through execution of the WAIT instruction. If the bus is idle at the time the WAIT instruction reaches the M stage of the pipeline the internal clocks are suspended and the pipeline is frozen. However, the internal timer and some of the input pins (*SL_Int*[5:0], *SL_NMI*, *SL_Reset*, *SL_ColdReset*, and *EJ_DINT*) continue to run. If the bus is not idle at the time the WAIT instruction reaches the M stage, the pipeline stalls until the bus becomes idle, at which time the clocks are stopped. When the CPU is in instruction controlled power management mode, any enabled interrupt, NMI, debug interrupt, or reset condition causes the CPU to exit this mode and resume normal operation. While the part is in this low-power mode, the *SL_SLEEP* signal is asserted to indicate to external agents what the state of the chip is.

EJTAG Debug Support in the M5100™ Core

The EJTAG debug logic in the M5100 processor core provides three optional modules:

1. Hardware breakpoints
2. Test Access Port (TAP) for a dedicated connection to a debug host
3. Tracing of program counter/data address/data value trace to On-chip memory or to a Trace probe

These features are covered in the following sections:

- [Section 11.1 “Debug Control Register”](#)
- [Section 11.2 “Hardware Breakpoints”](#)
- [Section 11.3 “Complex Breakpoint Usage”](#)
- [Section 11.4 “Test Access Port \(TAP\)”](#)
- [Section 11.5 “EJTAG TAP Registers”](#)
- [Section 11.6 “TAP Processor Accesses”](#)
- [Section 11.7 “SecureDebug”](#)
- [Section 11.8 “iFlowtrace™ Mechanism”](#)
- [Section 11.9 “PC/Data Address Sampling”](#)
- [Section 11.10 “Fast Debug Channel”](#)
- [Section 11.11 “cJTAG Interface”](#)

11.1 Debug Control Register

The *Debug Control Register (DCR)* register controls and provides information about debug issues, and is always provided with the CPU 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 none, some or all sources for soft reset. The soft reset masking may 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 running on the CPU if the probe expects to service dmseg accesses. The reset value in the table below takes effect on both hard and soft resets.

Figure 11.1 DCR Register Format

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
EJTAG_Brk_Override	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 11.1 DCR Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
EJTAG_Brk_Override	31	<p>Override <i>EjtagBrk</i> and <i>DINT</i> disable. Refer to Section 11.7 “SecureDebug”.</p> <p>Re-enable <i>EjtagBrk</i> and <i>DINT</i> signal during boot.</p> <p>Allows <i>EjtagBrk</i> to be asserted by an EJTAG probe (or assertion of <i>DINT</i> signal), resulting in a request for a Debug Interrupt exception from the processor. This provides a means of recovering the cpu from crash, hang, loop or low-power mode.</p> <p>This feature can allow a Debug Executive to communicate with the probe over the Fast Debug Channel (FDC) and provides a host-based debugger the ability to query the target processor via Debug Executive commands, useful for determining cause of hang.</p> <p>Software can write this bit and read back to determine if the Secure Debug feature is implemented.</p>	R/W	0						
ENM	29	<p>Endianness in which the processor is running in kernel and Debug Mode:</p> <table border="1"> <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	Externally Set
Encoding	Meaning									
0	Little endian									
1	Big endian									

Table 11.1 DCR 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 PCs captured</td> </tr> <tr> <td>1</td> <td>Capture only PCs that miss the cache.</td> </tr> </tbody> </table>	Encoding	Meaning	0	All PCs captured	1	Capture only PCs that miss the cache.	R	0
Encoding	Meaning									
0	All PCs captured									
1	Capture only PCs that miss the cache.									
PCnoASID	25	<p>Controls whether the PCSAMPLE scan chain includes or omits the <i>ASID</i> field. An 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									
FDCImpl	18	<p>Indicates if the fast debug channel is implemented:</p> <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									

Table 11.1 DCR Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bits									
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	Preset
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	Preset
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									
RDVec	11	Enables relocation of the debug exception vector. The value in the DebugVectorAddr register is used for EJTAG exceptions when <i>ProbTrap</i> =0, and <i>RDVec</i> =1.	R/W	0						
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	Preset
Encoding	Meaning									
0	No complex breakpoint block implemented									
1	Complex breakpoint block implemented									

Table 11.1 DCR Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bits									
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									
SRstE	1	Soft Reset Enable This bit allows the system to mask soft resets. The core does not internally mask soft resets. Rather the state of this bit appears on the <i>EJ_SRstE</i> external output signal, allowing the system to mask soft resets if desired.	R/W	1						

Table 11.1 DCR Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bits									
ProbEn	0	Probe Enable. This bit reflects the <i>ProbEn</i> bit in the <i>EJTAG Control</i> register: <table border="1" data-bbox="557 443 1078 590"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No accesses to dmseg allowed</td> </tr> <tr> <td>1</td> <td>Accesses to dmseg by EJTAG probe services allowed</td> </tr> </tbody> </table>	Encoding	Meaning	0	No accesses to dmseg allowed	1	Accesses to dmseg by EJTAG probe services allowed	R	Same value as <i>ProbEn</i> in <i>ECR</i> (see Table 9-4)
Encoding	Meaning									
0	No accesses to dmseg allowed									
1	Accesses to dmseg by EJTAG probe services allowed									
0	30, 28:27, 21:19, 13:12	Must be written as zeros; return zeros on reads.	0	0						

11.2 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 generally added to breakpoint, unless required to distinguish it from a software breakpoint.

There are two types of simple hardware breakpoints implemented in the M5100 core: Instruction breakpoints and Data breakpoints. The M5100 core may also contain a complex breakpoint unit.

A core may be configured with the following breakpoint options:

- No data or instruction breakpoints, without complex break support
- Two instruction and one data breakpoint, without complex break support
- Four instruction and two data breakpoints, without complex break support
- Six instruction and two data breakpoints, without support for complex breaks
- Six instruction and two data breakpoints, with support for complex breaks
- Eight instruction and four data breakpoints, without support for complex breaks
- Eight instruction and four data breakpoints, with support for complex breaks

Instruction breaks occur on instruction fetch operations, and the break is set on the virtual address on the bus between the CPU and the instruction cache. 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 (the value of PC) with the registers for each instruction breakpoint, including masking of address. When an instruction breakpoint matches, a debug

exception and/or a trigger is generated. An internal bit in the instruction breakpoint registers is set to indicate that the match occurred.

11.2.1 Data Breakpoints

Data breakpoints occur on load/store transactions. Breakpoints are set on virtual address 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), accessed bytes (BYTELANE) 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 debug exception and/or a trigger is generated, and an internal bit in the data breakpoint registers is set to indicate that the match occurred. The match is precise in that the debug exception or trigger occurs on the instruction that caused the breakpoint to match.

11.2.2 Complex Breakpoints

The complex breakpoint unit utilizes the instruction and data breakpoint hardware and looks for more specific matching conditions. There are several different types of enabling that allow more exact breakpoint specification. Tuples add an additional condition to data breakpoints of requiring an instruction breakpoint on the same instructions. Pass counters are counters that decrement each time a matching breakpoint condition is taken. When the counter reaches 0, the break or trigger effect of the breakpoint is enabled. Priming allows a breakpoint to only be enabled when another trigger condition has been detected. Data qualification allows instruction breakpoints to only be enabled when a corresponding load data triggerpoint has matched both address and data. Data qualified breakpoints are also disabled if a load is executed that matches on the address portion of the triggerpoint, but has a mismatching data value. The complex breakpoint features can be combined to create very complex sequences to match on.

In addition to the breakpoint logic, the complex break unit also includes a Stopwatch Timer block. This counter can be used to measure time spent in various sections. It can either be free-running, or it can be set up to start and stop counting based on a trigger from instruction breakpoints.

11.2.3 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, and the conditions for matching instruction and data breakpoints are described below. The breakpoints only match for instructions executed in non-debug mode, thus never on instructions executed in debug mode.

The match of an enabled breakpoint can either generate a debug exception or a trigger indication. 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.

11.2.3.1 Conditions for Matching Instruction Breakpoints

There are two methods for matching conditions: Equality and Mask or Address Range.

Equality and Mask

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 bit level. The registers for each instruction breakpoint have the values and mask used in the compare, and the equation that determines the match is shown below in C-like notation.

```
IB_match =
    ( <all 1's> == ( IBMnIBM | ~ ( PC ^ IBAnIBA ) ) )
```

The match indication for instruction breakpoints is always precise, i.e. indicated on the instruction causing the IB_match to be true.

Address Range

Cores may optionally support the address range triggered instruction breakpoints. When this feature is configured, the following changes are made to the instruction breakpoint registers:

- *IBAn* : represents the upper limit of a address range boundary
- *IBMn* : represents the lower limit of the address range boundary

In addition, the following bits must be supported:

IBCn[6].hwarts : a preset value of 1 indicates that the address range triggered instruction breakpoint feature is supported for this particular instruction breakpoint channel. This bit is read-only.

IBCn[5].excl : a value of 0 indicates that the breakpoint will match for addresses inclusive (within) the range defined by *IBMn* and *IBAn*. A value of 1 indicates that the breakpoint will match for addresses exclusive (outside) to the range defined by *IBMn* and *IBAn*. This bit is writeable.

IBCn[4].hwart : a value of 0 indicates that the breakpoint will match using the “Equality and Mask” equation as found in Section 11.2.3.1 “Conditions for Matching Instruction Breakpoints”. A value of 1 indicates that the breakpoint will match using address ranges using the equation below:

```
IB_match =
(!IBCnTCuse || ( TC == IBCnTC ) ) &&
( ! IBCnASIDuse || ( ASID == IBASIDnASID ) ) &&
( (~IBCnhwarts || ~IBCnhwart) &&
  (( IBMnIBM | ~ ( PC ^ IBAnIBA ) ) == ~0) ||
  (( IBCnhwarts && IBCnhwart) &&
  (~IBCnexcl && (IBM <= PC <= IBA)) ||
  ( IBCnexcl && (IBM > PC || PC > IBA)
)
)
```

Or if microMIPS is supported:

```
IB_range_match =
(!IBCnTCuse || ( TC == IBCnTC ) ) &&
( ! IBCnASIDuse || ( ASID == IBASIDnASID ) ) &&
( (~IBCnhwarts || ~IBCnhwart) &&
  (( IBMnIBM | ~ ( ( ( PC[MSB:1] << 1 ) + ISAmode ) ^ IBAnIBA ) ) == ~0 ) ||
  (( IBCnhwarts && IBCnhwart) &&
  ( IBMnIBM[0] | ~ ( ISAmode ^ IBAnIBA[0] ) ) == ~0) &&
  (~IBCnexcl && (IBM[MSB:1] <= PC[MSB:1] <= IBA[MSB:1])) ||
)
```

```
( IBCnexc1 && ( IBM[MSB:1] > PC[MSB:1] || PC[MSB:1] > IBA[MSB:1] ) )
```

Also note that addresses that overlap a boundary is considered for both exclusive and inclusive breakpoint matches.

11.2.3.2 Conditions for Matching Data Breakpoints

There are two methods for matching conditions, namely 1) by Equality and Mask or 2) by Address Range:

Equality and Mask

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 load/store for coprocessor, 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. The registers for each data breakpoint have 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 =
    ( ( ( TYPE == load ) && ! DBCn_NoLB ) ||
      ( ( TYPE == store ) && ! DBCn_NoSB ) ) &&
    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) 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 =
    ( <all 1's> == ( DBMn_DBM | ~ ( ADDR ^ DBAn_DBA ) ) ) &&
    ( <all 0's> != ( ~ BAI & BYTELANE ) )
```

The size of $DBCn_{BAI}$ and BYTELANE is 4 bits.

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> == ( DBCn_BLM | DBCn_BAI | ~ BYTELANE ) )
```

The size of $DBCn_{BLM}$, $DBCn_{BAI}$ and BYTELANE is 4 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 $DBCn_{VM}$ bit inverts the sense of the match - if set, the value match term will be high if the data value is not the same as the data in the $DBVn$ register. 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 =
```

```

DBCnIVM ^
(( (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] ) )

```

The match for a data breakpoint 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.

Address Range

Cores may optionally support the address range triggered data breakpoints. When this feature is configured, the following changes are made to the data breakpoint registers:

- *DBAn* : represents the upper limit of a address range boundary
- *DBMn* : represents the lower limit of the address range boundary

In addition, the following bits must be supported:

DBCn[10].hwarts: a preset value of 1 indicates that the address range triggered data breakpoint feature is supported for this particular data breakpoint channel. This bit is read-only.

DBCn[9].exc : a value of 0 indicates that the breakpoint will match for addresses inclusive (within) the range defined by *DBMn* and *DBAn*. A value of 1 indicates that the breakpoint will match for addresses exclusive (outside) to the range defined by *DBMn* and *DBAn*. This bit is writeable.

DBCn[8].hwart: a value of 0 indicates that the breakpoint will match using the “Equality and Mask” equation as found in Section 11.2.3.2 “Conditions for Matching Data Breakpoints”. A value of 1 indicates that the breakpoint will match using address ranges using the equation below:

```

DB_match =
(!DBCnTCuse || ( TC == DBCnTC ) ) &&
( ( ( TYPE == load ) && ! DBCnNoLB ) || ( ( TYPE == store ) && ! DBCnNoSB ) ) &&
DB_addr_range_match && ( DB_no_value_compare || DB_value_match )

DB_addr_range_match =
( ! DBCnASIDuse || ( ASID == DBASIDnASID ) ) &&
( ( (~DBCnhwarts || ~DBCnhwart) &&
  ( ( DBMnDBM | ~ ( ADDR ^ DBAnDBA ) ) == ~0 ) ||
  ( ( DBCnhwarts && DBCnhwart) &&
    ( (~DBCnexcl && (DBMn <= ADDR <= DBAn) ) ||
      ( DBCnexcl && (DBMn > ADDR || ADDR > DBAn)
    )
  )
)

```

When address range triggered data breakpoints is enabled, *DBCn.BLM[3:0]* must be set to 4'b1111 because value matching is not supported with this feature. Addresses that overlap a boundary is considered for both exclusive and inclusive breakpoint matches.

11.2.4 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.

11.2.4.1 Debug Exception by Instruction Breakpoint

If the breakpoint is enabled by *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 *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 can not 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.

11.2.4.2 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 debug data break exception occurs when a data breakpoint 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.
- A load transaction for a breakpoint with data value compare must occur from the memory system, since the value is required in order to evaluate the breakpoint.

The result of this is that the load or store instruction causing the debug data break exception appears as not executed, with the exception that a load from the memory system does occur for a breakpoint with data value compare, but the register file is not updated by the load.

If both data breakpoints without and with data value compare would match the same transaction and generate a debug exception, then the following rules apply with respect to updating the *BS[n]* bits.

- On both a load and store the *BS[n]* bits are required to be set for all matching breakpoints without a data value compare.

- On a store the $BS[n]$ bits are allowed but not required to be set for all matching breakpoints with a data value compare, but either all or none of the $BS[n]$ bits must be set for these breakpoints.
- On a load then none of the $BS[n]$ bits for breakpoints with data value compare are allowed to be set, since the load is not allowed to occur due to the debug exception from a breakpoint without a data value compare, and a valid data value is therefore not returned.

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. This re-execution may result in a repeated load from system memory, since the load may have occurred previously in order to evaluate the breakpoint as described above. I/O devices with side effects on loads may not be re-accessible without changing the system behavior. The Load Data Value register was introduced to capture the value that was read and allow debug software to synthesize the load instruction without re-accessing memory. Debug software is responsible for disabling breakpoints when returning to the instruction, otherwise the debug data break exception will reoccur.

11.2.5 Breakpoint Used as Triggerpoint

Both instruction and data hardware breakpoints can be setup by software so that a matching breakpoint does not generate a debug exception, but only an indication through the $BS[n]$ bit. The TE bit in the $IBCn$ or $DBCn$ register controls if an instruction or data breakpoint is used as a so-called triggerpoint. The triggerpoints are, like breakpoints, only compared for instructions executed in non-debug mode.

The $BS[n]$ bit in the IBS or DBS register is set when the respective IB_match or DB_match bit is true.

The triggerpoint feature can be used to start and stop tracing.

11.2.6 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 $drseg$, and the addresses are shown in [Table 11.2](#).

Table 11.2 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
0x1120 + n * 0x100	<i>IBCCn</i>	Instruction Breakpoint Complex Control n
0x1128 + n * 0x100	<i>IBPCn</i>	Instruction Breakpoint Pass Counter n
n is breakpoint number in range 0 to 5 (or 3 or 1, depending on the implemented hardware)		

An example of some of the registers; *IBA0* is at offset 0x1100 and *IBC2* is at offset 0x1138.

11.2.6.1 Instruction Breakpoint Status (IBS) Register (0x1000)

The *Instruction Breakpoint Status (IBS)* register holds implementation and status information about the instruction breakpoints. This register is required only if instruction breakpoints are implemented.

Figure 11.2 IBS Register Format




Table 11.3 IBS Register Field Descriptions

Fields		Description	Read/Wr ite	Reset State
Name	Bit(s)			
Res	31	Must be written as zero; returns zero on read.	R	0
ASIDsup	30	Indicates that ASID compare is supported in instruction breakpoints. 0: No ASID compare. 1: ASID compare (<i>IBASID_n</i> register implemented).	R	0
Res	29:28	Must be written as zero; returns zero on read.	R	0
BCN	27:24	Number of instruction breakpoints implemented.	R	0, 2, 4, 6 or 8 ^a
Res	23:8	Must be written as zero; returns zero on read.	R	0
BS	7:0	Break status for breakpoint n is at <i>BS[n]</i> , with n from 0 to 7 ^b . The bit is set to 1 when the condition for the corresponding breakpoint has matched and <i>IBCnTE</i> or <i>IBCnBE</i> are set	R/W	Undefined

[a] Based on actual hardware implemented.
 [b] In case of fewer than 8 Instruction breakpoints the upper bits become reserved.

11.2.6.2 Instruction Breakpoint Address n (IBAn) Register (0x1100 + n * 0x100)

The *Instruction Breakpoint Address n (IBAn)* register has the address used in the condition for instruction breakpoint n. This register is required only if instruction breakpoints are implemented.

Figure 11.3 IBAn Register Format




Table 11.4 IBAn Register Field Descriptions

Fields		Description	Read/Wr rite	Reset State
Name	Bit(s)			
IBA	31:0	Instruction breakpoint address for condition.	R/W	Undefined

11.2.6.3 Instruction Breakpoint Address Mask n (IBM n) Register (0x1108 + $n \times 0x100$)

The Instruction *Breakpoint Address Mask n* (IBM n) register has the mask for the address compare used in the condition for instruction breakpoint n . A 1 indicates that the corresponding address bit will not be considered in the match. A mask value of all 0's would require an exact address match, while a mask value of all 1's would match on any address. This register is required only if instruction breakpoints are implemented.

Figure 11.4 IBM n Register Format



Table 11.5 IBM n Register Field Descriptions

Fields		Description	Read/W rite	Reset State						
Name	Bit(s)									
IBM	31:0	Instruction breakpoint address mask for condition: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Corresponding address bit not masked.</td> </tr> <tr> <td>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.									

11.2.6.4 Instruction Breakpoint ASID n (IBASID n) Register (0x1110 + $n \times 0x100$)

For processors with a TLB-based MMU, this register is used to define an ASID value to be used in the match expression. On the M5100 processor, this register is reserved and reads as 0. This register is required only if instruction breakpoints are implemented.

Figure 11.5 IBASID n Register Format




Table 11.6 IBASID n Register Field Descriptions

Fields		Description	Read/Wr ite	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	0

11.2.6.5 Instruction Breakpoint Control n (IBC n) Register (0x1118 + $n \times 0x100$)

The *Instruction Breakpoint Control n* (IBC n) register controls the setup of instruction breakpoint n . This register is required only if instruction breakpoints are implemented.

Figure 11.6 IBCn Register Format




Table 11.7 IBCn Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bits									
Res	31:24	Must be written as zero; returns zero on read.	R	0						
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="text-align: center;">Encoding</th> <th style="text-align: center;">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	0
Encoding	Meaning									
0	Don't use ASID value in compare									
1	Use ASID value in compare									
Res	22:7	Must be written as zero; returns zero on read.	R	0						
hwarts	6	A preset value of 1 indicates that the address- range triggered instruction breakpoint feature is supported for this particular instruction breakpoint channel.	R	Preset						
excl	5	A value of 0 indicates that the breakpoint will match for addresses within (inclusive of) the range defined by <i>IBMn</i> and <i>IBAn</i> . A value of 1 indicates that the breakpoint will match for addresses outside (exclusive to) the range defined by <i>IBMn</i> and <i>IBAn</i> .	R/W	0						
hwart	4	A value of 0 indicates that the breakpoint will match using the “Equality and Mask” equation as found section under 11.2.3.1 “Conditions for Matching Instruction Breakpoints”. A value of 1 indicates that the breakpoint will match using the “Address Range” equation in section 11.2.3.1 “Conditions for Matching Instruction Breakpoints”	R/W	0						
Res	3	Must be written as zero; returns zero on read.	R	0						
TE	2	Use instruction breakpoint n as triggerpoint: <table border="1" style="margin: 5px auto; 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;">0</td> <td>Don't use it as triggerpoint</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Use it as triggerpoint</td> </tr> </tbody> </table>	Encoding	Meaning	0	Don't use it as triggerpoint	1	Use it as triggerpoint	R/W	0
Encoding	Meaning									
0	Don't use it as triggerpoint									
1	Use it as triggerpoint									
Res	1	Must be written as zero; returns zero on read.	R	0						
BE	0	Use instruction breakpoint n as breakpoint: <table border="1" style="margin: 5px auto; 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;">0</td> <td>Don't use it as breakpoint</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Use it as breakpoint</td> </tr> </tbody> </table>	Encoding	Meaning	0	Don't use it as breakpoint	1	Use it as breakpoint	R/W	0
Encoding	Meaning									
0	Don't use it as breakpoint									
1	Use it as breakpoint									

If complex breakpoints are implemented, there will be an 8b pass counter for each of the instruction breakpoints on the M5100 core.

Figure 11.8 IBPCn Register Format




Table 11.9 IBPCn Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
0	31:8	Ignored on write, returns zero on read.	R	0
PassCnt	7:0	Prevents a break/trigger action until the matching conditions on breakpoint n have been seen this number of times. Each time the matching condition is seen, this value will be decremented by 1. When the value reaches 0, subsequent matches will cause a break or trigger as requested and the counter will stay at 0. The break or trigger action is imprecise if the <i>PassCnt</i> register was last written to a non-zero value. It will remain imprecise until this register is written to 0 by software. The instruction pass counter should not be set on instruction breakpoints that are being used as part of a tuple breakpoint.	R/W	0

11.2.7 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 11.10](#).

Table 11.10 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
0x2128 + 0x100 * n	<i>DBCCn</i>	Data Breakpoint Complex Control n
0x2130 + 0x100 * n	<i>DBPCn</i>	Data Breakpoint Pass Counter n
0x2ff0	<i>DVM</i>	Data Value Match Register

n is breakpoint number as 0, 1, 2 or 3 (or just 0, depending on the implemented hardware)

An example of some of the registers; *DBM0* is at offset 0x2108 and *DBV1* is at offset 0x2220.

11.2.7.1 Data Breakpoint Status (DBS) Register (0x2000)

The *Data Breakpoint Status (DBS)* register holds implementation and status information about the data breakpoints. This register is required only if data breakpoints are implemented.

Figure 11.9 DBS Register Format




Table 11.11 DBS Register Field Descriptions

Fields		Description	Read/Wr ite	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. 0: Not supported 1: Supported	R	0
Res	29:28	Must be written as zero; returns zero on read.	R	0
BCN	27:24	Number of data breakpoints implemented.	R	4, 2, 1 or 0 ^a
Res	23:4	Must be written as zero; returns zero on read.	R	0
BS	3:0	Break status for breakpoint n is at $BS[n]$, with n from 0 to 1 ^b . The bit is set to 1 when the condition for the corresponding breakpoint has matched.	R/W0	Undefined

[a] Based on actual hardware implemented.
[b] In case of only 1 data breakpoint bit 1 become reserved.

11.2.7.2 Data Breakpoint Address n (DBAn) Register (0x2100 + 0x100 * n)

The *Data Breakpoint Address n (DBAn)* register has the address used in the condition for data breakpoint n . This register is required only if data breakpoints are implemented.

Figure 11.10 DBAn Register Format




Table 11.12 DBAn Register Field Descriptions

Fields		Description	Read/W rite	Reset State
Name	Bit(s)			
DBA	31:0	Data breakpoint address for condition.	R/W	Undefined

11.2.7.3 Data Breakpoint Address Mask n (DBMn) Register (0x2108 + 0x100 * n)

The *Data Breakpoint Address Mask n (DBMn)* register has the mask for the address compare used in the condition for data breakpoint n. A 1 indicates that the corresponding address bit will not be considered in the match. A mask value of all 0's would require an exact address match, while a mask value of all 1's would match on any address. This register is required only if data breakpoints are implemented.

Figure 11.11 DBMn Register Format




Table 11.13 DBMn Register Field Descriptions

Fields		Description	Read/W rite	Reset State
Name	Bit(s)			
DBM	31:0	Data breakpoint address mask for condition: 0: Corresponding address bit not masked 1: Corresponding address bit masked	R/W	Undefined

11.2.7.4 Data Breakpoint ASID n (DBASIDn) Register (0x2110 + 0x100 * n)

For processors with a TLB-based MMU, this register is used to define an ASID value to be used in the match expression. On the M5100 processor, this register is reserved and reads as 0. This register is required only if data breakpoints are implemented.

Figure 11.12 DBASIDn Register Format




Table 11.14 DBASIDn Register Field Descriptions

Fields		Description	Read/Wr ite	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	0

11.2.7.5 Data Breakpoint Control n (DBCn) Register (0x2118 + 0x100 * n)

The *Data Breakpoint Control n (DBCn)* register controls the setup of data breakpoint n. This register is required only if data breakpoints are implemented.

Figure 11.13 DBCn Register Format




Table 11.15 DBCn Register Field Descriptions

Fields		Description	Read/Write	Reset State						
Name	Bits									
Res	31:24	Must be written as zero; returns zero on reads.	R	0						
ASIDuse	23	Use ASID value in compare for data breakpoint n: <table border="1" data-bbox="482 417 984 537"> <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	0
Encoding	Meaning									
0	Don't use ASID value in compare									
1	Use ASID value in compare									
Res	22:18	Must be written as zero; returns zero on reads.	R	0						
BAI	17:14	Byte access ignore controls ignore of access to a specific byte. <i>BAI[0]</i> ignores access to byte at bits [7:0] of the data bus, <i>BAI[1]</i> ignores access to byte at bits [15:8], etc. <table border="1" data-bbox="482 699 984 877"> <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 not fulfilled on a store transaction: <table border="1" data-bbox="482 963 984 1142"> <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 not fulfilled on a load transaction: <table border="1" data-bbox="482 1228 984 1407"> <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									
Res	11	Must be written as zero; returns zero on reads.	R	0						
hwarts	10	A preset value of 1 indicates that the address range triggered data breakpoint feature is supported for this particular data breakpoint channel.	R	Preset						
excl	9	A value of 0 indicates that the breakpoint will match for addresses inclusive (within) the range defined by <i>DBMn</i> and <i>DBAn</i> . A value of 1 indicates that the breakpoint will match for addresses exclusive (outside) of the range defined by <i>DBMn</i> and <i>DBAn</i> .	R/W	0						

Table 11.15 DBCn Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bits									
hwart	8	A value of 0 indicates that the breakpoint will match using the “Equality and Mask” equation as found section under 11.2.3.2 “Conditions for Matching Data Breakpoints”. A value of 1 indicates that the breakpoint will match using the “Address Range” equation in section 11.2.3.2 “Conditions for Matching Data Breakpoints”	R/W	0						
BLM	7:4	Byte lane mask for value compare on data breakpoint. <i>BLM[0]</i> masks byte at bits [7:0] of the data bus, <i>BLM[1]</i> masks byte at bits [15:8], etc.: <table border="1" style="margin-left: 20px;"> <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									
Res	3	Must be written as zero; returns zero on reads.	R	0						
TE	2	Use data breakpoint n as triggerpoint: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Don't use it as triggerpoint</td> </tr> <tr> <td>1</td> <td>Use it as triggerpoint</td> </tr> </tbody> </table>	Encoding	Meaning	0	Don't use it as triggerpoint	1	Use it as triggerpoint	R/W	0
Encoding	Meaning									
0	Don't use it as triggerpoint									
1	Use it as triggerpoint									
IVM	1	Invert Value Match. When set, the data value compare will be inverted. i.e., a break or trigger will be taken if the value does not match the specified value	R/W	0						
BE	0	Use data breakpoint n as breakpoint: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Don't use it as breakpoint</td> </tr> <tr> <td>1</td> <td>Use it as breakpoint</td> </tr> </tbody> </table>	Encoding	Meaning	0	Don't use it as breakpoint	1	Use it as breakpoint	R/W	0
Encoding	Meaning									
0	Don't use it as breakpoint									
1	Use it as breakpoint									

11.2.7.6 Data Breakpoint Value n (DBVn) Register (0x2120 + 0x100 * n)

The *Data Breakpoint Value n (DBVn)* register has the value used in the condition for data breakpoint n. This register is required only if data breakpoints are implemented.

Figure 11.14 DBVn Register Format



Table 11.16 DBVn Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
DBV	31:0	Data breakpoint value for condition.	R/W	Undefined

11.2.7.8 Data Breakpoint Pass Counter n (DBPCn) Register (0x2130 + n*0x100)

The *Data Breakpoint Pass Counter n (DBPCn)* register controls the pass counter associated with data breakpoint n. This register is required only if complex breakpoints are implemented and only for implemented data breakpoints.

If complex breakpoints are implemented, there will be an 16b pass counter for each of the data breakpoints on the M5100 core.

Figure 11.16 DBPCn Register Format




Table 11.18 DBPCn Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
0	31:16	Ignored on write, returns zero on read.	R	0
PassCnt	15:0	Prevents a break/trigger action until the matching conditions on data breakpoint n have been seen this number of times. Each time the matching condition is seen, this value will be decremented by 1. When the value reaches 0, subsequent matches will cause a break or trigger as requested and the counter will stay at 0. The break or trigger action is imprecise if the <i>PassCnt</i> register was last written to a non-zero value. It will remain imprecise until this register is written to 0 by software.	R/W	0

11.2.7.9 Data Value Match (DVM) Register (0x2ff0)

The *Data Value Match (DVM)* register captures the data value of a load that takes a precise data value breakpoint. This allows debug software to synthesize the load instruction without re-executing it in case it is to a system register that has destructive reads. This register is required only if data breakpoints are implemented.

Figure 11.17 DVM Register Format



Table 11.19 DVM Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
LDV	31:0	Load data value for the last precise load data value breakpoint taken.	R	Undefined

11.2.8 Complex Breakpoint Registers

The registers for complex breakpoints are described [Table 11.20](#). These registers have implementation information and are used to setup the data breakpoints. All registers are in drseg.

Table 11.20 Addresses for Complex Breakpoint Registers

Offset in drseg	Register Mnemonic	Register Name and Description
$0x1120 + 0x100 * n$	<i>IBCCn</i>	Instruction Breakpoint Complex Control n - described above with instruction breakpoint registers
$0x1128 + 0x100 * n$	<i>IBPCn</i>	Instruction Breakpoint Pass Counter n - described above with instruction breakpoint registers
$0x2128 + 0x100 * n$	<i>DBCCn</i>	Data Breakpoint Complex Control n - described above with data breakpoint registers
$0x2130 + 0x100 * n$	<i>DBPCn</i>	Data Breakpoint Pass Counter n - described above with data breakpoint registers
0x8000	<i>CBTControl</i>	Complex Break and Triggerpoint Control - indicates which of the complex breakpoint features are implemented
$0x8300 + 0x20 * n$	<i>PrCndAIn</i>	Prime Condition Register A for Instruction breakpoint n
$0x84e0 + 0x20 * n$	<i>PrCndADn</i>	Prime Condition Register A for Data breakpoint n
0x8900	<i>STCtl</i>	Stopwatch Timer Control
0x8908	<i>STCnt</i>	Stopwatch Timer Count

n is breakpoint number from 0 to 7 (range dependent on implemented hardware)

11.2.8.1 Complex Break and Trigger Control (CBTC) Register (0x8000)

The *CBTC* register contains configuration bits that indicate which features of complex break are implemented as well as a control bit for the stopwatch timer. On the M5100 core, if complex break is implemented, all of the separate features will be present. This register is required only if complex breakpoints are implemented.

Figure 11.18 CBTC Register Format




Table 11.21 CBTC Register Field Descriptions

Fields		Description	Read/Write	Reset State
Name	Bits			
Res	31:9, 7:5	Reserved	R	0
STMode	8	Stopwatch Timer Mode: controls whether the stopwatch timer is free-running or controlled by triggerpoints: 0 - free-running 1 - started and stopped by instruction triggers	R/W	1

Table 11.21 CBTC Register Field Descriptions (Continued)

Fields		Description	Read/Write	Reset State
Name	Bits			
STP	4	Stopwatch Timer Present - indicates whether stopwatch timer is implemented.	R	1
PP	3	Priming Present - indicates whether primed breakpoints are supported	R	1
DQP	2	Data Qualify Present - indicates whether data qualified breakpoints are supported.	R	1
TP	1	Tuple Present - indicates whether any tuple breakpoints are implemented.	R	1
PCP	0	Pass Counters Present - indicates whether any break-points have pass counters associated with them.	R	1

11.2.8.2 Priming Condition A (PrCndA/Dn) Registers

The *Prime Condition* registers hold implementation specific information about which triggerpoints are used for the priming conditions for each breakpoint register. On the M5100 core, these connections are predetermined and these registers are read-only. This register is required only if complex breakpoints are implemented.

The architecture allows for up to 16 priming conditions to be specified and there can be up to 4 priming condition registers per breakpoint (A/B/C/D). The M5100 core only allows for 4 priming conditions and thus only implements the PrCndA registers. The general description is shown in [Table 11.22](#). The actual priming conditions for each of the breakpoints are shown in [Table 11.23](#).

Figure 11.19 PrCndA Register Format




Table 11.22 PrCndA Register Field Descriptions

Fields		Description	Read/Wr ite	Reset State
Name	Bit(s)			
CondN	31:24 23:16 15:8 7:0	Specifies which triggerpoint is connected to priming condition 3, 2, 1, or 0 ^a for the current breakpoint.	R	Preset
	31:30 23:22 15:14 7:6	Reserved	R	0
	29:28 21:20 13:12 5:4	Trigger type 00 - Special/Bypass 01 - Instruction 10 - Data 11 - Reserved	R	Preset
	27:24 19:16 11:8 3:0	Break Number, 0-14	R	Preset
[a] Condition 0 is always Bypass and will read as 8 b0				

Table 11.23 Priming Conditions and Register Values for 6I/2D Configuration

Break	Cond0	Cond1	Cond2	Cond3	PrCndA Value	drseg offset
Inst0	Bypass	Data0	Inst1	Inst2	0x1211_2000	0x8300
Inst1	Bypass	Data0	Inst0	Inst2	0x1210_2000	0x8320
Inst2	Bypass	Data0	Inst0	Inst1	0x1110_2000	0x8340
Inst3	Bypass	Data1	Inst4	Inst5	0x1514_2100	0x8360
Inst4	Bypass	Data1	Inst3	Inst5	0x1513_2100	0x8380
Inst5	Bypass	Data1	Inst3	Inst4	0x1413_2100	0x83a0
Data0	Bypass	Inst0	Inst1	Inst2	0x1211_1000	0x84e0
Data1	Bypass	Inst3	Inst4	Inst5	0x1514_1300	0x8500

Table 11.24 Priming Conditions and Register Values for 8I/4D Configuration

Break	Cond0	Cond1	Cond2	Cond3	PrCndA Value	drseg offset
Inst0	Bypass	Data0	Inst1	Inst2	0x1211_2000	0x8300
Inst1	Bypass	Data0	Inst0	Inst2	0x1210_2000	0x8320
Inst2	Bypass	Data1	Inst3	Inst4	0x1413_2100	0x8340
Inst3	Bypass	Data1	Inst2	Inst4	0x1412_2100	0x8360
Inst4	Bypass	Data2	Inst5	Inst6	0x1615_2200	0x8380

Break	Cond0	Cond1	Cond2	Cond3	PrCndA Value	drseg offset
Inst5	Bypass	Data2	Inst4	Inst6	0x1614_2200	0x83a0
Inst6	Bypass	Data3	Inst7	Inst0	0x1017_2300	0x83c0
Inst7	Bypass	Data3	Inst6	Inst0	0x1016_2300	0x83e0
Data0	Bypass	Inst0	Inst1	Data1	0x2111_1000	0x84e0
Data1	Bypass	Inst2	Inst3	Data2	0x2213_1200	0x8500
Data2	Bypass	Inst4	Inst5	Data3	0x2315_1400	0x8520
Data3	Bypass	Inst6	Inst7	Data0	0x2017_1600	0x8540

11.2.8.3 Stopwatch Timer Control (STCtl) Register (0x8900)

The *Stopwatch Timer Control (STCtl)* register gives configuration information about how the stopwatch timer register is controlled. On the M5100 core, the break channels that control the stopwatch timer are fixed and this register is read-only. This register is required only if stopwatch timer is implemented.

Figure 11.20 STCtl Register Format




Table 11.25 STCtl Register Field Descriptions

Fields		Description	Read/Wr ite	Reset State
Name	Bit(s)			
Res	31:18	Must be written as zero; returns zero on read.	R	0
StopChan1	17:14	Indicates the instruction breakpoint channel that will stop the counter if the timer is under pair1 breakpoint control	R	0
StartChan1	13:10	Indicates the instruction breakpoint channel that will start the counter if the timer is under pair1 breakpoint control	R	0
En1	9	Enables the second pair (pair1) of breakpoint registers to control the timer when under breakpoint control. If the stopwatch timer is configured to be under breakpoint control (by setting <i>CBTControlSTM</i>) and this bit is set, the breakpoints indicated in the StartChan1 and StopChan1 fields will control the timer. The M5100 core only supports 1 pair of stopwatch control breakpoints so this field is not writable and will read as 0.	R	0
StopChan0	8:5	Indicates the instruction breakpoint channel that will stop the counter if the timer is under pair0 breakpoint control.	R	0x4

Table 11.25 STCtl Register Field Descriptions

Fields		Description	Read/Wr ite	Reset State
Name	Bit(s)			
StartChan0	4:1	Indicates the instruction breakpoint channel that will start the counter if the timer is under pair0 breakpoint control.	R	0x1
En0	0	Enables the first pair (pair0) of breakpoint registers to control the timer when under breakpoint control. If the stopwatch timer is configured to be under breakpoint control (by setting <i>CBTControl/STM</i>) and this bit is set, the breakpoints indicated in the StartChan0 and StopChan0 fields will control the timer. The M5100 core only supports 1 pair of stopwatch control breakpoints so this field is not writable and will read as 1.	R	1

11.2.8.4 Stopwatch Timer Count (STCnt) Register (0x8908)

The *Stopwatch Timer Count (STCnt)* register is the count value for the stopwatch timer. This register is required only if the stopwatch timer is implemented.

Figure 11.21 STCnt Register Format




Table 11.26 STCtl Register Field Descriptions

Fields		Description	Read/Wr ite	Reset State
Name	Bit(s)			
Count	31:0	Current counter value	R/W	0

11.3 Complex Breakpoint Usage

11.3.1 Checking for Presence of Complex Break Support

Software should verify that the complex breakpoint hardware is implemented prior to attempting to use it. The full sequence of steps is shown below for general use. Spots where the M5100 core has restricted behavior are noted.

1. Read the *Config1EP* bit to check for the presence of EJTAG logic. EJTAG logic is always present on the M5100 core.
2. Read the *DebugNoDCR* bit to check for the presence of the *Debug Control Register (DCR)*. The *DCR* will always be implemented on the M5100 core.

3. Read the *DCRCBT* bit to check for the presence of any complex break and trigger features
4. Read the *CBTControl* register to check for the presence of each individual feature. If the M5100 core implements any complex break and trigger features, it will implement all of them
5. If Pass Counters are implemented, they may not be implemented for all break channels and may have different counter sizes. To determine the size and presence of each pass counter, software can write -1 to each of the *IBPCn* and *DBPCn* registers and read it back. If the M5100 core implements pass counters, it will implement an 8b counter for each instruction breakpoint and a 16b counter for each data breakpoint.
6. If tuples are implemented, they may only be supported on a subset of the data breakpoint channels. This can be checked by seeing if the *DBBCnTUP* bit can be set to 1. Additionally, some cores may support dynamically changing which instruction breakpoint is associated with a given data breakpoint. This can be checked by attempting to write the *DBCCnTIBrkNum* field. If the M5100 core implements tuple support, it will support it for all data breakpoint channels and the instruction breakpoint association will be fixed.
7. If Priming Conditions are supported, a core may only support a subset of the possible priming condition values. This can be checked by writing to the *xBCCnPrCnd* field. If only 1 or 2 bits can be written, the available priming conditions will be described in the *PrCndA* registers. If 3 bits are writable, *PrCndA* and *PrCndB* will describe the conditions, and if all 4 bits are writable, the *PrCndA*, *PrCndB*, *PrCndC*, and *PrCndD* registers will all exist. Some cores may also support changing the priming conditions and this can be checked by attempting to write to the *PrCnd* registers. If the M5100 core supports priming conditions, it will support 4 statically mapped priming conditions per breakpoint which will be described in the *PrCndA* registers.
8. If support for qualified breakpoints is indicated, it may only be supported for some of the breakpoints. Additionally, the data breakpoint used for the qualification may be configurable. Software can check this by writing to the *xBCCnDQ* and *xBCCnDQBrkNum* fields. If the M5100 core support qualified breakpoints, it will only support it on instruction breakpoints and the data break used for qualification will be fixed for each instruction breakpoint.
9. If the stopwatch timer is implemented, either one or two pairs of instruction breakpoints may be available for controlling it and it may be possible to dynamically select which instruction breakpoints are used. This can be tested by writing to the *STCtI* register.

11.3.2 General Complex Break Behavior

There is some general complex break behavior that is common to all complex breakpoints. This behavior is described below:

- Resets to a disabled state - when the core is reset, the complex break functionality will be disabled and debug software that is not aware of complex break should continue to function normally.
- Complex break state is not updated on exceptional instructions
- Complex breakpoints are evaluated at the end of the pipeline and complex breakpoint exceptions are taken imprecisely on the following instruction.
- There is no hazard between enabling and enabled events. When an instruction causes an enabling event, the following instruction sees the enabled state and reacts accordingly.

11.3.3 Usage of Pass Counters

Pass counters specify that the breakpoint conditions must match *N* times before the breakpoint action will be enabled.

- Controlled by writing to the per-breakpoint pass counter register
- Resets to 0
- Writing to a non-zero value enables the pass counter. When enabled, each time the breakpoint conditions match, the counter will be decremented by 1. After the counter value reaches 0, the breakpoint action (breakpoint exception, trigger, or complex break enable) will occur on any subsequent matches and the counter will not decrement further. The action does not occur on the match that causes the 1->0 counter decrement.
- If the breakpoint also has priming conditions and/or data qualified specified, the pass counter will only decrement when the priming and/or qualified conditions have been met.
- If a data breakpoint is configured to be a tuple breakpoint, the data pass counter will only decrement on instructions where both the instruction and data break conditions match. The pass counter for the instruction break involved in a tuple should not be enabled if the tuple is enabled.
- When a pass counter has been enabled, it will be treated as enabled until the pass counter is explicitly written to 0. Namely, breakpoint exceptions will continue to be taken imprecisely until the pass counter is disabled by writing to 0.
- The counter register will be updated as matches are detected. The current count value can be read from the register while operating in debug mode. Note that this behavior is architecturally recommended, but not required.

11.3.4 Usage of Tuple Breakpoints

A tuple breakpoint is the logical AND of a data breakpoint and an instruction breakpoint. Tuple breakpoints are specified as a condition on a data breakpoint. If the *DBCCnTUP* bit is set, the data breakpoint will not match unless there the corresponding instruction breakpoint conditions are also met.

- Uses the data breakpoint resources to specify the break action, break status, pass counters, and priming conditions.
- The instruction breakpoint involved in the tuple should be configured as follows:
 - $IBCCn_{CBE} = 1$
 - $IBCCn_{PrCnd} = IBCCn_{DQ} = IBCCn_{TE} = IBCCn_{BE} = IBCCn = 0$

11.3.5 Usage of Priming Conditions

Priming conditions provide a way to have one breakpoint enabled by another one. Prior to the priming condition being satisfied, any breakpoint matches are ignored.

- Priming condition resets to bypass which specifies that no priming is required
- 3 other priming conditions are available for each breakpoint. These condition vary from breakpoint to breakpoint (since it makes no sense for a breakpoint to prime itself). The conditions for each of the breakpoints are listed in [Table 11.23](#).
- The priming breakpoint must have $xBCn_{TE}$ or $xBCCn_{CBE}$ set.

- When the priming condition has been seen, the primed breakpoint will remain primed until its *xBCCn* register is written
- The primed state is stored with the breakpoint being primed and not with the breakpoint that is doing the priming.
- Each Prime condition is the comparator output after it has been qualified by its own Prime condition, data qualification, and pass counter. Using this, several stages of priming are possible (e.g. data cycle D followed by instruction A followed by instruction B N times followed by instruction C).

11.3.6 Usage of Data Qualified Breakpoints

Each of the instruction breakpoints can be set to be data qualified. In qualified mode, a breakpoint will recognize its conditions only after the specified data breakpoint matches both address and data. If the data breakpoint matches address, but has a mismatch on the data value, the instruction breakpoint will be unqualified and will not match until a subsequent qualifying match.

This feature can be used similarly to the ASID qualification that is available on cores with TLBs. If an RTOS loads a process ID for the current process, that load can be used as the qualifying breakpoint. When a matching process ID is loaded (entering the desired RTOS process), qualified instruction breakpoints will be enabled. When a different process ID is loaded (leaving the desired RTOS process), the qualified instruction breakpoints are disabled. Alternatively, with the InvertValueMatch feature of the data breakpoint, the instruction breakpoints could be enabled on any process ID other than the specified one.

- The qualifying data break must have *DBCn_{TE}* or *DBCCn_{CBE}* set.
- The qualifying data break should have data comparison enabled (via settings of *DBCn_{BLM}* and *DBCn_{BAI}*)
- The qualifying data break should not have pass counters, priming conditions, or tuples enabled.
- The qualifying data access can be either a load or store, depending on the settings of *DBCn_{NoSB}* and *DBCn_{NoLB}*
- The Qualified/Unqualified state is stored with the instruction breakpoint that is being qualified. Writing its *IBCCn* register will disqualify that breakpoint.
- Qualified instruction breakpoint can also have priming conditions and/or pass counters enabled. The pass counter will only decrement when the priming and qualifying conditions have been met. The instruction breakpoint action (break, trigger, or complex enable) will only occur when all priming, qualifying, and pass counter conditions have been met.
- Qualified instruction breakpoint can be used to prime another breakpoint

11.3.7 Usage of Stopwatch Timers

The stopwatch timer is a drseg memory mapped count register. It can be configured to be free running or controlled by instruction breakpoints. This could be used to measure the amount of time that is spent in a particular function by starting the counter upon function entry and stopping it upon exit.

- Count value is reset to 0
- Reset state has counter stopped and under breakpoint control so that the counter is not running when the core is not being debugged.

- Bit in *CBTControl* register controls whether the counter is free-running or breakpoint controlled.
- Counter does not count in debug mode
- When breakpoint controlled, the involved instruction breakpoints must have *IBCn_{TE}* or *IBCCn_{CBE}* set in order to start or stop the timer.

11.4 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.
- 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.

11.4.1 EJTAG Internal and External Interfaces

The external interface of the EJTAG module consists of the 5 signals defined by the IEEE standard.

Table 11.27 EJTAG Interface Pins

Pin	Type	Description
TCK	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> .
TMS	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> .
TDI	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> .
TDO	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> .

Table 11.27 EJTAG Interface Pins (Continued)

Pin	Type	Description
TRST_N	I	<p>Test Reset Input (Optional pin)</p> <p>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></p> <p>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.</p>

11.4.2 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 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 11.22](#). 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.

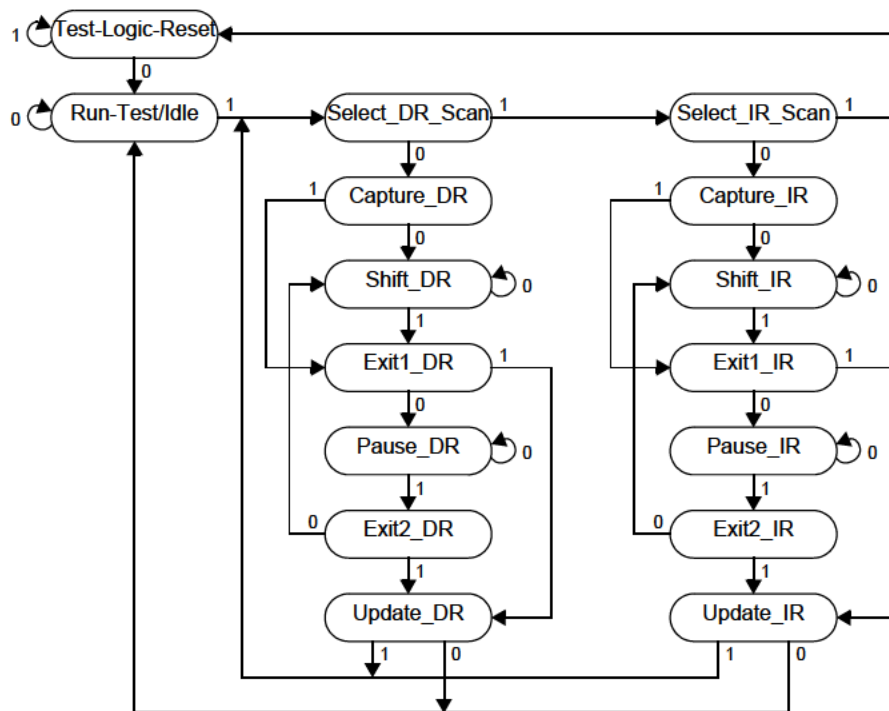
When test access is required, a protocol is applied via the *TMS* and *TCK* inputs, causing the TAP to exit the *Test-Logic-Reset* state and move through the appropriate states. From the *Run-Test/Idle* state, an Instruction register scan or a data register scan can be issued to transition the TAP through the appropriate states shown in [Figure 11.22](#).

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 11.22 TAP Controller State Diagram



11.4.2.1 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.

11.4.2.2 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.

11.4.2.3 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.

11.4.2.4 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.

11.4.2.5 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.

11.4.2.6 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.

11.4.2.7 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.

11.4.2.8 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.

11.4.2.9 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.

11.4.2.10 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.

11.4.2.11 Capture_IR State

In this state the shift register contained in the Instruction register loads a fixed pattern (00001₂) 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.

11.4.2.12 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.

11.4.2.13 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.

11.4.2.14 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.

11.4.2.15 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.

11.4.2.16 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.

11.4.3 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 11.28 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
0x14	PCSAMPLE	Selects the PCsample register
0x17	FDC	Selects Fast Debug Channel.
0x1F	BYPASS	Bypass mode

11.4.3.1 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.

11.4.3.2 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.

11.4.3.3 IMPCODE Instruction

This instruction selects the Implementation register for output, which is always 32 bits.

11.4.3.4 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.

11.4.3.5 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.


11.4.3.6 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*.

11.4.3.7 ALL Instruction

This instruction is used to select the concatenation of the Address and Data register, and the *EJTAG Control* register between *TDI* and *TDO*. It can be used in particular if switching instructions in the instruction register takes too many *TCK* cycles. The first bit shifted out is bit 0.

Figure 11.23 Concatenation of the EJTAG Address, Data and Control Registers



11.4.3.8 EJTAGBOOT Instruction

When the EJTAGBOOT instruction is given and the Update-IR state is left, then the reset values of the ProbTrap, ProbEn and *EjtagBrk* bits in the *EJTAG Control* register are set to 1 after a hard or soft reset.

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.

It is possible to make the CPU go into debug mode just after a hard or soft reset, without fetching or executing any instructions from the normal memory area. This can be used for download of code to a system which have no code in ROM.

The Bypass register is selected when the EJTAGBOOT instruction is given.

11.4.3.9 NORMALBOOT Instruction


When the NORMALBOOT instruction is given and the Update-IR state is left, then the reset value of the ProbTrap, ProbEn and *EjtagBrk* bits in the *EJTAG Control* register are set to 0 after hard or soft reset.

The Bypass register is selected when the NORMALBOOT instruction is given.

11.4.3.10 FASTDATA Instruction

This selects the Data and the Fastdata registers at once, as shown in [Figure 11.24](#).

Figure 11.24 TDI to TDO Path When in Shift-DR State and FASTDATA Instruction is Selected



11.4.3.11 PCsample Register (PCSAMPLE Instruction)

This selects the PCsample Register. The use of the PCsample Register is described in more detail in [Section 11.9 “PC/Data Address Sampling”](#)

11.4.3.12 FDC Instruction

This selects the Fast Debug Channel. The use of the FDC is described in more detail in [Section 11.10 “Fast Debug Channel”](#).

11.4.3.13 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.

11.4.3.14 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.

11.4.3.15 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.

11.5 EJTAG TAP Registers

The EJTAG TAP Module has one Instruction register and a number of data registers, all accessible through the TAP:

11.5.1 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 00001₂, 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 11.28](#).

11.5.2 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)*
- Processor Access Address Register
- Processor Access Data Register
- FastData Register

11.5.2.1 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.

11.5.2.2 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 11.29](#) 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.

Figure 11.25 Device Identification Register Format

31	28 27	12 11	1 0
Version	PartNumber	ManufID	R

Table 11.29 Device Identification Register

Fields		Description	Read/Write	Reset State
Name	Bit(s)			
Version	31:28	Version (4 bits) This field identifies the version number of the processor derivative.	R	<i>EJ_Version[3:0]</i>
PartNumber	27:12	Part Number (16 bits) This field identifies the part number of the processor derivative.	R	<i>EJ_PartNumber[15:0]</i>
ManufID	11:1	Manufacturer Identity (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

11.5.2.3 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.

Figure 11.26 Implementation Register Format




Table 11.30 Implementation Register Descriptions

Fields		Description	Read/Wr ite	Reset State						
Name	Bit(s)									
EJTAGver	31:29	EJTAG Version. 2: Version 2.6	R	5						
Reserved	28:25	Reserved	R	0						
DINTsup	24	<i>DINT</i> Signal Supported from Probe This bit indicates if the <i>DINT</i> signal from the probe is supported: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td><i>DINT</i> signal from the probe is not supported</td> </tr> <tr> <td>1</td> <td>Probe can use <i>DINT</i> signal to make debug interrupt.</td> </tr> </tbody> </table>	Encoding	Meaning	0	<i>DINT</i> signal from the probe is not supported	1	Probe can use <i>DINT</i> signal to make debug interrupt.	R	<i>EJ_DINTsup</i>
Encoding	Meaning									
0	<i>DINT</i> signal from the probe is not supported									
1	Probe can use <i>DINT</i> signal to make debug interrupt.									

Table 11.30 Implementation Register Descriptions

Fields		Description	Read/Wr ite	Reset State										
Name	Bit(s)													
ASIDsize	23:21	Size of ASID field in implementation: <table border="1" style="margin-left: 20px;"> <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>6-bit ASID</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	6-bit ASID	2	8-bit ASID	3	Reserved	R	0
Encoding	Meaning													
0	No ASID in implementation													
1	6-bit ASID													
2	8-bit ASID													
3	Reserved													
Reserved	20:17	Reserved	R	0										
MIPS16	16	Indicates whether MIPS16 is implemented: <table border="1" style="margin-left: 20px;"> <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	0				
Encoding	Meaning													
0	No MIPS16 support													
1	MIPS16 implemented													
Reserved	15	Reserved	R	0										
NoDMA	14	No EJTAG DMA Support	R	1										
Reserved	13:0	Reserved	R	0										

11.5.2.4 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 both hard and soft CPU resets, but not on TAP controller resets by e.g. *TRST_N*. *TCK* clock is not required when the hard or soft CPU reset occurs, but the bits are still updated to the reset value when the *TCK* applies. The first 5 *TCK* clocks after hard or soft CPU resets may result in reset of the bits, due to synchronization between clock domains.

Figure 11.27 EJTAG Control Register Format

31	30	29	28	23	22	21	20	19	18	17	16	15	14	13	12	11	4	3	2	0
Rocc	Psz	Res	Doze	Halt	PerRst	PRnW	PrAcc	Res	PrRst	ProbEn	ProbTrap	Res	Ejtag-Brk	Res	DM	Rs				

Table 11.31 EJTAG Control Register Descriptions

Fields		Description	Read/Write	Reset State																																	
Name	Bit(s)																																				
Rocc	31	<p>Reset Occurred The bit indicates if a CPU reset has occurred:</p> <table border="1"> <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 <i>Rocc</i> 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 <i>EJTAG Control</i> register is not updated in the <i>Update-DR</i> state unless <i>Rocc</i> is 0, or written to 0. This is in order to ensure proper handling of processor access.</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.																																				
Psz[1:0]	30:29	<p>Processor Access Transfer Size 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	28:23	Reserved	R	0																																	

Table 11.31 EJTAG Control Register Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
Doze	22	<p>Doze state</p> <p>The Doze bit indicates any kind 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.									
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 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_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									

Table 11.31 EJTAG Control Register Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
PrAcc	18	<p>Processor Access (PA) 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. Write of 1 is ignored. 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. Finishing a Processor Access is not accepted while the <i>ROCC</i> bit is set. This is to avoid that a Processor Access occurring after the reset is finished due to indication of a Processor Access that occurred before the reset. 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						
PrRst	16	<p>Processor Reset (implementation-dependent behavior) 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. 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. 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 hard or soft reset.</p>	R/W	0						

Table 11.31 EJTAG Control Register Descriptions (Continued)

Fields		Description	Read/ Write	Reset State												
Name	Bit(s)															
ProbEn	15	<p>Probe Enable This bit indicates to the CPU if the EJTAG memory is handled by the probe so processor accesses are answered:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>The probe does not handle EJTAG memory transactions</td> </tr> <tr> <td>1</td> <td>The 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. The <i>ProbEn</i> bit is reflected as a read-only bit in the <i>ProbEn</i> bit, bit 0, in the <i>Debug Control Register (DCR)</i>.</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 <i>ProbEn</i> 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 EJTAG-BOOT 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>Processor is in non-debug mode (No EJTAGBOOT indication given)</td> </tr> <tr> <td>1</td> <td>Processor is in debug mode (EJTAG-BOOT indication given)</td> </tr> </tbody> </table>	Encoding	Meaning	0	The probe does not handle EJTAG memory transactions	1	The probe does handle EJTAG memory transactions	Encoding	Meaning	0	Processor is in non-debug mode (No EJTAGBOOT indication given)	1	Processor is in debug mode (EJTAG-BOOT indication given)	R/W	0 or 1 from EJTAGBOOT
Encoding	Meaning															
0	The probe does not handle EJTAG memory transactions															
1	The probe does handle EJTAG memory transactions															
Encoding	Meaning															
0	Processor is in non-debug mode (No EJTAGBOOT indication given)															
1	Processor is in debug mode (EJTAG-BOOT indication given)															

Table 11.31 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 0xBFC0.0480</td> </tr> <tr> <td>1</td> <td>In EJTAG memory at 0xFF20.0200 in dmseg</td> </tr> </tbody> </table> <p>Valid setting of the <i>ProbTrap</i> bit depends on the setting of the <i>ProbEn</i> bit, as described for the <i>ProbEn</i> bit. The <i>ProbTrap</i> should not be set to 1, for debug exception vector in EJTAG memory, unless the <i>ProbEn</i> bit is also set to 1 to indicate that the EJTAG memory may be accessed.</p> <p>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 <i>ProbTrap</i> bit prior to setting the <i>EjtagBrk</i> bit will have effect for the <i>EjtagBrk</i>.</p> <p>The reset value of the bit depends on whether the EJTAG-BOOT 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>Processor is in non-debug mode (No EJTAGBOOT indication given)</td> </tr> <tr> <td>1</td> <td>Processor is in debug mode (EJTAG-BOOT indication given)</td> </tr> </tbody> </table>	Encoding	Meaning	0	In normal memory 0xBFC0.0480	1	In EJTAG memory at 0xFF20.0200 in dmseg	Encoding	Meaning	0	Processor is in non-debug mode (No EJTAGBOOT indication given)	1	Processor is in debug mode (EJTAG-BOOT indication given)	R/W	0 or 1 from EJTAGBOOT
Encoding	Meaning															
0	In normal memory 0xBFC0.0480															
1	In EJTAG memory at 0xFF20.0200 in dmseg															
Encoding	Meaning															
0	Processor is in non-debug mode (No EJTAGBOOT indication given)															
1	Processor is in debug mode (EJTAG-BOOT indication given)															
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.</p> <p>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.</p> <p>The reset value of the bit depends on whether the EJTAG-BOOT 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>Processor is in non-debug mode (No EJTAGBOOT indication given)</td> </tr> <tr> <td>1</td> <td>Processor is in debug mode (EJTAG-BOOT indication given)</td> </tr> </tbody> </table>	Encoding	Meaning	0	Processor is in non-debug mode (No EJTAGBOOT indication given)	1	Processor is in debug mode (EJTAG-BOOT indication given)	R/W1	0 or 1 from EJTAGBOOT						
Encoding	Meaning															
0	Processor is in non-debug mode (No EJTAGBOOT indication given)															
1	Processor is in debug mode (EJTAG-BOOT indication given)															
Res	11:4	Reserved	R	0												

Table 11.31 EJTAG Control Register Descriptions (Continued)

Fields		Description	Read/Write	Reset State						
Name	Bit(s)									
DM	3	Debug Mode This bit indicates the debug or non-debug mode: <table border="1" data-bbox="553 415 1055 533"> <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> The bit is sampled in the <i>Capture-DR</i> state of the TAP controller.	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						

11.5.3 Processor Access Address Register

The Processor Access Address (*PAA*) 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.

11.5.3.1 Processor Access Data Register

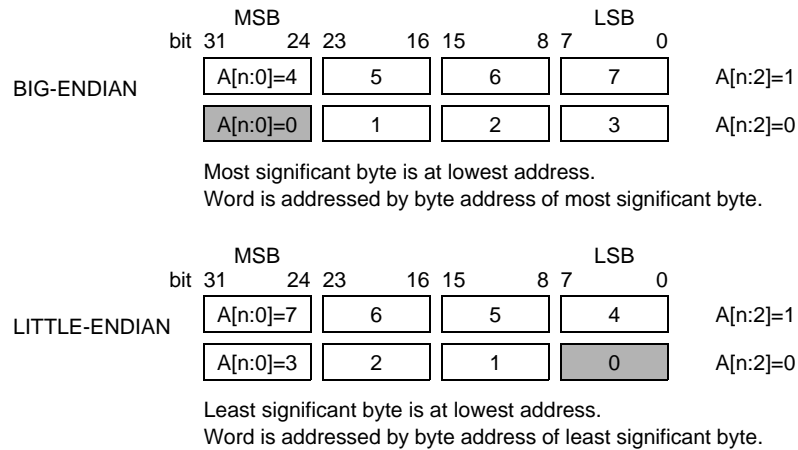
The *Processor Access Data (PAD)* register is used to provide data value to and from a processor access. The length of the Data register is 32 bits, and this register is selected by shifting in the DATA instruction.

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, and the register should only be updated with a new value when a processor access write is pending.

The *PAD* register is 32 bits wide. Data alignment is not used for this register, so the value in the *PAD* register matches data on the internal bus. The undefined bytes for a PA write are undefined, and for a *PAD* read then 0 (zero) must be shifted in for the unused bytes.

The organization of bytes in the *PAD* register depends on the endianness of the core, as shown in [Figure 11.28](#). The endian mode for debug/kernel mode is determined by the state of the *SI_Endian* input at power-up.

Figure 11.28 Endian Formats for the PAD Register



The size of the transaction and thus the number of bytes available/required for the *PAD* register is determined by the *Psz* field in the *ECR*.

11.5.4 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 11.29 Fastdata Register Format



Table 11.32 Fastdata Register Field Description

Fields		Description	Read/Write	Power-up State
Name	Bits			
SPrAcc	0	Shifting in a zero value requests completion of the <i>Fastdata</i> access. The <i>PrAcc</i> bit in the <i>EJTAG Control</i> register is overwritten with zero when the access succeeds. (The access succeeds if <i>PrAcc</i> is one and the operation address is in the legal <i>dmseg</i> <i>Fastdata</i> area.) When successful, a one is shifted out. Shifting out a zero indicates a <i>Fastdata</i> access failure. Shifting in a one does not complete the <i>Fastdata</i> access and the <i>PrAcc</i> 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

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 11.33 shows the values of the *PrAcc* and *SPrAcc* bits and the results of a Fastdata access.

Table 11.33 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 FASTDATA	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
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.

11.6 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 soft or hard reset.

11.6.1 Fetch/Load and Store from/to EJTAG Probe Through dmseg

1. The internal hardware latches the requested address into the PA 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 PA Address register and shifts out the requested address.
6. The EJTAG Probe selects the PA 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 PA Address register
2. The internal hardware latches the data to be written into the PA 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 PA Address register and shifts out the requested address.
7. The EJTAG Probe selects the PA Data register and shifts out the data to be 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 requested 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.

Note: Probe accesses and external bus accesses are serialized by the core. A probe access will not begin until all external bus requests have completed. Similarly, a new probe or external bus access will not begin until a pending probe access has completed.

11.7 SecureDebug

For security reasons, users can optionally disable certain EJTAG capabilities via the SecureDebug feature in order to prevent untrusted access to the core through debug mode.

11.7.1 Disabling EJTAG Debugging

11.7.1.1 EJ_DisableProbeDebug Signal

An input signal to the core is defined, *EJ_DisableProbeDebug*, which when asserted, forces *ProbEn*=0 and *ProbTrap*=0. *EJ_DisableProbeDebug* overrides any other *ProbEn* or *ProbTrap* settings.

Suggested implementation of the *EJ_DisableProbeDebug* signal is for a microcontroller to provide a bit within non-volatile memory (outside the core) that is pre-programmed to set or clear this control signal.

Table 11.34 EJ_DisableProbeDebug Signal Overview

Signal	Description	Direction	Compliance
EJ_DisableProbeDebug	When asserted: <ul style="list-style-type: none"> • <i>ProbEn</i> = 0 • <i>ProbTrap</i> = 0 • <i>EjtagBrk</i> is disabled¹. • EJTAGBOOT is disabled. • PC Sampling is disabled. • <i>DINT</i> signal is ignored¹. 	Input	Required for SecureDebug

1. An override is provided.

11.7.1.2 Override for EjtagBrk and DINT Disable

An override for the *EjtagBrk* and *DINT* disable caused by the *EJ_DisableProbeDebug* signal is provided by the Memory Protection Unit (MPU) *Config* register field *EjtagBrk_Override*. This override feature is only available if the Memory Protection Unit is implemented.

The override can be asserted by the CPU during the trusted boot process. Its purpose is to allow a probe to assert *EjtagBrk* or the assertion of the *DINT* signal, which requests a Debug Interrupt exception, thus providing a means of recovering the CPU from a crash or hang. This feature allows a Debug Executive, if one is provided in target firmware, to communicate with the probe over the Fast Debug Channel (FDC) in order to get the attention of the target by causing a debug exception. It also allows a host-based debugger to query the target via Debug Executive commands to determine the cause of the hang.

11.7.2 EJTAG Features Unmodified by SecureDebug

SecureDebug will not modify the following EJTAG features:

- FDC (Fast Debug Channel) over EJTAG. This is required to provide a path for an EJTAG probe to send and receive messages via the Debug Executive when one is included in the target code. The physical EJTAG serial connection, pins, and protocol must function correctly as well as a cJTAG (2-wire) connection for FDC.
- *RST** signal. This is the hardware signal on the EJTAG connector that connects to the target system reset circuit. It can be asserted by an EJTAG probe.

11.8 iFlowtrace™ Mechanism

There is only one optional trace mechanism that is available to extract additional information about program execution. iFlowtrace is a light-weight instruction-only tracing scheme that is sufficient to reconstruct the execution flow in the core and it can only be controlled by debug software. This tracing scheme has been kept very simple to minimize the impact on die size.

The iFlowtrace tracing scheme is not a strict subset of the PDtrace tracing methodology, and its trace format outputs differ from those of PDtrace. Trace formats, using simplified instruction state descriptors, were designed for the iFlowtrace trace to simplify the trace mechanism and to obtain better compression.

Tracing is disabled if the processor enters Debug Mode (refer to the EJTAG specification for description of Debug Mode). This is true for both Normal Trace Mode as well as Special Trace Mode.

The presence of the iFlowtrace mechanism is indicated by the *CPO Config3_{ITL}* register bit.

11.8.1 A Simple Instruction-Only Tracing Scheme

A trace methodology can often be mostly defined by its inputs and outputs. Hence this basic scheme is described by the inputs to the core tracing logic and by the trace output format from the core. We assume here that the execution flow of the program is traced at the end of the execution path in the core similar to PDtrace.

11.8.1.1 Trace Inputs

1. *In_TraceOn*: when on, legal trace words are coming from the core and at the point when it is turned on, that is for the first traced instruction, a full PC value is output. When off, it cannot be assumed that legal trace words are available at the core interface.
2. *In_Stall*: This says, stall the processor to avoid buffer overflow that can lose trace information. When off, a buffer overflow will simply throw away trace data and start over again. When on, the processor is signalled from the tracing logic to stall until the buffer is sufficiently drained and then the pipeline is restarted.

11.8.1.2 Normal Trace Mode Outputs

1. Stall cycles in the pipe are ignored by the tracing logic and are not traced. This is indicated by the signal *Out_Valid* that is turned off when no valid instruction is being traced. When *Out_Valid* is asserted, instructions are traced out as described in the rest of this section. The traced instruction PC is a virtual address.
2. In the output format, every sequentially executed instruction is traced as 1'b0.
3. Every instruction that is not sequential to the previous one is traced as either a 10 or an 11 (read this as a serial bitstream from left to right). This implies that the target instruction of a branch or jump is traced this way, not the actual branch or jump instruction (this is similar to PDtrace):
4. A 10 instruction implies a taken branch for a conditional branch instruction whose condition is unpredictable statically, but whose branch target can be computed statically and hence the new PC does not need to be traced out. Note that if this branch was not taken, it would have been indicated by a 0 bit, that is sequential flow.
5. A 11 instruction implies a taken branch for an indirect jump-like instruction whose branch target could not be computed statically and hence the taken branch address is now given in the trace. This includes, for example, instructions like jr, jalr, and interrupts:
 - 11 00 - followed by 8 bits of 1-bit shifted offset from the last PC. The bit assignments of this format on the bus between the core tracing logic and the ITCB is:

$$\begin{aligned} [3:0] &= 4'b0011 \\ [11:4] &= PC\delta[8:1] \end{aligned}$$
 - 11 01 - followed by 16 bits of 1-bit shifted offset from the last PC. The bit assignments of this format on the bus between the core tracing logic and the ITCB is:

[3:0] = 4'b1011
[19:4] = PCdelta[16:1]

- 11 10 - followed by 31 of the most significant bits of the PC value, followed by a bit (NCC) that indicates no code compression. Note that for a MIPS32 or MIPS64 instruction, NCC=1, and for microMIPS instruction NCC=0. This trace record will appear at all transition points between MIPS32/MIPS64 and microMIPS instruction execution.

This form is also a special case of the 11 format and it is used when the instruction is not a branch or jump, but nevertheless the full PC value needs to be reconstructed. This is used for synchronization purposes, similar to the Sync in PDtrace. In iFlowtrace rev 2.0 onwards, the sync period is user-defined, and is counted down and when an internal counter runs through all the values, this format is used. The bit assignments of this format on the bus between the core tracing logic and the ITCB is:

[3:0] = 4'b0111
[34:4] = PC[31:1]
[35] = NCC

- 11 11 - Used to indicate trace resumption after a discontinuity occurred. The next format is a 1110 that sends a full PC value. A discontinuity might happen due to various reasons, for example, an internal buffer overflow, and at trace-on/trace-off trigger action.

11.8.2 Special Trace Modes

iFlowtrace 2.0 adds special trace modes which can only be active when the normal tracing mode is disabled. Software can determine which modes are supported by attempting to write the enable bits in the *IFCTL* register. Software can check the Illegal bit in the *IFCTL* register—if an unsupported combination of modes is requested, the bit will be set and the trace contents will be unpredictable. The special trace modes are described below.

11.8.2.1 Mode Descriptions

Delta Cycle Mode

This mode is specified in combination with the other special trace modes. It is enabled via the CYC bit in the Control/Status Register. When delta cycle reporting is enabled, each trace message will include a 10b delta cycle value which reports the number of cycles that have elapsed since the last message was generated. A value of 0 indicates that the two messages were generated in the same cycle. A value of 1 indicates that they were generated in consecutive cycles. If 1023 cycles elapse without an event being traced, a counter rollover message is generated.

Note: If the processor clocks stop due to execution of the WAIT instruction, the delta cycle counter will also stop and will report 'active' cycles between events rather than 'total' cycles.

Breakpoint Match Mode

This mode uses EJTAG data and instruction breakpoint hardware to enable a trace of PC values. Instead of starting or stopping trace, a triggerpoint will cause a single breakpoint match trace record. This record indicates that there was a triggerpoint match, the breakpoint ID of the matching breakpoint, and the PC value of an instruction that matched the instruction of data breakpoint. This mode can only be used when normal tracing mode is turned off. This mode can not be used in conjunction with other special trace modes. This mode is enabled or disabled via the BM field in the Control/Status register (see [Section 11.8.6 “ITCB Register Interface for Software Configurability”](#)).

The breakpoints used in this mode must have the TE bit set to enable the match condition.

Software should avoid setting up overlapping breakpoints. The behavior when multiple matches occur on the same instruction is to report a BreakpointID of 7.

Filtered Data Tracing Mode

This mode uses EJTAG data breakpoint hardware to enable a trace of data values. Rather than starting or stopping trace as in normal trace mode, a data triggerpoint will cause a filtered data trace record. This record indicates that there was a data triggerpoint match, the breakpoint ID of the matching breakpoint, whether it was a load or store, the size of the request, low order address bits, and the data value. This mode can only be used when normal tracing mode is turned off. This mode can not be used in conjunction with other special trace modes. This mode can be enabled or disabled via the FDT bit in the Control/Status register (see [Section 11.8.6 “ITCB Register Interface for Software Configurability”](#)).

The corresponding data breakpoint must have the TE bit set to enable the match condition.

Software should avoid setting up overlapping data breakpoints. The behavior when multiple matches on one load or store are detected is to report a BreakpointID of 7.

Extended Filtered Data Tracing Mode

Extends Filtered Data Tracing Mode by adding the virtual address of the load/store instruction to the generated trace information. (see [Section “Filtered Data Tracing Mode”](#) above).

This behavior is enabled/disabled by the FDT_CAUSE field in the *IFCTL* Control/Status register (see [Section 11.8.6 “ITCB Register Interface for Software Configurability”](#)). FDT_CAUSE only has effect if the FDT field is also set.

The extended trace sequence is a FDT trace message followed by the Breakpoint Match (BM) trace message. If the *IFCTL_CYC* field is set, the FDTtrace message will have a DeltaCycle Message value of ‘0’ directly followed by the Breakpoint Match message. This message sequence (FDT, delta cycle of 0, and BM) indicates to the trace disassembler that Extended Filtered Data Tracing mode is enabled (*IFCTL_FDT_CAUSE*=1).

Function Call/Return and Exception Tracing Mode

In this mode, the PC value of function calls and returns and/or exceptions and returns are traced out. This mode can only be used when normal tracing mode is turned off. This mode cannot be used in conjunction with other special trace modes. The function call/return and exception/return are independently enabled or disabled via the FCR and ER bits in the Control//Status register (see [Section 11.8.6 “ITCB Register Interface for Software Configurability”](#)).

These events are reported for the following instructions:

- MIPS32 function calls: JAL, JALR, JALR.HB, JALX
- microMIPS function calls: JAL, JALR, JALR.HB, JALX, JALR16, JALRS16, JALRS, JALRS.HB, JALS
- MIPS32 function returns: JR, JR.HB
- microMIPS function returns: JR, JR.HB, JRC, JRADDIUSP, JR16
- Exceptions: Reported on the first instruction of the exception handler
- Exception returns: ERET

- MCU ASE Interrupt returns: IRET

Other Trace Messages

In any of the special trace modes, it is possible to embed messages into the trace stream directly from a program. This is done by writing to the *UserTraceData1* or *UserTraceData2* Cop0 registers. When *UserTraceData1* register is written, a trace message of type “User Triggered Message 1” (UTM1) is generated. When *UserTraceData2* register is written, a trace message of type “User Triggered Message 2” (UTM2) is generated. Please refer to 8.2.49 “User Trace Data1 Register (CP0 Register 23, Select 3)/User Trace Data2 Register (CP0 Register 24, Select 3)” on page 229.

Overflow messages can also be generated when tracing off-chip if the IO control bit is 0 and trace data is generated faster than it is consumed. No overflow will be generated when using on-chip trace.

11.8.2.2 Special Trace Mode Outputs

The normal and special trace modes cannot be enabled at the same time because the trace message encoding is not unique between the two modes. The software reading the trace stream must be aware of which mode is selected to know how to interpret the bits in the trace stream. The message types for each type of special trace message are unique.

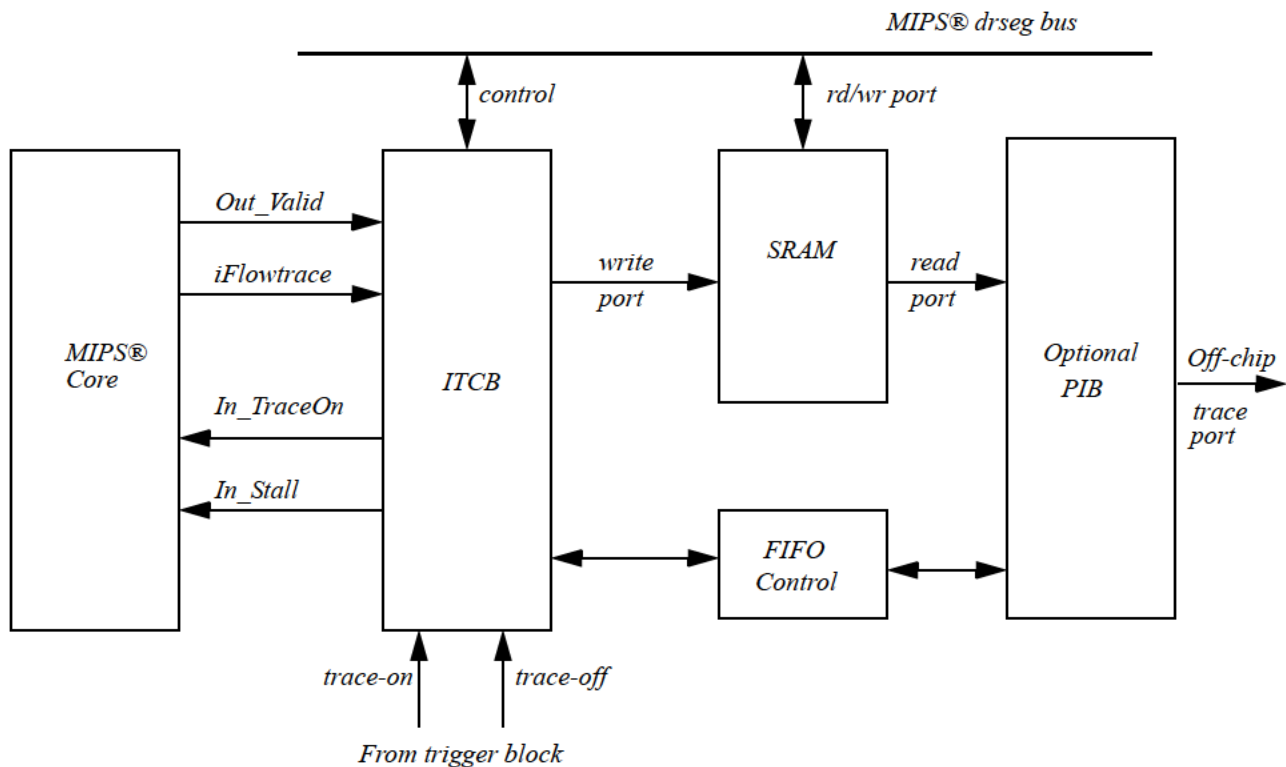
- 00 (as above, read a bitstream from left to right) - Delta Cycle Rollover message. The output format is:
[1:0] = 2'b00
- 010 - User Trace Message. The format of this type of message is:
[2:0] = 3'b010
[34:3] = Data[31:0]
[35] = UTM2/UTM1 (1=UTM2, 0=UTM1)
[44:36] = DeltaCycle (if enabled)
- 011 - Reserved
- 10 - Breakpoint Match Message. The output format during this trace mode is:
[1:0] = 2'b01
[5:2] = BreakpointID
[6] = Instruction Breakpoint
[37:7] = MatchingPC[31:1]
[38] = NCC
[48:39] = DeltaCycle (if enabled)
Note that for a MIPS32 or MIPS64 instruction, NCC=1, and for microMIPS instruction NCC=0.
- 110 - Filtered Data Message. The output format during this trace mode is:
[2:0] = 3'b011
[6:3] = BreakpointID
[7] = Load/Store (1=Load, 0=Store)
[8] = FullWord (1=32b data, 0=<32b)
[14:5] = Addr[7:2]
[46:15] = {32b data value} OR
[46:15] = {BE[3:0], 4'b0, 24b data value} OR
[46:15] = {BE[3:0], 12'b0, 16b data value} OR
[46:15] = {BE[3:0], 20'b0, 8b data value}
[56:47] = DeltaCycle (if enabled)

- 1110 - Function Call/Return/Exception Tracing. The output format during this trace mode is:
 [3:0] = 4'b0111
 [4] = FC
 [5] = Ex
 [6] = R
 [37:8] = PC[31:1]
 [38] = NCC
 [48:39] = Delta Cycle (if enabled)
 Note that for a MIPS32 or MIPS64 instruction, NCC=1, and for microMIPS instruction NCC=0. FC=1 implies a function call, Ex=1 implies the start of an exception handler, and R=1 implies a function or exception return.
- 1111- Overflow message. The format of this type of message is:
 [3:0] = 4'b1111

11.8.3 ITCB Overview

The iFlowtrace Control Block (ITCB) is responsible for accepting trace signals from the CPU core, formatting them, and storing them into an on-chip FIFO. The figure also shows the Probe Interface Block (PIB) which reads the FIFO and outputs the memory contents through a narrow off-chip trace port.

Figure 11.30 Trace Logic Overview



11.8.4 ITCB iFlowtrace Interface

The iFlowtrace interface consists of 57 data signals plus a valid signal. The 57 data signals encode information about what the CPU is doing in each clock cycle. Valid indicates that the CPU is executing an instruction in this cycle and

therefore the 57 data signals carry valid execution information. The iFlowtrace data bus is encoded as shown in [Table 11.35](#). Note that all the non-defined upper bits of the bus are zeroes.

Table 11.35 Data Bus Encoding

Valid	Data (LSBs)	Description
0	X	No instructions executed in this cycle
1	0	Normal Mode: Sequential instruction executed
1	01	Normal Mode: Branch executed, destination predictable from code
1	<8>0011	Normal Mode: Discontinuous instruction executed, PC offset is 8 bit signed offset
1	<16>1011	Normal Mode: Discontinuous instruction executed, PC offset is 16 bit signed offset
1	<NCC><31>0111	Normal Mode: Discontinuous instruction or synchronization record, No Code Compression (NCC) bit included as well as 31 MSBs of the PC value
1	00	Special Mode: Delta Cycle Rollover message
1	<10><32>010	Special Mode: User add-in Trace Message. 32 bit user data as well as 10 bit delta cycle if enabled.
1	<10><NCC><31><1><4>01	Special Mode: Breakpoint Match Message. 4-bit breakpoint ID, 1 bit indicates breakpoint type, 31 MSBs of the PC value, NCC bit included as well as 10-bit delta cycle if enable.
1	<10><32><6><1><1><4>011	Special Mode: Filtered Data Message. 4 bit breakpoint ID, 1 bit load or store indication, 1 bit full word indication, 6 bit of addr[7:2], 32 bit of the data information included as well as 10 bit delta cycle if enabled.
1	<10><NCC><31><R><Ex><FC>011	Special Mode: Function Call/Return/Exception Tracing. 1 bit function call indication, 1 bit exception indication, 1 bit function or exception return indication, 31 MSBs of the PC value, NCC bit included as well as 10 bit delta cycle if enabled.
1	1111	Internal overflow

11.8.5 TCB Storage Representation

Records from iFlowtrace are inserted into a memory stream exactly as they appear in the iFlowtrace data output. Records are concatenated into a continuous stream starting at the LSB. When a trace word is filled, it is written to memory along with some tag bits. Each record consists of a 64-bit word, which comprises 58 message bits and 6 tag bits or header bits that clarify information about the message in that word.

The ITCB includes a 58-bit shift register to accumulate trace messages. When 58 or more bits are accumulated, the 58 bits and 6 tag bits are sent to the memory write interface. Messages may span a trace word boundary; in this case, the 6 tag bits indicate the bit number of the first full trace message in the 58-bit data field.

The tag bits are slightly encoded so they can serve a secondary purpose of indicating to off-chip trace hardware when a valid trace word transmission begins. The encoding ensures that at least one of the 4 LSBs of the tag is always a 1 for a valid trace message. The tag values are shown in [Table 11.36](#). The longest trace message is 57 bits (filtered data trace in special trace mode with delta cycle), so the starting position indicated by the tag bits is always between 0 and 56.

Table 11.36 Tag Bit Encoding

Starting Bit of First Full Trace Message	Encoding (decimal)
0	58

Table 11.36 Tag Bit Encoding

Starting Bit of First Full Trace Message	Encoding (decimal)
16	59
32	60
48	61
Unused	0,16,32,48
Reserved	62,63
Others	StartingBit

When trace stops (ON set to zero), any partially filled trace words are written to memory. Any unused space above the final message is filled with 1's. The decoder distinguishes 1111 patterns used for fill in this position from an 1111 overflow message by recognizing that it is the last trace word.

These trace formats are written to a trace memory that is either on-chip or off-chip. No particular size of SRAM is specified; the size is user selectable based on the application needs and area trade-offs. Each trace word can typically store about 20 to 30 instructions in normal trace mode, so a 1 KWord trace memory could store the history of 20K to 30K executed instructions.

The on-chip SRAM or trace memory is written continuously as a circular buffer. It is accessible via drseg address mapped registers. There are registers for the read pointer, write pointer, and trace word. The write pointer register includes a wrap bit that indicates that the pointer has wrapped since the last time the register was written. Before starting trace, the write pointer would typically be set to 0. To read the trace memory, the read pointer should be set to 0 if there has not been a wrap, or to the value of the write pointer if there has been. Reading the trace word register will read the entry pointed to by the read pointer and will automatically increment the read pointer. Software can continue reading until all valid entries have been read out.

11.8.6 ITCB Register Interface for Software Configurability

The ITCB includes a drseg memory interface to allow software to set up tracing and read the current status. If an on-chip trace buffer is also implemented, there are additional registers included for accessing it.

11.8.6.1 iFlowtrace Control/Status (IFCTL) Register (offset 0x3fc0)

The Control/Status register provides the mechanism for turning on the different trace modes. [Figure 11.31](#) has the format of the register and [Table 11.37](#) describes the register fields.

Figure 11.31 Control/Status Register

31	30	16	15	14	13	12	11	10	9	8	5	4	3	2	1	0
Illegal	0	FDT_CAUSE	CYC	FDT	BM	ER	FCR	EST	SyP	OfClk	OfC	IO	En	On		

Table 11.37 Control/Status Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
0	30:16	Reserved for future use. Read as zeros, must be written as zeros	R	0	Required
Illegal	31	This bit is set by hardware and indicates if the currently enabled trace output modes are an illegal combination. A value of 1 indicates an unsupported setting. A value of 0 indicates that the currently selected settings are legal.	R	0	Required
FDT_CAUSE	15	Extended Filtered Data Trace mode (FDT). Adds causing load/store virtual address to Filtered Data Trace. FDT_CAUSE only has effect if FDT is set. The extended trace sequence is a FDT trace message followed by the Breakpoint Match (BM) trace message. If CYC is set, the FDT trace message will have a DeltaCycle Message value of '0' directly followed by the Breakpoint match (BM) message. This message sequence (FDT, delta cycle of 0, and BM) indicates to the trace disassembler that Extended Filtered Data Tracing mode is enabled.	R/W	0	Optional for iFlowtrace rev 2.0+
CYC	14	Delta Cycle Mode: This mode can be set in combination with the EST special trace modes. When set, a delta cycle value is included in each of the trace messages and indicates the number of cycles since the last message was generated. If this tracing mode is not implemented, the field is read-only and read as zero.	R/W	0	Optional for iFlowtrace rev 2.0+
FDT	13	Filtered Data Trace mode. If set, on a data breakpoint match, the data value of the matching breakpoint is traced. Normal tracing is inhibited when this mode is active. If this tracing mode is not implemented, the field is read-only and read as zero.	R/W	0	Optional for iFlowtrace rev 2.0+
BM	12	Breakpoint Match. If set, only instructions that match instruction or data breakpoints are traced. Normal tracing is inhibited when this mode is active. If this tracing mode is not implemented, the field is read-only and read as zero.	R/W	0	Optional for iFlowtrace rev 2.0+
ER	11	Trace exceptions and exception returns. If set, trace includes markers for exceptions and exception returns. Can be used in conjunction with the FCR bit. Inhibits normal tracing. If this tracing mode is not implemented, the field is read-only and read as zero.	R/W	0	Optional for iFlowtrace rev 2.0+
FCR	10	Trace Function Calls and Returns. If set, trace includes markers for function calls and returns. Can be used in conjunction with the ER bit. If this tracing mode is not implemented, the field is read-only and read as zero.	R/W	0	Optional for iFlowtrace rev 2.0+
EST	9	Enable Special Tracing Modes. If set, normal tracing is inhibited, allowing the user to choose one of several special tracing modes. Setting this bit inhibits normal trace mode. If no special tracing modes are implemented, this field is read-only, and read as zero.	R/W	0	Optional for iFlowtrace rev 2.0+
SyP	8:5	Synchronization Period. The synchronization period is set to $2^{(SyP+8)}$ instructions. Thus a value of 0x0 implies 256 instructions, and a value of 0xF implies 8M instructions.	R/W	0	Required for iFlowtrace rev 2.0+

Table 11.37 Control/Status Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
OfClk	4	Controls the Off-chip clock ratio. When the bit is set, this implies 1:2, that is, the trace clock is running at 1/2 the core clock, and when the bit is clear, implies 1:4 ratio, that is, the trace clock is at 1/4 the core clock. Ignored unless OfC is also set.	R/W	0	Required
OfC	3	Off-chip. 1 enables the PIB (if present) to unload the trace memory. 0 disables the PIB and would be used when on-chip storage is desired or if a PIB is not present. This bit is settable only if the design supports both on-chip and off-chip modes. Otherwise is a read-only bit indicating which mode is supported.	R/W or R	Preset	Required
IO	2	Inhibit overflow. If set, the CPU is stalled whenever the trace memory is full. Ignored unless OfC is also set.	R/W	0	Required
En	1	Trace enable. This bit may be set by software or by Trace-on/Trace-off action bits from the Complex Trigger block. Software writes EN with the desired initial state of tracing when the ITCB is first turned on and EN is controlled by hardware thereafter. EN turning on and off does not flush partly filled trace words.	R/W	0	Required
On	0	Software control of trace collection. 0 disables all collection and flushes out any partially filled trace words.	R/W	0	Required

11.8.6.2 ITCBTW Register (offset 0x3F80)

The *ITCBTW* register is used to read Trace Words from the on-chip trace memory. The TW read is the TW pointed to by the *ITCBRDP* register. A side effect of reading the *ITCBTW* register is that the *ITCBRDP* register increments to the next TW in the on-chip trace memory. If *ITCBRDP* is at the max size of the on-chip trace memory, the increment wraps back to address zero.

Note that this is a 64b register. On a 32b processor, software must read the upper word (offset 0x3F84) first as the address increment takes place on a read of the lower word (0x3F80).

The format of the *ITCBTW* register is shown below, and the field is described in [Table 11.38](#).

Figure 11.32 ITCBTW Register Format




Table 11.38 ITCBTW Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Names	Bits				
Data	63:0	Trace Word	R	Undefined	Required

11.8.6.3 ITCBRDP Register (Offset 0x3f88)

The *ITCBRDP* register is the address pointer to on-chip trace memory. It points to the TW read when reading the *ITCBTW* register. This value will be automatically incremented after a read of the *ITCBTW* register.

The format of the *ITCBRDP* register is shown below, and the field is described in [Table 11.39](#). 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 11.33 ITCBRDP Register Format




Table 11.39 ITCBRDP Register Field Descriptions

Fields		Description	Read/Write	Reset State	Compliance
Names	Bits				
Data	31:(n+1)	Reserved. Must be written zero, reads back zero.	0	0	Required
Address	n:0	Byte address of on-chip trace memory word.	R/W	Undefined	Required

11.8.6.4 ITCBWRP Register (Offset 0x3f90)

The *ITCBWRP* 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. The top bit in the register indicates whether the pointer has wrapped. If it has, then the write pointer will also point to the oldest trace word, and the read pointer can be set to that to read the entire array in order. If it is cleared, then the read pointer can be set to 0 to read up to the write pointer position.

The format of the *ITCBWRP* register is shown below, and the field is described in [Table 11.40](#). 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 11.34 ITCBWRP Register Format




Table 11.40 ITCBWRP Register Field Descriptions

Fields		Description	Read/Write	Reset State	Compliance
Names	Bits				
Wrap	31	Indicates that the entire array has been written at least once	R/W	Undefined	Required
0	30:(n+1)	Reserved. Must be written zero, reads back zero.	0	0	Required
Address	n:0	Byte address of the next on-chip trace memory word to be written	R/W	Undefined	Required

11.8.7 ITCB iFlowtrace Off-Chip Interface

The off-chip interface consists of a 4-bit data port (*TR_DATA*) and a trace clock (*TR_CLK*). *TR_CLK* can be a DDR clock; that is, both edges are significant. *TR_DATA* and *TR_CLK* follow the same timing and have the same output structure as the PDtrace TCB described in MIPS specifications. The trace clock is synchronous to the system clock but running at a divided frequency. The *OfClk* bit in the *Control/Status* register indicates the ratio between the trace clock and the core clock. The Trace clock is always 1/2 of the trace port data rate, hence the “full speed” ITCB outputs data at the CPU core clockrate but the trace clock is half that, hence the 1:2 *OfClk* value is the full speed, and the 1:4 *OfClk* ratio is half-speed.

When a 64-bit trace word is ready to transmit, the PIB reads it from the FIFO and begins sending it out on *TR_DATA*. It is sent in 4-bit increments starting at the LSBs. In a valid trace word, the 4 LSBs are never all zero, so a probe listening on the *TR_DATA* port can easily determine when the transmission begins and then count 15 additional cycles to collect the whole 64-bit word. Between valid transmissions, *TR_DATA* is held at zero and *TR_CLK* continues to run.

TR_CLK runs continuously whenever a probe is connected. An optional signal *TR_PROBE_N* may be pulled high when a probe is not connected and could be used to disable the off-chip trace port. If not present, this signal must be tied low at the Probe Interface Block (PIB) input.

The following encoding is used for the 6 tag bits to tell the PIB receiver that a valid transmission is starting:

```
// if (srcount == 0), EncodedSrCount = 111010 = 58
// else if (srcount == 16) EncodedSrCount = 111011 = 59
// else if (srcount == 32) EncodedSrCount = 111100 = 60
// else if (srcount == 48) EncodedSrCount = 111101 = 61
// else EncodedSrCount = srcount
```

11.8.8 Breakpoint-Based Enabling of Tracing

Each hardware breakpoint in the EJTAG block (see the MIPS EJTAG Specification, MD00047, revision 4.14) has a control bit associated with it that enables a trigger signal to be generated on a break match condition. In special trace mode, this trigger can be used to insert an event record into the trace stream. In normal trace mode, this trigger signal can be used to turn trace on or off, thus allowing a user to control the trace on/off functionality using breakpoints. Similar to the TraceIBPC and TraceDBPC registers in PDtrace, registers are defined to control the start and stop of iFlowtrace. The details on the actual register names and drseg addresses are shown in [Table 11.41](#).

Table 11.41 drseg Registers that Enable/Disable Trace from Breakpoint-Based Triggers

Register Name	drseg Address	Reset Value	Description
ITrigiFlowTrcEn	0x3FD0	0	Register that controls whether or not hardware instruction breakpoints can trigger iFlowtrace tracing functionality
DTrigiFlowTrcEn	0x3FD8	0	Register that controls whether or not hardware data and tuple breakpoints can trigger iFlowtrace tracing functionality

The bits in each register are defined as follows:

- Bit 28 (IE/DE): Used to specify whether the trigger signal from EJTAG simple or complex instruction (data or tuple) break should trigger iFlowtrace tracing functions or not. A value of 0 disables trigger signals from EJTAG instruction breaks, and 1 enables triggers for the same.

- Bits 14:0 (IBrk/DBrk): Used to explicitly specify which instruction (data or tuple) breaks enable or disable iFlowtrace. A value of 0 implies that trace is turned off (unconditional trace stop) and a value of 1 specifies that the trigger enables trace (unconditional trace start).

11.9 PC/Data Address Sampling

It is often useful for program profiling and analysis to periodically sample the value of the PC. This information can be used for statistical profiling akin to gprof, and is also very useful for detecting hot-spots in the code. In a multi-threaded environment, this information can be used to understand thread behavior, and to verify thread scheduling mechanisms in the absence of a full-fledged tracing facility like PDtrace.

The PC sampling feature is optional within EJTAG, but EJTAG and the TAP controller must be implemented if PC Sampling is required. When implemented, PC sampling can be turned on or off using an enable bit; when the feature is enabled, the PC value is continually sampled.


The presence or absence of the PC Sampling feature is indicated by the *PCS* (PC Sample) bit in the Debug Control Register. If PC sampling is implemented, and the *PCSe* (PC Sample Enable) bit in the Debug Control Register is also set to one, then the PC values are constantly sampled at the defined rate (*DCR_{PCR}*) and written to a TAP register. The old value in the TAP register is overwritten by the new value, even if this register has not been read out by the debug probe.

The presence or absence of Data Address Sampling is indicated by the *DAS* (Data Address Sample) bit in the Debug Control Register and enabled by the *DASe* (Data Address Sampling Enable) bit in the Debug Control Register.

The sample rate is specified by the 3-bit *PCR* (PC Sample Rate) field (bits 8:6) in the Debug Control Register (*DCR*). These three bits encode a value 2^5 to 2^{12} in a manner similar to the specification of SyncPeriod. When the implementation allows these bits to be written, the internal PC sample counter will be reset by each write, so that counting for the requested sample rate is immediately restarted.

The sample format includes a New data bit, the sampled value, the ASID of the sampled value (if not disabled by PCnoASID, bit 25 in *DCR*). Figure 11.35 shows the format of the sampled values in the PCSAMPLE TAP register for MIPS32. The New data bit is used by the probe to determine if the sampled data just read out is new or has already been read and must be discarded.

Figure 11.35 PCSAMPLE TAP Register Format (MIPS32)



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.

Note that some of the smaller sample periods can be shorter than the time needed to read out the sampled value. That is, it might take 41 (TCK) clock ticks to read a MIPS32 sample, while the smallest sample period is 32 (processor) clocks. While the sample is being read out, multiple samples may be taken and discarded, needlessly wasting power. To reduce unnecessary overhead, the TAP register includes only those fields that are enabled. If both PC Sampling and Data Sampling are enabled, then both samples are included in the PCSample scan register. PC Sample is in the least significant bits followed by a Data Address Sample. If either PC Sampling or Data Address Sampling is disabled, then the TAP register does not include that sample. The total scan length is $49 * 2 = 82$ bits if all fields are present and enabled.

11.9.1 PC Sampling in Wait State

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. 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 each 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.

11.9.2 Data Address Sampling

EJTAG revision 5.0 extends the PC sampling mechanism to allow sampling of data (load and store) addresses. This feature is enabled with DASE, bit 23 in the Debug Control Register. When enabled, the PCSAMPLE scan register includes a data address sample. All load and store addresses can be captured, or they can be qualified using a data breakpoint trigger. DASQ=1 configures data sampling to record a data address only when it triggers data breakpoint 0. To be used for Data Address Sampling qualification, data breakpoint 0 must be enabled using its TE (trigger enable) bit.

PCSR controls how often data addresses are sampled. When the PCSR counter triggers, the most recent load/store address generated is accepted and made available to shift out through PCSAMPLE.

11.10 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 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.

Refer to the EJTAG Specification [12] for the general details on FDC. The remainder of this section describes implementation specific behavior and register values.

The FDC memory mapped registers are located in the common device memory map (CDMM) region. FDC has a device ID of 0xFD.

11.10.1 Common Device Memory Map

Software on the core accesses FDC through memory-mapped registers, 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, as described in [8.2.35 “CDMMBase Register \(CP0 Register 15, Select 2\)” on page 207](#).

Refer to *MIPS® Architecture For Programmers Volume III* [9] for full details on the CDMM.

11.10.2 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 *CauseFDCI* bit indicates that the interrupt is pending. Traditionally, this interrupt is also sent to the core output *SI_FDCI* where it is combined with one of the *SI_Int* pins. However, this is no longer needed as the core will internally route the interrupt to the IP number set by the *IntCtl.IPFDCI* field. Note that this interrupt is a regular interrupt and not a debug interrupt.

The FDC Configuration Register (see [Section 11.10.6.2 “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 they are ORed together to form a single interrupt.

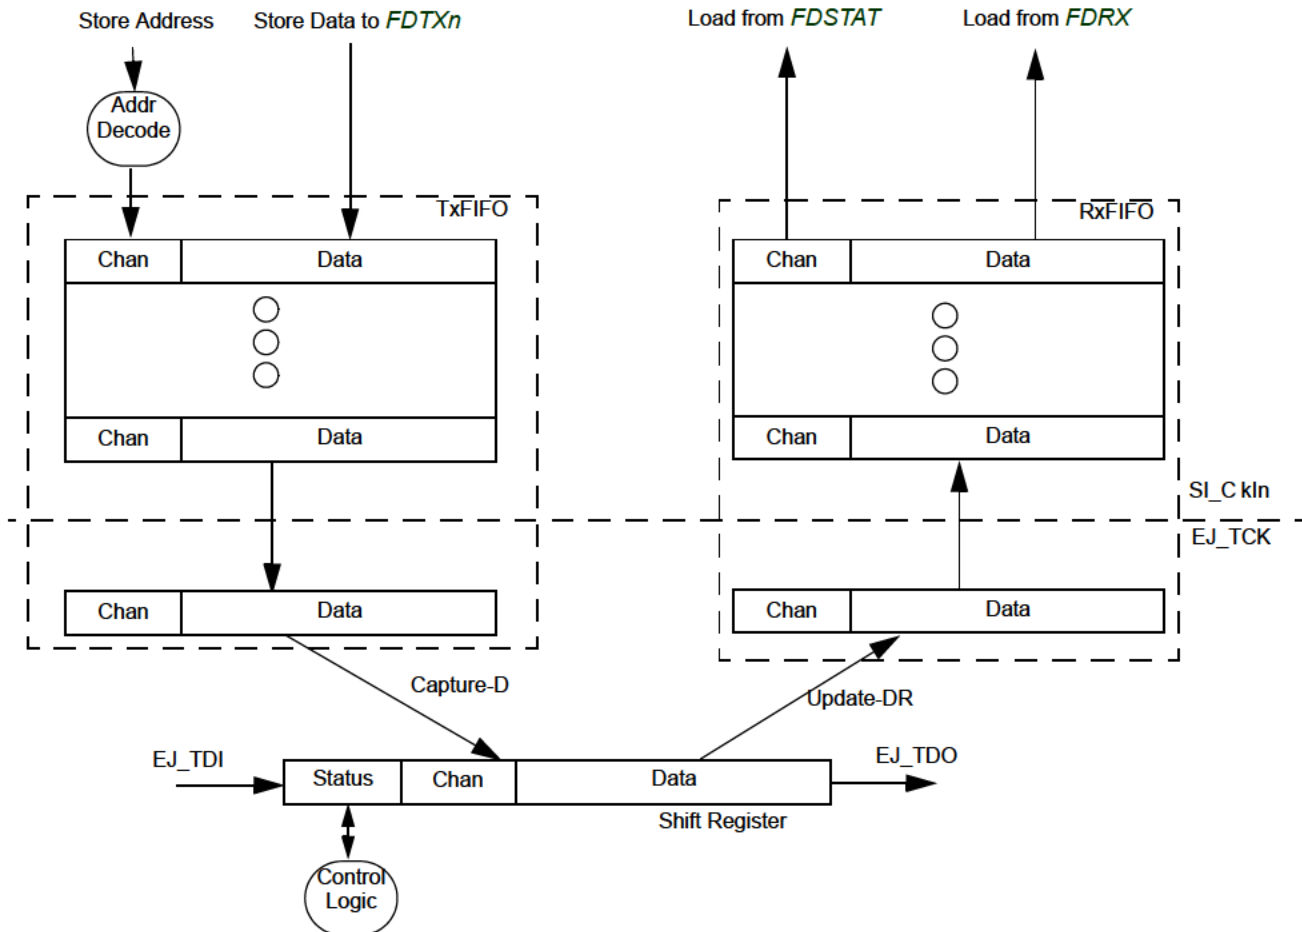
The following interrupt thresholds are supported:

- **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)

11.10.3 M5100™ Core FDC Buffers

[Figure 11.36](#) shows the general organization of the transmit and receive buffers on the M5100 core.

Figure 11.36 Fast Debug Channel Buffer Organization



One particular thing to note is the asynchronous crossings between the EJ_TCK and SI_CkIn clock domains. This crossing is handled with a handshake 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_CkIn FIFO entries, then into the EJ_TCK FIFO entry, requiring several EJ_TCK cycles to complete the handshake and move the data. EJ_TCK is generally much slower than SI_CkIn , and may even be stopped (although that would be uncommon when this feature is in use). This can result in not enough space for new data, even though there are only $N-1$ data values queued up. To prevent the loss of data, the TxF flag in $FDSTAT$ is set when all of the SI_CkIn FIFO entries are full. Software writes to the FIFO should always check the TxF bit before attempting the write and should not make any assumptions about being able to use all entries arbitrarily. i.e., software seeing the FxE bit set should not assume that it can write $TxCnt$ data words without checking for full.
- TxFIFO Almost Empty Interrupt:** As transmit data moves from SI_CkIn 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_CkIn 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 Rx FIFO has similar characteristics, but these are even less visible to software since *Sl_ClkIn* must be running to access the FDC registers.

11.10.4 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.

11.10.5 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 11.37](#) and the fields are described in [Figure 11.42](#)

Figure 11.37 FDC TAP Register Format




Table 11.42 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	0x0
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

Table 11.42 FDC TAP Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
Data	31:0	Data value being scanned in or out	R/W	Undefined

11.10.6 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 11.43](#)

Table 11.43 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)

11.10.6.1 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 11.38](#) has the format of an Access Control and Status register (shown as a 64-bit register), and [Table 11.44](#) describes the register fields.

Figure 11.38 FDC Access Control and Status Register




Table 11.44 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

Table 11.44 FDC Access Control and Status Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
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

11.10.6.2 FDC Configuration (FDCFG) Register (Offset 0x8)

The FDC configuration register holds information about the current configuration of the Fast Debug Channel mechanism. [Figure 11.39](#) has the format of the FDC Configuration register, and [Table 11.45](#) describes the register fields.

Figure 11.39 FDC Configuration Register




Table 11.45 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

Table 11.45 FDC Configuration Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State										
Name	Bits													
TxIntThresh	19:18	Controls whether transmit interrupts are enabled and the state of the TxFIFO needed to generate an interrupt. <table border="1"> <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 11.10.2 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 11.10.2 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 11.10.2 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"> <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										

11.10.6.3 FDC Status (FDSTAT) Register (Offset 0x10)

The FDC Status register holds up to date state information for the FDC mechanism. [Figure 11.40](#) shows the format of the FDC Status register, and [Table 11.46](#) describes the register fields.

Figure 11.40 FDC Status Register

31	24	23	16	15	8	7	4	3	2	1	0
Tx_Count		Rx_Count			0		RxChan	RxE	RxF	TxE	TxF

Table 11.46 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
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

Table 11.46 FDC Status Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
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

11.10.6.4 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 *RxE* flag in *FDSTAT* prior to reading. [Figure 11.41](#) shows the format of the *FDC Receive* register, and [Table 11.47](#) describes the register fields.

Figure 11.41 FDC Receive Register





Table 11.47 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

11.10.6.5 FDC Transmit n (FDTXn) Registers (Offset 0x20 + 0x8*n)

These sixteen registers 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 *TxF* flag in *FDSTAT* to ensure that there is space for the write. [Figure 11.42](#) shows the format of the FDC Transmit register, and [Table 11.48](#) describes the register fields.

Figure 11.42 FDC Transmit Register



Instruction Set Overview

This chapter provides a general overview on the three CPU instruction set formats of the MIPS architecture: Immediate, Jump, and Register. Refer to [Chapter 13, “M5100™ Processor Core Instructions”](#) on page 328 for a complete listing and description of instructions.

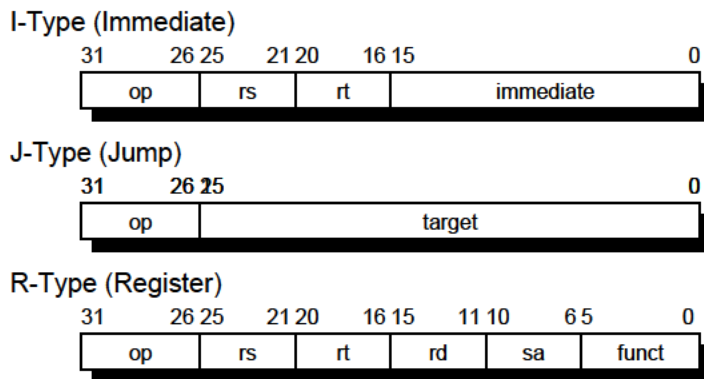
This chapter discusses the following topics

- [Section 12.1 “CPU Instruction Formats”](#)
- [Section 12.2 “Load and Store Instructions”](#)
- [Section 12.3 “Computational Instructions”](#)
- [Section 12.4 “Jump and Branch Instructions”](#)
- [Section 12.5 “Control Instructions”](#)
- [Section 12.6 “Coprocesor Instructions”](#)
- [Section 12.7 “Enhancements to the MIPS Architecture”](#)
- [Section 12.8 “MCU ASE Instructions”](#)

12.1 CPU Instruction Formats

Each 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) (shown in [Figure 12.1](#)). 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.

Figure 12.1 Instruction Formats



op	6-bit operation code
rs	5-bit source register specifier
rt	5-bit target (source/destination) register or branch condition
immediate	16-bit immediate value, branch displacement or address displacement
target	26-bit jump target address
rd	5-bit destination register specifier
sa	5-bit shift amount
funct	6-bit function field

12.2 Load and Store Instructions

12.2.1 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 the M5100 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.

12.2.2 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 the addressed field. 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.

Instruction Set Overview

The access type, together with the three low-order bits of the address, define the bytes accessed within the addressed word as shown in Table 12.1. Only the combinations shown in Table 12.1 are permissible; other combinations cause address error exceptions.

Table 12.1 Byte Access Within a Word

Access Type	Low Order Address Bits			Bytes Accessed							
	2	1	0	Big Endian (31-----0)				Little Endian (31-----0)			
				Byte				Byte			
Word	0	0	0	0	1	2	3	3	2	1	0
Triplebyte	0	0	0	0	1	2			2	1	0
	0	0	1		1	2	3	3	2	1	
Halfword	0	0	0	0	1					1	0
	0	1	0			2	3	3	2		
Byte	0	0	0	0							0
	0	0	1		1					1	
	0	1	0			2			2		
	0	1	1				3	3			

12.3 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
- Multiply
- Divide

These operations fit in the following four categories of computational instructions:

- ALU Immediate instructions
- Three-operand Register-type Instructions
- Shift Instructions
- Multiply And Divide Instructions

12.3.1 Cycle Timing for Multiply and Divide Instructions

Any multiply instruction in the integer pipeline is transferred to the multiplier as 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. Refer to [Chapter 2, “Pipeline of the M5100™ Core” on page 40](#) for more information on instruction latency and repeat rates.

12.4 Jump and Branch Instructions

Jump and branch instructions change the control flow of a program. All jump and branch instructions occur with a delay of one instruction: that is, the instruction immediately following the jump or branch (this is known as the instruction in the *delay slot*) always executes while the target instruction is being fetched from storage.

12.4.1 Overview of Jump Instructions

Subroutine calls in high-level languages are usually implemented with Jump or Jump and Link instructions, both of which are J-type instructions. 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, dispatches, and large cross-page jumps are usually implemented with the Jump Register or Jump and Link Register instructions. Both are R-type instructions that take 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 [Chapter 13, “M5100™ Processor Core Instructions” on page 328](#).

12.4.2 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 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.

12.5 Control Instructions

Control instructions allow the software to initiate traps; they are always R-type.

12.6 Coprocessor Instructions

CP0 instructions perform operations on the System Control Coprocessor registers to manipulate the memory management and exception handling facilities of the processor. Refer to [Chapter 13, “M5100™ Processor Core Instructions” on page 328](#) for a listing of CP0 instructions.

12.7 Enhancements to the MIPS Architecture

The core execution unit implements the MIPS32 architecture, which includes the following instructions.

Instruction Set Overview

- CLOCount Leading Ones
- CLZCount Leading Zeros
- MADDMultiply and Add Word
- MADDUMultiply and Add Unsigned Word
- MSUBMultiply and Subtract Word
- MSUBUMultiply and Subtract Unsigned Word
- MULMultiply Word to Register
- SSNOPSuperscalar Inhibit NOP

12.7.1 CLO - Count Leading Ones

The CLO instruction counts the number of leading ones in a word. The 32-bit word in the GPR *rs* is scanned from most-significant to least-significant bit. The number of leading ones is counted and the result is written to the GPR *rd*. If all 32 bits are set in the GPR *rs*, the result written to the GPR *rd* is 32.

12.7.2 CLZ - Count Leading Zeros

The CLZ instruction counts the number of leading zeros in a word. The 32-bit word in the GPR *rs* is scanned from most-significant to least-significant bit. The number of leading zeros is counted and the result is written to the GPR *rd*. If all 32 bits are cleared in the GPR *rs*, the result written to the GPR *rd* is 32.

12.7.3 MADD - Multiply and Add Word

The MADD instruction multiplies two words and adds the result to the HI/LO register pair. The 32-bit word value in the GPR *rs* is multiplied by the 32-bit value in the GPR *rt*, treating both operands as signed values, to produce a 64-bit result. The product is added to the 64-bit concatenated values in the HI and LO register pair. The resulting value is then written back to the HI and LO registers. No arithmetic exception occurs under any circumstances.

12.7.4 MADDU - Multiply and Add Unsigned Word

The MADDU instruction multiplies two unsigned words and adds the result to the HI/LO register pair. The 32-bit word value in the GPR *rs* is multiplied by the 32-bit value in the GPR *rt*, treating both operands as unsigned values, to produce a 64-bit result. The product is added to the 64-bit concatenated values in the HI and LO register pair. The resulting value is then written back to the HI and LO registers. No arithmetic exception occurs under any conditions.

12.7.5 MSUB - Multiply and Subtract Word

The MSUB instruction multiplies two words and subtracts the result from the HI/LO register pair. The 32-bit word value in the GPR *rs* is multiplied by the 32-bit value in the GPR *rt*, treating both operands as signed values, to produce a 64-bit result. The product is subtracted from the 64-bit concatenated values in the HI and LO register pair. The resulting value is then written back to the HI and LO registers. No arithmetic exception occurs under any circumstances.

12.7.6 MSUBU - Multiply and Subtract Unsigned Word

The MSUBU instruction multiplies two unsigned words and subtracts the result from the HI/LO register pair. The 32-bit word value in the GPR *rs* is multiplied by the 32-bit value in the GPR *rt*, treating both operands as unsigned values, to produce a 64-bit result. The product is subtracted from the 64-bit concatenated values in the HI and LO register pair. The resulting value is then written back to the HI and LO registers. No arithmetic exception occurs under any circumstances.

12.7.7 MUL - Multiply Word

The MUL instruction multiplies two words and writes the result to a GPR. The 32-bit word value in the GPR *rs* is multiplied by the 32-bit value in the GPR *rt*, treating both operands as signed values, to produce a 64-bit result. The least-significant 32-bits of the product are written to the GPR *rd*. The contents of the HI and LO register pair are not defined after the operation. No arithmetic exception occurs under any circumstances.

12.7.8 SSNOP- Superscalar Inhibit NOP

The MIPS32 M5100 processor cores treat this instruction as a regular NOP.

12.8 MCU ASE Instructions

The MCU ASE includes some new instructions which are particularly useful in microcontroller applications.

12.8.1 ACLR

This instruction allows a bit within an uncached I/O control register to be atomically cleared; that is, the read-modify byte write sequence performed by this instruction cannot be interrupted.

12.8.2 ASET

This instruction allows a bit within an uncached I/O control register to be atomically set; that is, the read-modify byte write sequence performed by this instruction cannot be interrupted.

12.8.3 IRET

This instruction can be used as a replacement for the ERET instruction when returning from an interrupt. This instruction implements the Automated Interrupt Epilogue feature, which automates restoring some of the COP0 registers from the stack and updating the CO_Status register in preparation for returning to non-exception mode. This instruction also implements the optional Interrupt Chaining feature, which allows a subsequent interrupt to be handled without returning to non-exception mode.

M5100™ Processor Core Instructions

This chapter supplements the *MIPS32® Architecture Reference Manual, Volume II* by describing instruction behavior that is specific to a MIPS32 M5100 processor core. For complete descriptions of all MIPS32 and microMIPS32 instructions, refer to *MIPS® Architecture For Programmers, Volume II: The MIPS32® Instruction Set* [7] and *MIPS® Architecture For Programmers, Volume II: The microMIPS32® Instruction Set* [8].

This chapter is divided into the following sections:

- [Section 13.1 “Understanding the Instruction Descriptions”](#)
- [Section 13.2 “M5100™ Core Opcode Map”](#)
- [Section 13.3 “MIPS32® Instruction Set for the M5100™ Core”](#)

The M5100 processor core also supports the MIPS32 microMIPS architecture. The microMIPS instruction set is described in [Chapter 14, “microMIPS™ Instruction Set Architecture”](#) on page 357.

The M5100 processor core also supports the instructions in the MIPS DSP Module Revision 2. The MIPS DSP Module Revision 2 instruction set is described in [Chapter 4, “The MIPS® DSP Module”](#) on page 101.

13.1 Understanding the Instruction Descriptions

Refer to *Volume II* of the *MIPS32 Architecture Reference Manual* for detailed information about the instruction descriptions, namely, the instruction fields, definition of terms, and functional notation. This section provides basic information

13.2 M5100™ Core Opcode Map

Key

- CAPITALIZED text indicates an opcode mnemonic
- *Italicized* text refers the reader to indicates to the specified opcode submap for further instruction bit decode.
- Entries containing the α symbol indicate that a reserved instruction fault occurs if the core executes this instruction.
- Entries containing the β symbol indicate that a coprocessor unusable exception occurs if the core executes this instruction

Table 13.1 Encoding of the Opcode Field

opcode		bits 28..26							
		0	1	2	3	4	5	6	7
bits 31..29		000	001	010	011	100	101	110	111
0	000	Special	RegImm	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	001	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	010	COP0	β	COP2	β	BEQL	BNEL	BLEZL	BGTZL
3	011	α	α	α	α	Special2	ΘΑΛΞ	α	Σπεχιαλ3
4	100	LB	LH	LWL	LW	LBU	LHU	LWR	α
5	101	SB	SH	SWL	SW	α	α	SWR	CACHE
6	110	LL	β	LWC2	PREF	α	β	α	α
7	111	SC	β	SWC2	α	α	β	α	α

Table 13.2 Special Opcode Encoding of Function Field

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	SLL	β	SRL/ ROTR	SRA	SLLV	α	SRLV/ ROTRV	SRAV
1	001	JR	JALR	MOVZ	MOVN	SYSCALL	BREAK	α	SYNC
2	010	MFHI	MTHI	MFLO	MTLO	α	α	α	α
3	011	MULT	MULTU	DIV	DIVU	α	α	α	α
4	100	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	101	α	α	SLT	SLTU	α	α	α	α
6	110	TGE	TGEU	TLT	TLTU	TEQ	α	TNE	α
7	111	α	α	α	α	α	α	α	α

Table 13.3 Special2 Opcode Encoding of Function Field

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	MADD	MADDU	MUL	α	MSUB	MSUBU	α	α
1	001		α	α	α	α	α	α	α
2	010	<i>UDI</i> ¹ or α							
3	011								
4	100	CLZ	CLO	α	α	α	α	α	α
5	101	α	α	α	α	α	α	α	α
6	110	α	α	α	α	α	α	α	α
7	111	α	α	α	α	α	α	α	SDBBP

1. CorExtend instructions are a build-time option of the M5100 Pro core, if not implemented this instructions space will cause a reserved instruction exception. If assembler support exists, the mnemonics for CorExtend instructions are most likely UDI0, UDI1, ..., UDI15.

Table 13.4 Special3 Opcode Encoding of Function Field

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	EXT	α	α	α	INS	α	α	α
1	001	α	α	α	α	α	α	α	α
2	010	α	α	α	α	α	α	α	α
3	011	α	α	α	α	α	α	α	α
4	100	BSHFL	α	α	α	α	α	α	α
5	101	α	α	α	α	α	α	α	α
6	110	α	α	α	α	α	α	α	α
7	111	α	α	α	ΡΔΗΩΡ	α	α	α	α

Table 13.5 RegImm Encoding of rt Field

rt		bits 18..16							
		0	1	2	3	4	5	6	7
bits 20..19		000	001	010	011	100	101	110	111
0	00	BLTZ	BGEZ	BLTZL	BGEZL	α	α	α	α
1	01	TGEI	TGEIU	TLTI	TLTIU	TEQI	α	TNEI	α
2	10	BLTZAL	BGEZAL	BLTZALL	BGEZALL	α	α	α	α
3	11	α	α	α	α	α	α	α	ΣΨNXI

Table 13.6 COP2 Encoding of rs Field

rs		bits 23..21							
		0	1	2	3	4	5	6	7
bits 25..24		000	001	010	011	100	101	110	111
0	00	MFC2	α	CFC2	MΦHX2	MTC2	α	CTC2	MTHX2
1	01	BC2	BC2 ¹						
2	10	CO							
3	11								

1. The core will treat the entire row as a *BC2* instruction. However compiler and assembler support only exists for the first one. Some compiler and assembler products may allow the user to add new instructions.

Table 13.7 COP2 Encoding of rt Field When rs=BC2

rt		bits 16	
bits 17		0	1
0		BC2F	BC2T
1		BC2FL	BC2TL

Table 13.8 COP0 Encoding of rs Field

rs		bits 23..21							
		0	1	2	3	4	5	6	7
bits 25..24		000	001	010	011	100	101	110	111
0	00	MFC0	α	α	α	MTC0	α	α	α
1	01	α	α	ΡΔΠΓΠΡ	ΜΦΜΧ0	α	α	ΩΡΠΓΠΡ	α
2	10	CO							
3	11								

Table 13.9 COP0 Encoding of Function Field When rs=CO

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	α	α	α	α	α	α	α	α
1	001	α	α	α	α	α	α	α	α
2	010	α	α	α	α	α	α	α	α
3	011	ERET	IAXK	α	α	α	α	α	DERET
4	100	WAIT	α	α	α	α	α	α	α
5	101	α	α	α	α	α	α	α	α
6	110	α	α	α	α	α	α	α	α
7	111	α	α	α	α	α	α	α	α

Table 13.10 MIPS32 COP1 Encoding of rs Field

rs		bits 23..21							
		0	1	2	3	4	5	6	7
bits 25..24		000	001	010	011	100	101	110	111
0	00	MFC1	*	CFC1	MFHC1	MTC1	*	CTC1	MTHC1
1	01	$BC1\delta$	*	*	*	*	*	*	*
2	10	$S\delta$	$D\delta$	*	*	$W\delta$	$L\delta$	*	*
3	11	*	*	*	*	*	*	*	*

Table 13.11 MIPS32 COP1 Encoding of Function Field When rs=S

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	ROUND.L ▽	TRUNC.L ▽	CEIL.L ▽	FLOOR.L ▽	ROUND.W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	MOVCF δ	MOVZ	MOVN	*	RECIP ▽	RSQRT ▽	*
3	011	*	*	*	*	*	*	*	*
4	100	*	CVT.D	*	*	CVT.W	CVT.L ▽	*	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

Table 13.12 MIPS32 COP1 Encoding of Function Field When rs=D

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	ROUND.L ▽	TRUNC.L ▽	CEIL.L ▽	FLOOR.L ▽	ROUND.W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	MOVCF δ	MOVZ	MOVN	*	RECIP ▽	RSQRT ▽	*
3	011	*	*	*	*	*	*	*	*
4	100	CVT.S	*	*	*	CVT.W	CVT.L ▽	*	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

Table 13.13 MIPS32 COP1 Encoding of Function Field When rs=W or L¹

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	*	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

1. Format type L is legal only if 64-bit floating point operations are enabled.

Table 13.14 MIPS32 COP1 Encoding of tf Bit When rs=S or D, Function=MOVCF

tf	bit 16	
	0	1
	MOVf fmt	MOVt fmt

13.3 MIPS32® Instruction Set for the M5100™ Core

This section provides a summary of the MIPS32 instructions for the M5100 cores (microMIPS32 instructions are described in [Chapter 14, “microMIPS™ Instruction Set Architecture”](#) on page 357).

[Table 13.15](#) lists the instructions in alphabetical order. Instructions that have implementation-dependent behavior are described in subsequent sections; all other MIPS32 instructions are described in detail in the *MIPS® Architecture For Programmers, Volume II: The MIPS32® Instruction Set* [7] and are not duplicated here.

Table 13.15 Instruction Set

Instruction	Description	Function
ADD	Integer Add	$Rd = Rs + Rt$
ADDI	Integer Add Immediate	$Rt = Rs + \text{Immed}$
ADDIU	Unsigned Integer Add Immediate	$Rt = Rs +_{\text{U}} \text{Immed}$
ADDU	Unsigned Integer Add	$Rd = Rs +_{\text{U}} Rt$
AND	Logical AND	$Rd = Rs \& Rt$
ANDI	Logical AND Immediate	$Rt = Rs \& (0_{16} \parallel \text{Immed})$
ACLR	Atomic Bit Clear	See MCU ASE Instructions
ASET	Atomic Bit Set	See MCU ASE Instructions
B	Unconditional Branch (Assembler idiom for: BEQ r0, r0, offset)	$PC += (\text{int})\text{offset}$
BAL	Branch and Link (Assembler idiom for: BGEZAL r0, offset)	$GPR[31] = PC + 8$ $PC += (\text{int})\text{offset}$
BC2F	Branch On COP2 Condition False	if $COP2Condition(cc) == 0$ $PC += (\text{int})\text{offset}$
BC2FL	Branch On COP2 Condition False Likely	if $COP2Condition(cc) == 0$ $PC += (\text{int})\text{offset}$ else Ignore Next Instruction
BC2T	Branch On COP2 Condition True	if $COP2Condition(cc) == 1$ $PC += (\text{int})\text{offset}$
BC2TL	Branch On COP2 Condition True Likely	if $COP2Condition(cc) == 1$ $PC += (\text{int})\text{offset}$ else Ignore Next Instruction
BEQ	Branch On Equal	if $Rs == Rt$ $PC += (\text{int})\text{offset}$

Table 13.15 Instruction Set (Continued)

Instruction	Description	Function
BEQL	Branch On Equal Likely	if $R_s == R_t$ PC += (int)offset else Ignore Next Instruction
BGEZ	Branch on Greater Than or Equal To Zero	if $!R_s[31]$ PC += (int)offset
BGEZAL	Branch on Greater Than or Equal To Zero And Link	GPR[31] = PC + 8 if $!R_s[31]$ PC += (int)offset
BGEZALL	Branch on Greater Than or Equal To Zero And Link Likely	GPR[31] = PC + 8 if $!R_s[31]$ PC += (int)offset else Ignore Next Instruction
BGEZL	Branch on Greater Than or Equal To Zero Likely	if $!R_s[31]$ PC += (int)offset else Ignore Next Instruction
BGTZ	Branch on Greater Than Zero	if $!R_s[31] \ \&\& \ R_s != 0$ PC += (int)offset
BGTZL	Branch on Greater Than Zero Likely	if $!R_s[31] \ \&\& \ R_s != 0$ PC += (int)offset else Ignore Next Instruction
BLEZ	Branch on Less Than or Equal to Zero	if $R_s[31] \ \ R_s == 0$ PC += (int)offset
BLEZL	Branch on Less Than or Equal to Zero Likely	if $R_s[31] \ \ R_s == 0$ PC += (int)offset else Ignore Next Instruction
BLTZ	Branch on Less Than Zero	if $R_s[31]$ PC += (int)offset
BLTZAL	Branch on Less Than Zero And Link	GPR[31] = PC + 8 if $R_s[31]$ PC += (int)offset
BLTZALL	Branch on Less Than Zero And Link Likely	GPR[31] = PC + 8 if $R_s[31]$ PC += (int)offset else Ignore Next Instruction
BLTZL	Branch on Less Than Zero Likely	if $R_s[31]$ PC += (int)offset else Ignore Next Instruction
BNE	Branch on Not Equal	if $R_s != R_t$ PC += (int)offset

Table 13.15 Instruction Set (Continued)

Instruction	Description	Function
BNEL	Branch on Not Equal Likely	if Rs != Rt PC += (int)offset else Ignore Next Instruction
BREAK	Breakpoint	Break Exception
CACHE	Cache Operation	NOP
CFC2	Move Control Word From Coprocessor 2	Rt = CCR[2, n]
CLO	Count Leading Ones	Rd = NumLeadingOnes(Rs)
CLZ	Count Leading Zeroes	Rd = NumLeadingZeroes(Rs)
COP0	Coprocessor 0 Operation	See Coprocessor Description
COP2	Coprocessor 2 Operation	See Coprocessor 2 Description
CTC2	Move Control Word To Coprocessor 2	CCR[2, n] = Rt
DERET	Return from Debug Exception	PC = DEPC Exit Debug Mode
DI	Disable Interrupts	Rt=Status Status _{IE} =0
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	Stall until execution hazards are cleared
EI	Enable Interrupts	Rt=Status Status _{IE} =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,msbd,lsb)
HYPCALL	Trigger Hypercall exception	See Instruction Description Chapter of [14] or [15].
INS	Insert Bit Field	Rt=InsertField(Rt,Rs,msb,lsb)
IRET	Return from Exception	See MCU ASE Instructions
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

Table 13.15 Instruction Set (Continued)

Instruction	Description	Function
JALR.HB	Jump and Link Register with Hazard Barrier	Rd = PC + 8 PC = Rs Stall until all execution and instruction hazards are cleared
JR	Jump Register	PC = Rs
JR.HB	Jump Register with Hazard Barrier	PC = Rs Stall until all execution and instruction hazards are cleared
LB	Load Byte	Rt = (byte)Mem[Rs+offset]
LBU	Unsigned Load Byte	Rt = (ubyte)Mem[Rs+offset]
LH	Load Halfword	Rt = (half)Mem[Rs+offset]
LHU	Unsigned Load Halfword	Rt = (uhalf)Mem[Rs+offset]
LL	Load Linked Word	Rt = Mem[Rs+offset] LL = 1 LLAdr = Rs + offset
LUI	Load Upper Immediate	Rt = immediate << 16
LW	Load Word	Rt = Mem[Rs+offset]
LWC2	Load Word To Coprocessor 2	CPR[2, n, 0] = Mem[Rs+offset]
LWL	Load Word Left	See LWL instruction.
LWR	Load Word Right	See LWR instruction.
MADD	Multiply-Add	HI, LO += (int)Rs * (int)Rt
MFC0	Move From Coprocessor 0	Rt = CPR[0, n, sel]
MFC2	Move From Coprocessor 2	Rt = CPR[2, n, sel] _{31 0}
MFGC0	Move from Guest Coprocessor 0	See Instruction Description Chapter of [14] or [15].
MFHC2	Move From High Word Coprocessor2	Rt= CPR[2,n,sel] _{63 32}
MFHI	Move From HI	Rd = HI
MFLO	Move From LO	Rd = LO
MOVN	Move Conditional on Not Zero	if GPR[rt] ≠ 0 then GPR[rd] = GPR[rs]
MOVZ	Move Conditional on Zero	if GPR[rt] = 0 then GPR[rd] = GPR[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
MTC2	Move To Coprocessor 2	CPR[2, n, sel] _{31 0} = Rt
MTGC0	Move To Guest Coprocessor 0	See Instruction Description Chapter of [14] or [15].
MTHC2	Move To High Word Coprocessor 2	CPR[2, n, sel] _{63 32} = Rt
MTHI	Move To HI	HI = Rs

Table 13.15 Instruction Set (Continued)


Instruction	Description	Function
MTLO	Move To LO	LO = Rs
MUL	Multiply with register write	HI LO = Unpredictable Rd = LO
MULT	Integer Multiply	HI LO = (int)Rs * (int)Rd
NOP	No Operation (Assembler idiom for: SLL r0, r0, r0)	
NOR	Logical NOR	Rd = ~(Rs Rt)
OR	Logical OR	Rd = Rs Rt
ORI	Logical OR Immediate	Rt = Rs Immed
PREF	Prefetch	NOP
RDHWR	Read HardWare Register	Rt = HWR[Rd]
RDPGPR	Read GPR from Previous Shadow Set	Rd = SGPR[SRSCtl _{pSS} , Rt]
ROTR	Rotate Word Right	Rd = Rt _{sa-1:0} Rt _{31:sa}
ROTRV	Rotate Word Right Variable	Rd = Rt _{Rs-1:0} Rt _{31:Rs}
SB	Store Byte	(byte)Mem[Rs+offset] = Rt
SC	Store Conditional Word	if LL = 1 mem[Rxoffs] = Rt Rt = LL
SDBBP	Software Debug Breakpoint	Trap to SW Debug Handler
SEB	Sign Extend Byte	Rd = SignExtend(Rt _{7:0})
SEH	Sign Extend Half	Rd = SignExtend(Rt _{15:0})
SH	Store Halfword	(half)Mem[Rs+offset] = Rt
SLL	Shift Left Logical	Rd = Rt << sa
SLLV	Shift Left Logical Variable	Rd = Rt << Rs[4:0]
SLT	Set on Less Than	if (int)Rs < (int)Rt Rd = 1 else Rd = 0
SLTI	Set on Less Than Immediate	if (int)Rs < (int)Immed Rt = 1 else Rt = 0
SLTIU	Set on Less Than Immediate Unsigned	if (uns)Rs < (uns)Immed Rt = 1 else Rt = 0
SLTU	Set on Less Than Unsigned	if (uns)Rs < (uns)Immed Rd = 1 else Rd = 0
SRA	Shift Right Arithmetic	Rd = (int)Rt >> sa
SRAV	Shift Right Arithmetic Variable	Rd = (int)Rt >> Rs[4:0]

Table 13.15 Instruction Set (Continued)

Instruction	Description	Function
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	$Mem[Rs+offset] = Rt$
SWC2	Store Word From Coprocessor 2	$Mem[Rs+offset] = CPR[2, n, 0]$
SWL	Store Word Left	See SWL instruction description.
SWR	Store Word Right	See SWR instruction description.
SYNC	Synchronize	See SYNC instruction below.
SYNCI	Synchronize Caches to Make Instruction Writes Effective	NOP
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 $(int)Rs \geq (int)Rt$ TrapException
TGEI	Trap if Greater Than or Equal Immediate	if $(int)Rs \geq (int)Immed$ TrapException
TGEIU	Trap if Greater Than or Equal Immediate Unsigned	if $(uns)Rs \geq (uns)Immed$ TrapException
TGEU	Trap if Greater Than or Equal Unsigned	if $(uns)Rs \geq (uns)Rt$ TrapException
TLBGP	Probe Guest TLB	See Instruction Description Chapter of [14] or [15].
TLBGR	Read Guest TLB	See Instruction Description Chapter of [14] or [15].
TLBGWI	Write Guest TLB	See Instruction Description Chapter of [14] or [15].
TLBGWR	Write Guest TLB	See Instruction Description Chapter of [14] or [15].
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

Table 13.15 Instruction Set (Continued)

Instruction	Description	Function
TNE	Trap if Not Equal	if Rs != Rt TrapException
TNEI	Trap if Not Equal Immediate	if Rs != (int)Immed TrapException
WAIT	Wait for Interrupts	Stall until interrupt occurs
WRPGPR	Write to GPR in Previous Shadow Set	SGPR[SRSCtl _{PSS} ,Rd]=Rt
WSBH	Word Swap Bytes within Halfwords	Rd=SwapBytesWithinHalfs(Rt)
XOR	Exclusive OR	Rd = Rs ^ Rt
XORI	Exclusive OR Immediate	Rt = Rs ^ (uns)Immed



Format: ACLR bit, offset(base)

microMIPS and MCU ASE

Purpose: Atomically Clear Bit within Byte

Description: Disable interrupts; $temp \leftarrow memory[GPR[base] + offset]$; $temp \leftarrow (temp \text{ and } \sim(1 \ll bit))$; $memory[GPR[base] + offset] \leftarrow temp$; Enable Interrupts

The contents of the byte at the memory location specified by the effective address are fetched. The specified bit within the byte is cleared to zero. The modified byte is stored in memory at the location specified by the effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address. The read-modify-write sequence cannot be interrupted.

Transactions with locking semantics occur in some memory interconnects/busses. It is implementation-specific whether this instruction uses such locking transactions.

Restrictions:

The operation of the processor is **UNDEFINED** if an ACLR instruction is executed in the delay slot of a branch or jump instruction.

Operation:

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..2 || (pAddr_1..0 xor ReverseEndian2)
TempIE ← Status_IE
Status_IE ← 0
memword ← LoadMemory(CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr_1..0 xor BigEndianCPU2
temp ← memword7+8*byte..8*byte
temp ← temp and ((1 || 0bit) xor 0xFF)
dataword ← temp || 08*byte
StoreMemory(CCA, BYTE, dataword, pAddr, vAddr, DATA)
Status_IE ← TempIE
    
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Programming Notes:

Upon a TLB miss, a TLBS exception is signalled in the ExcCode field of the *Cause* register. For address error, a ADES exception is signalled in the ExcCode field of the *Cause* register. For other data-stream related exceptions such as Debug Data Break exceptions and Watch exceptions, it is implementation-specific whether this instruction is treated as a load or as a store.



Format: ASET bit, offset(base)

MIPS and MCU ASE

Purpose: Atomically Set Bit within Byte

Description: Disable interrupts; temp ← memory[GPR[base] + offset]; temp ← (temp or (1 << bit)) ; memory[GPR[base] + offset] ← temp; Enable Interrupts

The contents of the 8-bit byte at the memory location specified by the effective address are fetched. The specified bit within the byte is set to one. The modified byte is stored in memory at the location specified by the effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address. The read-modify-write sequence cannot be interrupted.

Transactions with locking semantics occur in some memory interconnects/busses. It is implementation-specific whether this instruction uses such locking transactions.

Restrictions:

The operation of the processor is **UNDEFINED** if an ASET instruction is executed in the delay slot of a branch or jump instruction.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1.. || (pAddr..0 xor ReverseEndian)
TempIE ← Status_IE
Status_IE ← 0
memword ← LoadMemory (CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr..0 xor BigEndianCPU
temp ← memword7+8*byte..8*byte
temp ← temp or ( 1 || 0bit )
dataword ← temp || 08*byte
StoreMemory (CCA, BYTE, dataword, pAddr, vAddr, DATA)
Status_IE ← TempIE
    
```


Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Programming Notes:

Upon a TLB miss, a TLBS exception is signalled in the ExcCode field of the *Cause* register. For address error, a ADES exception is signalled in the ExcCode field of the *Cause* register. For other data-stream related exceptions such as Debug Data Break exceptions and Watch exceptions, it is implementation-specific whether this instruction is treated as a load or as a store.

Atomically Set Bit within Byte



Format: ASET bit, offset (base)

microMIPS AND MCU ASE

Purpose: Atomically Set Bit within Byte

Description: Disable interrupts; $temp \leftarrow memory[GPR[base] + offset]$; $temp \leftarrow (temp \text{ or } (1 \ll bit))$; $memory[GPR[base] + offset] \leftarrow temp$; Enable Interrupts

The contents of the byte at the memory location specified by the effective address are fetched. The specified bit within the byte is set to one. The modified byte is stored in memory at the location specified by the effective address. The 12-bit signed *offset* is added to the contents of GPR *base* to form the effective address. The read-modify-write sequence cannot be interrupted.

Transactions with locking semantics occur in some memory interconnects/busses. It is implementation-specific whether this instruction uses such locking transactions.

Restrictions:

The operation of the processor is **UNDEFINED** if an ASET instruction is executed in the delay slot of a branch or jump instruction.

Operation:

```
vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1.. || (pAddr..0 xor ReverseEndian)
TempIE ← Status_IE
Status_IE ← 0
memword ← LoadMemory (CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr..0 xor BigEndianCPU
temp ← memword7+8*byte..8*byte
temp ← temp or (1 || 0bit)
dataword ← temp || 08*byte
StoreMemory (CCA, BYTE, dataword, pAddr, vAddr, DATA)
Status_IE ← TempIE
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Programming Notes:

Upon a TLB miss, a TLBS exception is signalled in the ExcCode field of the *Cause* register. For address error, a ADES exception is signalled in the ExcCode field of the *Cause* register. For other data-stream related exceptions such as Debug Data Break exceptions and Watch exceptions, it is implementation-specific whether this instruction is treated as a load or as a store.

31	26	25	6	5	0
COP0 010000	C0 1	0 00 0000 0000 0000 0000			IRET 111000
6	1	20			6

Format: IRET

MIPS and MCU ASE

Purpose: Interrupt Return with Automated Interrupt Epilogue

Optionally jump directly to another interrupt vector without returning to original return address.

Description:

IRET is used to automate some of the operations that are required when returning from an interrupt handler. It can be used in place of the ERET instruction at the end of interrupt handlers. The IRET instruction is only appropriate when using Shadow Register Sets and EIC Interrupt mode. The automated operations of this instruction can be used to reverse the effects of the automated operations of the Auto-Prologue feature.

If the EIC mode of interrupts and the Interrupt Chaining feature are used, the IRET instruction can be used to shorten the time between returning from the current interrupt handler and handling the next requested interrupt.

If Automated Prologue feature is disabled, then IRET behaves exactly as ERET.

If either *Status_{ERL}* or *Status_{BEV}* bits are set, then IRET behaves exactly as ERET.

If Interrupt Chaining is disabled:

- Interrupts are disabled. COP0 *Status*, *SRSCtl*, and *EPC* registers are restored from the stack. GPR 29 is incremented for the stack frame size. IRET then clears execution and instruction hazards, conditionally restores *SRSCtl_{CSS}* from *SRSCtl_{PSS}*, and returns to the interrupted instruction pointed by the *EPC* register at the completion of interrupt processing.

If Interrupt Chaining is enabled:

- Interrupts are disabled. COP0 *Status* register is restored from the stack. The priority output of the External Interrupt Controller is compared with the *IPL* field of the *Status* register.
- If *Status_{IPL}* has a higher priority than that of the External Interrupt Controller value:

COP0 *SRSCtl* and *EPC* registers are restored from the stack. GPR 29 is incremented for the stack frame size. IRET then clears execution and instruction hazards, conditionally restores *SRSCtl_{CSS}* from *SRSCtl_{PSS}*, and returns to the interrupted instruction pointed by the *EPC* register at the completion of interrupt processing.

- If *Status_{IPL}* field has a lower priority than that of the External Interrupt Controller value:

The value of GPR 29 is first saved to a temporary register then GPR 29 is incremented for the stack frame size. The EIC is signalled that the next pending interrupt has been accepted. This signalling will update the *Cause_{R IPL}* and *SRSCtl_{EICSS}* fields from the EIC output values. The *SRSCtl_{EICSS}* field is copied to the *SRSCtl_{CSS}* field while the *Cause_{R IPL}* field is copied to the *Status_{IPL}* field. The saved temporary register is copied to the GPR 29 of the current SRS. The *KSU*, *ERL* and *EXL* fields of the *Status* register are optionally set to zero. No barrier for execution hazards nor instruction hazards is created. IRET finishes by jumping to the interrupt vector driven by the EIC.

IRET does not execute the next instruction (i.e., it has no delay slot).

Restrictions:

The operation of the processor is **UNDEFINED** if an IRET is executed in the delay slot of a branch or jump instruction.

The operation of the processor is **UNDEFINED** if an IRET is executed when either Shadow Register Sets are not enabled or when EIC interrupt mode is not enabled.

An IRET placed between an LL and SC instruction will always cause the SC to fail.

The effective addresses used for the stack memory transactions must be naturally-aligned. If either of the two least-significant bits of the address is non-zero, an Address Error exception occurs.

IRET implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the IRET returns.

In a Release 2 implementation, IRET does not restore $SRSCtl_{CSS}$ from $SRSCtl_{PSS}$ if $Status_{BEV} = 1$, or if $Status_{ERL} = 1$ because any exception that sets $Status_{ERL}$ to 1 (Reset, Soft Reset, NMI, or cache error) does not save $SRSCtl_{CSS}$ in $SRSCtl_{PSS}$. If software sets $Status_{ERL}$ to 1, it must be aware of the operation of an IRET that may be subsequently executed.

The stack memory transactions behave as individual LW operations with respect to exception reporting. BadVAddr would report the faulting address for unaligned access and the faulting word address for unprivileged access, TLB Refill and TLB Invalid exceptions. For TLB exceptions, the faulting word address would be reflected in the *Context*, and *EntryHi* registers. The *CacheError* register would reflect the faulting word address for Cache Errors.

Operation:

```

if (( IntCtlAPE == 0 ) | ( StatusERL == 1 ) | ( StatusBEV == 1 ))
  Act as ERET // read Operation section of ERET description
else
  if (ISAMode)
    EPC ← PC..1 || 1 // in case of memory exception
  else
    EPC ← PC // in case of memory exception
  endif
  temp ← 0x4 + GPR[29]
  tempStatus ← LoadStackWord(temp)
  ClearHazards()
  if ( (IntCtlICE == 0) | ((IntCtlICE == 1) &
    (tempStatusIPL > EICRIPL)) )
    temp ← 0x8 + GPR[29]
    tempSRCtl ← LoadStackWord(temp)
    temp ← 0x0 + GPR[29]
    tempEPC ← LoadStackWord(temp)
  endif
  Status ← tempStatus
  if ( (IntCtlICE == 0) | ((IntCtlICE == 1) &
    (tempStatusIPL > EICRIPL)) )
    GPR[29] ← GPR[29] + DecodedValue(IntCtlstkDec)
    SRCtl ← tempSRCtl
    EPC ← tempEPC
    temp ← EPC
    StatusEXL ← 0
    if (ArchitectureRevision ≥ 2) and (SRCtlHSS > 0)
      and (StatusBEV = 0) then
        SRCtlCSS ← SRCtlPSS
    endif
  endif

```

```

endif
if IsMicroMIPSImplemented() then
    PC ← temp31..1 || 0
    ISAMode ← temp0
else
    PC ← temp
endif
LLbit ← 0
CauseIC ← 0
ClearHazards()
else
    Signal_EIC_for_Next_Interrupt()
    (wait for EIC outputs to update)
    CauseR IPL ← EICR IPL
    SRSCtlEICSS ← EICSS
    temp29 ← GPR[29]
    GPR[29] ← GPR[29] + DecodedValue(IntCtlstkDec)
    StatusIPL ← CauseR IPL
    SRSCtlCSS ← SRSCtlEICSS
    NewShadowSet ← SRSCtlEICSS
    GPR[29] ← temp29
    if (IntCtlClrEXL == 1)
        StatusEXL ← 0
        StatusKSU ← 0
    endif
    CauseIC ← 1
    ClearHazards()
    PC ← CalcIntrptAddress()
endif
endif

function 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
endfunction LoadStackWord

function CalcIntrptAddress()
    if StatusBEV = 1
        vectorBase ← 0xBFC0.0200
    else
        if ( ArchitectureRevision ≥ 2)
            vectorBase ← EBase31..12 || 011)
        else
            vectorBase ← 0x8000.0000
        endif
    endif
    if (CauseIV = 0)
        vectorOffset = 0x180
    else
        if (StatusBEV = 1) or (IntCtlvS = 0)
            vectorOffset = 0x200
        endif
    endif
endfunction CalcIntrptAddress()

```

```
else
  if ( Config3VEIC = 1 and EIC_Option=1)
    VectorNum = CauseRIPL
  elseif (Config3VEIC = 1 and EIC_Option=2)
    VectorNum = EIC_VectorNum
  elseif (Config3VEIC = 0 )
    VectorNum = VIntPriorityEncoder()
  endif
  if (Config3VEIC = 1 and EIC_Option=3)
    vectorOffset = EIC_VectorOffset
  else
    vectorOffset = 0x200 + (VectorNum x (IntCtlVS || 05))
  endif
endif

endif
CalcIntrptAddress = vectorBase | vectorOffset
endfunction CalcIntrptAddress
```

Exceptions:

Coprocessor Unusable Exception, TLB Refill, TLB Invalid, Address Error, Watch, Cache Error, Bus Error
Exceptions

31	26	25	6	5	0
POOL32A 000000	000 0000 0011 0100 1101			POOL32AXf 111100	
6	20			6	

Format: IRET

microMIPS and MCU ASE

Purpose: Interrupt Return with Automated Interrupt Epilogue

Optionally jump directly to another interrupt vector without returning to original return address.

Description:

IRET automates some of the operations that are required when returning from an interrupt handler and can be used in place of the ERET instruction at the end of interrupt handlers. IRET is only appropriate when using Shadow Register Sets and the EIC Interrupt mode. The automated operations of this instruction can be used to reverse the effects of the automated operations of the Auto-Prologue feature.

If the EIC interrupt mode and the Interrupt Chaining feature are used, the IRET instruction can be used to shorten the time between returning from the current interrupt handler and handling the next requested interrupt.

If the Automated Prologue feature is disabled, then IRET behaves exactly like ERET.

If either the *Status_{ERL}* or *Status_{BEV}* bits are set, then IRET behaves exactly like ERET.

If Interrupt Chaining is disabled:

Interrupts are disabled. COP0 *Status*, *SRSCtl*, and *EPC* registers are restored from the stack. GPR 29 is incremented for the stack frame size. IRET then clears execution and instruction hazards, conditionally restores *SRSCtl_{CSS}* from *SRSCtl_{PSS}*, and returns at the completion of interrupt processing to the interrupted instruction pointed to by the *EPC* register. If Interrupt Chaining is enabled:

Interrupts are disabled. COP0 *Status* register is restored from the stack. The priority output of the External Interrupt Controller is compared with the IPL field of the *Status* register.

If *Status_{IPL}* has a higher priority than the External Interrupt Controller value:

COP0 *SRSCtl* and *EPC* registers are restored from the stack. GPR 29 is incremented for the stack frame size. IRET then clears execution and instruction hazards, conditionally restores *SRSCtl_{CSS}* from *SRSCtl_{PSS}*, and returns to the interrupted instruction pointed to by the *EPC* register at the completion of interrupt processing.

If *Status_{IPL}* has a lower priority than the External Interrupt Controller value:

The value of GPR 29 is first saved to a temporary register and then GPR 29 is incremented for the stack frame size. The EIC is signalled that the next pending interrupt has been accepted. This signalling will update the *Cause_{R IPL}* and *SRSCtl_{EICSS}* fields from the EIC output values. The *SRSCtl_{EICSS}* field is copied to the *SRSCtl_{CSS}* field, while the *Cause_{R IPL}* field is copied to the *Status_{IPL}* field. The saved temporary register is copied to the GPR 29 of the current SRS. The KSU and EXL fields of the *Status* register are optionally set to zero. No barrier for execution hazards or instruction hazards is created. IRET finishes by jumping to the interrupt vector driven by the EIC.

IRET does not execute the next instruction (i.e., it has no delay slot).

Restrictions:

The operation of the processor is **UNDEFINED** if IRET is executed in the delay slot of a branch or jump instruction.

The operation of the processor is **UNDEFINED** if IRET is executed when either Shadow Register Sets are not enabled, or the EIC interrupt mode is not enabled.

An IRET placed between an LL and SC instruction will always cause the SC to fail.

The effective addresses used for stack transactions must be naturally-aligned. If either of the two least-significant bits of the address is non-zero, an Address Error exception occurs.

IRET implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SY NCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier begin with the instruction fetch and decode of the instruction at the PC to which the IRET returns.

In a Release 2 implementation, IRET does not restore $SRSCtl_{CSS}$ from $SRSCtl_{PSS}$ if $Status_{BEV} = 1$ or $Status_{ERL} = 1$, because any exception that sets $Status_{ERL}$ to 1 (Reset, Soft Reset, NMI, or cache error) does not save $SRSCtl_{CSS}$ in $SRSCtl_{PSS}$. If software sets $Status_{ERL}$ to 1, it must be aware of the operation of an IRET that may be subsequently executed.

The stack transactions behave as individual L W operations with respect to exception reporting. BadVAddr would report the faulting address for an unaligned access, and the faulting word address for unprivileged access, TLB Refill, and TLB Invalid exceptions. For TLB exceptions, the faulting word address would be reflected in the *Context* and *EntryHi* registers. The *CacheError* register would reflect the faulting word address for Cache Errors.

Operation:

```

if (( IntCtlAPE == 0 ) | ( StatusERL == 1 ) | ( StatusBEV == 1 ))
    Act as ERET // read Operation section of ERET description
else
    if (ISAMode)
        EPC ← PC31..1 || 1 // in case of memory exception
    else
        EPC ← PC // in case of memory exception
    endif
    temp ← 0x4 + GPR[29]
    tempStatus ← LoadStackWord(temp)
    ClearHazards()
    if ( (IntCtlICE == 0) | ((IntCtlICE == 1) &
        (tempStatusIPL > EICRIPL)) )
        temp ← 0x8 + GPR[29]
        tempSRSCtl ← LoadStackWord(temp)
        temp ← 0x0 + GPR[29]
        tempEPC ← LoadStackWord(temp)
    endif
    Status ← tempStatus
    if ( (IntCtlICE == 0) | ((IntCtlICE == 1) &
        (tempStatusIPL > EICRIPL)) )
        GPR[29] ← GPR[29] + DecodedValue(IntCtlstkDec)
        SRSCtl ← tempSRSCtl
        EPC ← tempEPC
        temp ← EPC
        StatusEXL ← 0
        if (ArchitectureRevision ≥ 2) and (SRSCtlHSS > 0) and (StatusBEV = 0) then
            SRSCtlCSS ← SRSCtlPSS
        endif
        if IsMicroMIPSImplemented() then

```

```

        PC ← temp31..1 || 0
        ISAMode ← temp0
    else
        PC ← temp
    endif
    LLbit ← 0
    CauseIC ← 0
    ClearHazards()
else
    Signal_EIC_for_Next_Interrupt()
    (wait for EIC outputs to update)
    CauseR IPL ← EICR IPL
    SRSCtlEICSS ← EICSS
    temp29 ← GPR[29]
    GPR[29] ← GPR[29] + DecodedValue(IntCtlStkDec)
    StatusIPL ← CauseR IPL
    SRSCtlCSS ← SRSCtlEICSS
    NewShadowSet ← SRSCtlEICSS
    GPR[29] ← temp29
    if (IntCtlClrEXL == 1)
        StatusEXL ← 0
        StatusKSU ← 0
    endif
    CauseIC ← 1
    ClearHazards()
    PC ← CalcIntrptAddress()
endif
endif

function 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
endfunction LoadStackWord


function CalcIntrptAddress()
    if StatusBEV = 1
        vectorBase ← 0xBFC0.0200
    else
        if (ArchitectureRevision □ 2)
            vectorBase ← EBase31..12 || 011
        else
            vectorBase ← 0x8000.0000
        endif
    endif
    if (CauseIV = 0)
        vectorOffset = 0x180
    else
        if (StatusBEV = 1) or (IntCtlVS = 0)
            vectorOffset = 0x200
        else
            if (Config3VEIC = 1 and EIC_Option=1)

```

```
        VectorNum = CauseRIPL
    elseif (Config3VEIC = 1 and EIC_Option=2)
        VectorNum = EIC_VectorNum
    elseif (Config3VEIC = 0 )
        VectorNum = VIntPriorityEncoder()
    endif
    if (Config3VEIC = 1 and EIC_Option=3)
        vectorOffset = EIC_VectorOffset
    else
        vectorOffset = 0x200 + (VectorNum x (IntCtlVS || 05))
    endif
    endif
    CalcIntrptAddress = vectorBase | vectorOffset
endfunction CalcIntrptAddress
```

Exceptions:

Coprocessor Unusable Exception, TLB Refill, TLB Invalid, Address Error, Watch, Cache Error, Bus Error
Exceptions



Format: LL *rt*, *offset* (*base*)

MIPS32

Purpose: Load Linked Word

To load a word from memory for an atomic read-modify-write

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 processor. There can be only one active RMW sequence per processor. When an LL is executed it starts an active RMW sequence replacing any other sequence that was active. The 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 processor does not cause an action that, by itself, causes an SC for the same block to fail on another processor.

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 processors 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 addressed location may be uncached for the *M5100* core.

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: SC *rt*, *offset* (*base*)

MIPS32

Purpose: Store Conditional Word

To store a word to memory to complete an atomic read-modify-write

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*. On the M5100 core, the SRAM interface supports a lock protocol and the success or failure can be indicated by external hardware.

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] ← 031 || 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.

31	26	25	21	20	16	15	11	10	6	5	0
SPECIAL 000000			0 00 0000 0000 0000 0				stype			SYNC 001111	
6			15				5			6	

Format: SYNC (stype = 0 implied)

MIPS32

Purpose: Synchronize Shared Memory

To order loads and stores.

Description:

Simple Description:

- SYNC affects only *uncached* and *cached coherent* loads and stores. The loads and stores that occur before the SYNC must be completed before the loads and stores after the SYNC 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.
- SYNC 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 SYNC is required on entry to and exit from Debug Mode to guarantee that memory affects are handled correctly.

Detailed Description:

- SYNC does not guarantee the order in which instruction fetches are performed. The *stype* values 1-31 are reserved for future extensions to the architecture. A value of zero will always be defined such that it performs all defined synchronization operations. Non-zero values may be defined to remove some synchronization operations. As such, software should never use a non-zero value of the *stype* field, as this may inadvertently cause future failures if non-zero values remove synchronization operations.
- The SYNC instruction is externalized on the SRAM interface of the *M5100* core. External logic can use this information in a system-dependent manner to enforce memory ordering between various memory elements in the system.

Restrictions:

The effect of SYNC on the global order of loads and stores for memory access types other than *uncached* and *cached coherent* is **UNPREDICTABLE**.

Operation:

SyncOperation(stype)

Exceptions:

None

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:

The WAIT instruction forces the core into low power mode. The pipeline is stalled and when all external requests are completed, the processor's main clock is stopped. The processor will restart when reset (*SL_Reset* or *SL_ColdReset*) is signaled, or a non-masked interrupt is taken (*SI_NMI*, *SI_Int*, or *EJ_DINT*). Note that the *M5100* core does not use the code field in this instruction.

If the pipeline restarts as the result of an enabled interrupt, that interrupt is taken between the WAIT instruction and the following instruction (EPC for the interrupt points at 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

microMIPS™ Instruction Set Architecture

The microMIPS™ architecture minimizes the code footprint of applications, thus reducing the cost of memory, which is particularly high for embedded memory. At the same time, the high performance of MIPS cores is maintained. Using this technology, the customer can generate best results without spending time to profile its application. The smaller code footprint typically leads to reduced power consumption per executed task because of the smaller number of memory accesses.

microMIPS is a replacement for the existing MIPS16e ASE. It is also an alternative to the MIPS32 instruction encoding and can be implemented in parallel or stand-alone.

Overview of changes from the existing MIPS 32ISA:

- 16-bit and 32-bit opcodes; for MIPS64, also includes 48-bit opcodes
- Optimized opcode/operand field definitions based on statistics
- Branch and jump delay slots are retained for maximum compatibility and lowest risk
- Removal of branch likely instructions, emulation by assembler
- Fine-tuned register allocation algorithm in compilers for smallest code size

14.1 Overview

14.1.1 MIPSr3™ Architecture

MIPSr3 is a family of architectures which includes Release 3.0 of the MIPS32 Architecture and the first release of the microMIPS32 architecture. Enhancements included in the MIPSr3 Architecture are:

- MIPS32 Release 3 ISA and microMIPS ISA.
- The MIPS16e ASE is phased out and is replaced by microMIPS. Therefore these two ASEs never co-exist within the same processor core.
- Branch likely instructions are phased out in microMIPS and are emulated by the assembler. They remain available in the MIPS32 encoding.

Unless otherwise described in this document, all other aspects of the MIPS32r3 architecture are identical to MIPS32 Release 2.

14.1.2 Default ISA Mode

The instruction sets available in an implementation are reported in the *Config3_{ISA}* register field (bits 15:14). *Config1_{CA}* (bit 2) is not used for microMIPS.

For implementations that support both microMIPS and MIPS32, the selected ISA mode following reset is determined by the setting of the *Config3_{ISA}* register field, which is a read-only field set by a hardware signal external to the processor core.

For implementations that support both microMIPS and MIPS32, the selected ISA mode of an exception handler is determined by the setting of the *Config3_{ISAOnExc}* register field (bit 16). The *Config3_{ISAOnExc}* register field is writeable by software and has a reset value that is set by a hardware signal external to the processor core. This register field allows privileged software to change the ISA mode to be used for subsequent exceptions. All exception types whose vectors are offsets of the *EBase* register have this capability.

For implementations that support both microMIPS and MIPS32, the selected ISA mode of a debug exception is determined by the setting of the *ISAonDebug* register field in the *EJTAG TAP Control* register. This register field is writeable by EJTAG probe software and has a reset value that is set by a hardware signal external to the processor core.

14.1.3 Software Detection

Software can determine if microMIPS is implemented by checking the state of the *ISA* (Instruction Set Architecture) field in the *Config3* CP0 register. *Config1_{CA}* (bit 2) is not used for microMIPS.

Software can determine if the MIPS32 ISA is implemented by checking the state of the *ISA* (Instruction Set Architecture) register field in the *Config3* CP0 register.

Software can determine which ISA is used when handling an exception by checking the state of the *ISAOnExc* (ISA on Exception) field in the *Config3* CP0 register.

Debug Probe Software can determine which ISA is used when handling a debug exception by checking the state of the *ISAonDebug* field in the *EJTAG TAP Control* register.

14.1.4 Compliance and Subsetting

This document does not change the instruction subsets as defined by the other MIPS32 architecture reference manuals, including the subsets defined by the various ASEs.

14.1.5 Mode Switch

The MIPS32 architecture defines an ISA mode for each processor. An ISA mode value of 0 indicates MIPS32 instruction decoding. In processors implementing microMIPS, an ISA mode value of 1 selects microMIPS instruction decoding.

In microMIPS implementations, the ISA mode is not directly visible to normal software. When EJTAG is implemented, the ISA mode is reflected in the *EJTAG TAP Control* register.

Mode switching between MIPS32 and microMIPS uses the same mechanism used by MIPS16e, namely, the JALX, JR, JR.HB, JALR, and JALR.HB instructions, as described below.

- The JALX instruction executes a JAL and switches to the other mode.

- The JR and JALR instructions interpret bit 0 of the source registers as the target ISA mode (0=MIPS32, 1=microMIPS) and therefore set the ISA Mode bit according to the contents of bit 0 of the source register. For the actual jump operation, the PC is loaded with the value of the source register with bit 0 set to 0. The same applies to JR.HB and JALR.HB. The instructions JALR and JALR.HB save the ISA mode into bit 0 of the destination register.
- When exceptions or interrupts occur and the processor writes to *EPC*, *DEPC*, or *ErrorEPC*, the ISA Mode bit is saved into bit 0 of these registers. Then the ISA Mode bit is set according to the *Config3*_{ISA} register field. On return from an exception, the processor loads the ISA Mode bit based on the value from either *EPC*, *DEPC*, or *ErrorEPC*.

If only one ISA mode exists (either MIPS32 or microMIPS), then this mode switch mechanism does not exist, and the ISA mode has a fixed value (0=MIPS32, 1=microMIPS). Executing the JALX instruction will cause a Reserved Instruction exception. JR and JALR instructions cause an Address exception on the target instruction fetch when bit 0 of the source register is different from the ISA mode. The same applies to JR.HB and JALR.HB. Exception handlers must be encoded in the instruction format supported by the processor.

14.1.6 Branch and Jump Offsets

In the MIPS32 architecture, because instructions are always 32 bits in size, the jump and branch target addresses are word (32-bit) aligned. Jump/branch offset fields are shifted left by two bits to create a word-aligned effective address.

In the microMIPS architecture, because instructions can be either 16 or 32 bits in size, the jump and branch target addresses are halfword (16-bit) aligned. Branch/jump offset fields are shifted left by only one bit to create halfword-aligned effective addresses.

To maintain the existing MIPS32 ABIs, link unit/object file entry points are restricted to 32-bit word alignments. In the future, a microMIPS-only ABI can be created to remove this restriction.

14.1.7 Coprocessor Unusable Behavior

If an instruction associated with a non-implemented coprocessor is executed, it is implementation-specific whether a processor executing in microMIPS mode raises an RI exception or a coprocessor unusable exception. While in microMIPS mode, the M5100 has the same behavior as in MIPS32 mode; coprocessor unusable exceptions will be raised.

14.2 Instruction Formats

This section defines the formats of microMIPS instructions. The microMIPS variable-length encoding comprises 16-bit and 32-bit wide instructions.

The 6-bit major opcode is left-aligned within the instruction encoding. Instructions can have 0 to 4 register fields. For 32-bit instructions, the register field width is 5 bits, while for most 16-bit instructions, the register field width is 3 bits, utilizing instruction-specific register encoding. All 5-bit register fields are located at a constant position within the instruction encoding.

The immediate field is right-aligned in the following instructions:

- some 16-bit instructions with 3-bit register fields
- 32-bit instructions with 16-bit or 26-bit immediate field

The name ‘immediate field’ as used here includes the address offset field for branches and load/store instructions as well as the jump target field.

Other instruction-specific fields are typically located between the immediate and minor opcode fields. Instructions that have multiple “other” fields are listed in alphabetical order according to the name of the field, with the first name of the order located at the lower bit position. An empty bit field that is not explicitly shown in the instruction format is located next to the minor opcode field.

Figure 14.1 and Figure 14.2 show the 16-bit and 32-bit instruction formats.

Figure 14.1 16-Bit Instruction Formats

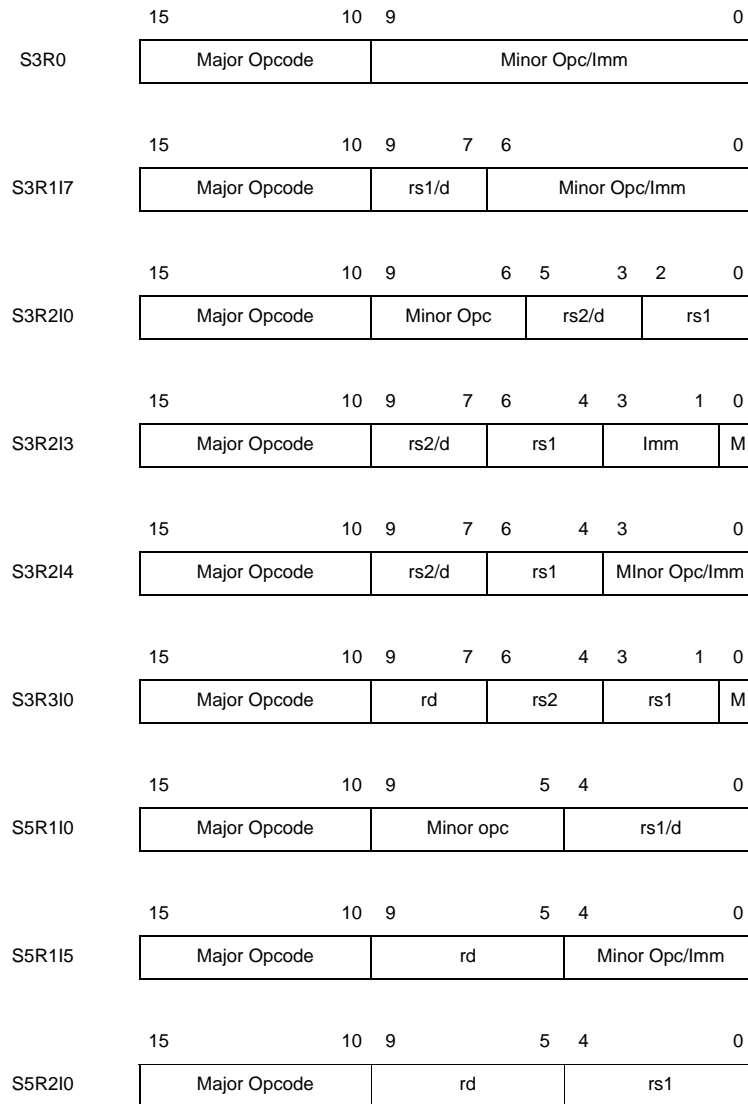


Figure 14.2 32-Bit Instruction Formats






Figure 14.3 Immediate Fields within 32-Bit Instructions


32-bit instruction formats with 26-bit immediate fields:



32-bit instruction formats with 16-bit immediate fields:



32-bit instruction formats with 12-bit immediate fields:



The instruction size can be completely derived from the major opcode. For 32-bit instructions, the major opcode also defines the position of the minor opcode field and whether or not the immediate field is right-aligned.

Instructions formats are named according to the number of the register fields and the size of the immediate field. The names have the structure R<x>I<y>. For example, an instruction based on the format R2I16 has 2 register fields and a 16-bit immediate field.

14.2.1 Instruction Stream Organization and Endianness

16-bit instructions are placed within the 32-bit (or 64-bit) memory element according to system endianness.

- On a 32-bit processor in big-endian mode, the first instruction is read from bits 31..16, and the second instruction is read from bits 15..0.
- On a 32-bit processor in little-endian mode, the first instruction is read from bits 15..0, and the second instruction is read from bits 31..16.

The above rule also applies to the halfwords of 32-bit instructions. This means that a 32-bit instruction is not treated as a word data type; instead, the halfwords are treated in the same way as individual 16-bit instructions. The halfword containing the major opcode is always the first in the sequence.

Example:

```

SRL r1, r1, 7    binary opcode fields: 000000 00001 00001 00111 00001 000000
                  hex representation:   0021 3840

                  Address:   3  2  1  0
Little Endian:    Data:     38 40 00 21

                  Address:   0  1  2  3
Big Endian:      Data:     00 21 38 40
    
```

Instructions are placed in memory such that they are in-order with respect to the address.

14.3 microMIPS Re-encoded Instructions

This section lists all microMIPS re-encoded instructions, sorted into 16-bit and 32-bit categories.

In the 16-bit category:

- Frequent MIPS32 instructions and macros, re-encoded as 16-bit. Register and immediate fields are reduced in size by using encodings of frequently occurring values.

In the 32-bit category:

- All MIPS32 instructions, including all application-specific extensions except MIPS16e, re-encoded: MIPS32, MIPS-3D ASE, MIPS DSP Module, MIPS MT ASE, and SmartMIPS ASE.
- Opcode space for user-defined instructions (UDIs).
- New instructions designed primarily to reduce code size.

To differentiate between 16-bit and 32-bit encoded instructions, the instruction mnemonic can be optionally extended with the suffix “16” or “32” respectively. This suffix is placed at the end of the instruction before the first ‘.’ if there is one. For example:

```
ADD16, ADD32, ADD32.PS
```

If these suffixes are omitted, the assembler automatically chooses the smallest instruction size.

For each instruction, the tables in this chapter provide all necessary information about the bit fields. The formats of the instructions are defined in [Section 12.1 “CPU Instruction Formats”](#). Together with the major and minor opcode encodings, which can be derived from the tables in [Section 13.2 “M5100™ Core Opcode Map”](#), the complete instruction encoding is provided.

Most register fields have a width of 5 bits. 5-bit register fields use linear encoding (r0=’00000’, r1=’00001’, etc.). For 16-bit instructions, whose register field size is variable, the register field width is explicitly stated in the instruction table ([Table 14.1](#) and [Table 14.2](#)), and the individual register and immediate encodings are shown in [Table 14.3](#). The ‘other fields’ are defined by the respective column, with the order of these fields in the instruction encoding defined by the order in the tables.

14.3.1 16-Bit Category

14.3.1.1 Frequent MIPS32 Instructions

These are frequent MIPS32 instructions with reduced register and immediate fields containing frequently used registers and immediate values.

MOVE is a very frequent instruction. It therefore supports full 5-bit unrestricted register fields for maximum efficiency. In fact, MOVE used to be a simplified macro of an existing MIPS32 instruction.

There are 2 variants of the LW and SW instructions. One variant implicitly uses the SP register to allow for a larger offset field. The value in the offset field is shifted left by 2 before it is added to the base address.

There are four variants of the ADDIU instruction:

1. A variant with one 5-bit register specifier that allows any GPR to be the source and destination register
2. A variant that uses the stack pointer as the implicit source and destination register
3. A variant that has separate 3-bit source and destination register specifiers
4. A variant that has the stack pointer as the implicit source register and one 3-bit destination register specifier

A 16-bit NOP instruction is needed because of the new 16-bit instruction alignment and the need in specific cases to align instructions on a 32-bit boundary. It can save code size as well. NOP is not shown in the table because it is realized as a macro (as is NEGU).

```
NOP16 = MOVE16 r0, r0
```

```
NEGU16 rt, rs = SUBU16 rt, r0, rs
```

Because microMIPS instructions are 16-bit aligned, the 16-bit branch instructions support 16-bit aligned branch target addresses. The offset field is left shifted by 1 before it is added to the PC.

The compact instruction JRC is to be used instead of JR, when the jump delay slot after JR cannot be filled. This saves code size. Because JRC may execute as fast as JR with a NOP in the delay slot, JR is preferred if the delay slot can be filled.

The breakpoint instructions, BREAK and SDBBP, include a 16-bit variant that allows a breakpoint to be inserted at any instruction address without overwriting more than a single instruction.

Table 14.1 16-Bit Re-encoding of Frequent MIPS32 Instructions

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
ADDIU5	POOL16D	5bit:1	4	5		0	1	Add Immediate Unsigned Word Same Register
ADDIUSP	POOL16D	0	9	0		0	1	Add Immediate Unsigned Word to Stack Pointer
ADDIUR2	POOL16E	2	3	3		0	1	Add Immediate Unsigned Word Two Registers
ADDIUR1SP	POOL16E	1	6	3		0	1	Add Immediate Unsigned Word One Registers and Stack Pointer
ADDU16	POOL16A	3	0	3		0	1	Add Unsigned Word
AND16	POOL16C	2	0	3		0	4	AND
ANDI16	ANDI16	2	4	3		0	0	AND Immediate
B16	B16	0	10			0	0	Branch
BREAK16	POOL16C	0	0		4	0	6	Cause Breakpoint Exception
JALR16	POOL16C	1	0	5		0	5	Jump and Link Register, 32-bit delay-slot
JALRS16	POOL16C	1	0	5		0	5	Jump and Link Register, 16-bit delay-slot
JR16	POOL16C	1	0	5		0	5	Jump Register
LBU16	LBU16	2	4	3		0	0	Load Byte Unsigned
LHU16	LHU16	2	4	3		0	0	Load Halfword
LI16	LI16	1	7	3		0	0	Load Immediate
LW16	LW16	2	4	3		0	0	Load Word
LWGP	LWGP16	1	7	3		0	0	Load Word GP
LWSP	LWSP16	5bit:1	5	5		0	0	Load Word SP

Table 14.1 16-Bit Re-encoding of Frequent MIPS32 Instructions (Continued)

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
MFHI16	POOL16C	1	0	5		0	5	Move from HI Register
MFLO16	POOL16C	1	0	5		0	5	Move from LO Register
MOVE16	MOVE16	2	0	5		0	0	Move
NOT16	POOL16C	2	0	3		0	4	NOT
OR16	POOL16C	2	0	3		0	4	OR
SB16	SB16	2	4	3		0	0	Store Byte
SDBBP16	POOL16C	0	0		4	0	6	Cause Debug Breakpoint Exception
SH16	SH16	2	4	3		0	0	Store Halfword
SLL16	POOL16B	2	3	3		0	1	Shift Word Left Logical
SRL16	POOL16B	2	3	3		0	1	Shift Word Right Logical
SUBU16	POOL16A	3	0	3		0	1	Sub Unsigned
SW16	SW16	2	4	3		0	0	Store Word
SWSP	SWSP16	5bit:1	5	5		0	0	Store Word SP
XOR16	POOL16C	2	0	3		0	4	XOR

14.3.1.2 Frequent MIPS32 Instruction Sequences

These 16-bit instructions are equivalent to frequently-used short sequences of MIPS32 instructions. The instruction-specific register and immediate value selection are shown in [Table 14.3](#).

Table 14.2 16-Bit Re-encoding of Frequent MIPS32 Instruction Sequences

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
BEQZ16	BEQZ16	1	7	3		0	0	Branch on Equal Zero
BNEZ16	BNEZ16	1	7	3		0	0	Branch on Not Equal Zero
JRADDIUSP	POOL16C	0	5				5	Jump Register; ADDIU SP
JRC	POOL16C	1	0	5		0	5	Jump Register Compact

Table 14.2 16-Bit Re-encoding of Frequent MIPS32 Instruction Sequences (Continued)

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
LWM16	POOL16C	0	4		2	0	4	Load Word Multiple
SWM16	POOL16C	0	4		2	0	4	Store Word Multiple

14.3.1.3 Instruction-Specific Register Specifiers and Immediate Field Encodings

Table 14.3 Instruction-Specific Register Specifiers and Immediate Field Values

Instruction	Number of Register Fields	Immediate Field Size (bit)	Register 1 Decoded Value	Register 2 Decoded Value	Register 3 Decoded Value	Immediate Field Decoded Value
ADDIU5	5bit:1	4	rd: 5-bit field			-8..0..7
ADDIU16	0	9				$(-258..-3, 2..257) \ll 2$
ADDIUR2	2	3	rs1:2-7,16, 17	rd:2-7,16, 17		-1, 1, 4, 8, 12, 16, 20, 24
ADDIUR16	1	6	rd:2-7,16, 17			$(0..63) \ll 2$
ADDU16	3	0	rs1:2-7,16, 17	rs2:2-7,16, 17	rd:2-7,16, 17	
AND16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		
ANDI16	2	4	rs1:2-7,16, 17	rd:2-7,16, 17		1, 2, 3, 4, 7, 8, 15, 16, 31, 32, 63, 64, 128, 255, 32768, 65535
B16	0	10				$(-512..511) \ll 1$
BEQZ16	1	7	rs1:2-7,16, 17			$(-64..63) \ll 1$
BNEZ16	1	7	rs1:2-7,16, 17			$(-64..63) \ll 1$
BREAK16	0	4				0..15
JALR16	5bit:1	0	rs1:5-bit field			
JALRS16	5bit:1	0	rs1:5-bit field			
JRADDIU16	0	5				$(0..31) \ll 2$
JR16	5bit:1	0	rs1:5 bit field			
JRC	5bit:1	0	rs1:5 bit field			
LBU16	2	4	rb:2-7,16,17	rd:2-7,16, 17		-1,0..14
LHU16	2	4	rb:2-7,16,17	rd:2-7,16, 17		$(0..15) \ll 1$
LI16	1	7	rd:2-7,16, 17			-1,0..126
LW16	2	4	rb:2-7,16,17	rd:2-7,16, 17		$(0..15) \ll 2$
LWM16	2bit list:1	4				$(0..15) \ll 2$
LWGP	1	7	rd:2-7,16,17			$(-64..63) \ll 2$
LWSP	5bit:1	5	rd:5-bit field			$(0..31) \ll 2$
MFHI16	5bit:1	0	rd:5-bit field			
MFLO16	5bit:1	0	rd:5-bit field			
MOVE16	5bit:2	0	rd:5-bit field	rs1:5-bit field		
NOT16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		
OR16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		
SB16	2	4	rb:2-7,16,17	rs1:0, 2-7, 17		0..15
SDBBP16	0	0				0..15
SH16	2	4	rb:2-7,16,17	rs1:0, 2-7, 17		$(0..15) \ll 1$
SLL16	2	3	rs1:2-7,16, 17	rd:2-7,16, 17		1..8 (see encoding tables)

Table 14.3 Instruction-Specific Register Specifiers and Immediate Field Values (Continued)

Instruction	Number of Register Fields	Immediate Field Size (bit)	Register 1 Decoded Value	Register 2 Decoded Value	Register 3 Decoded Value	Immediate Field Decoded Value
SRL16	2	3	rs1:2-7,16, 17	rd:2-7,16, 17		1..8 (see encoding tables)
SUBU16	3	0	rs1:2-7,16, 17	rs2:2-7,16, 17	rd:2-7,16, 17	
SW16	2	4	rb:2-7,16,17	rs1:0, 2-7, 17		(0..15) << 2
SWSP	5bit:1	5	rs1: 5 bit field			(0..31) << 2
SWM16	2-bit list:1	4				(0..15)<<2
XOR16	2	0	rs1:2-7,16, 17	rd:2-7,16, 17		

14.3.2 16-bit Instruction Register Set

Many of the 16-bit instructions use 3-bit register specifiers in their binary encodings. The register set used for most of these 3-bit register specifiers is listed in Table 14.4. The register set used for SB16, SH16, SW16 source register is listed in Table 14.5. These register sets are a true subset of the register set available in 32-bit mode; the 3-bit register specifiers can directly access 8 of the 32 registers available in 32-bit mode (which uses 5-bit register specifiers).

In addition, specific instructions in the 16-bit instruction set implicitly reference the stack pointer register (*sp*), global pointer register (*gp*), the return address register (*ra*), the integer multiplier/divider output registers (*HI/LO*) and the program counter (*PC*). Of these, Table 14.6 lists *sp*, *gp* and *ra*. Table 14.7 lists the microMIPS special-purpose registers, including *PC*, *HI* and *LO*.

The microMIPS also contains some 16-bit instructions that use 5-bit register specifiers. Such 16-bit instructions provide access to all 32 general-purpose registers.

Table 14.4 16-Bit Instruction General-Purpose Registers - \$2-\$7, \$16, \$17

16-Bit Register Encoding ¹	32-Bit MIPS Register Encoding ²	Symbolic Name (From <i>ArchDefs.h</i>)	Description
0	16	s0	General-purpose register
1	17	s1	General-purpose register
2	2	v0	General-purpose register
3	3	v1	General-purpose register
4	4	a0	General-purpose register
5	5	a1	General-purpose register
6	6	a2	General-purpose register
7	7	a3	General-purpose register

1. “0-7” correspond to the register’s 16-bit binary encoding and show how that encoding relates to the MIPS registers. “0-7” never refer to the registers, except within the binary microMIPS instructions. From the assembler, only the MIPS names (\$16, \$17, \$2, etc.) or the symbolic names (s0, s1, v0, etc.) refer to the registers. For example, to access register number 17 in the register file, the programmer references \$17 or s1, even though the microMIPS binary encoding for this register is 001.
2. General registers not shown in the above table are not accessible through the 16-bit instruction using 3-bit register specifiers. The Move instruction can access all 32 general-purpose registers.

Table 14.5 SB16, SH16, SW16 Source Registers - \$0, \$2-\$7, \$17

16-Bit Register Encoding ¹	32-Bit MIPS Register Encoding ²	Symbolic Name (From <i>ArchDefs.h</i>)	Description
0	0	zero	Hard-wired Zero
1	17	s1	General-purpose register
2	2	v0	General-purpose register
3	3	v1	General-purpose register
4	4	a0	General-purpose register
5	5	a1	General-purpose register
6	6	a2	General-purpose register
7	7	a3	General-purpose register

1. “0-7” correspond to the register’s 16-bit binary encoding and show how that encoding relates to the MIPS registers. “0-7” never refer to the registers, except within the binary microMIPS instructions. From the assembler, only the MIPS names (\$16, \$17, \$2, etc.) or the symbolic names (s0, s1, v0, etc.) refer to the registers. For example, to access register number 17 in the register file, the programmer references \$17 or s1, even though the microMIPS binary encoding for this register is 001.
2. General registers not shown in the above table are not accessible through the 16-bit instructions using 3-bit register specifier. The Move instruction can access all 32 general-purpose registers.

Table 14.6 16-Bit Instruction Implicit General-Purpose Registers

16-Bit Register Encoding	32-Bit MIPS Register Encoding	Symbolic Name (From <i>ArchDefs.h</i>)	Description
Implicit	28	gp	Global pointer register
Implicit	29	sp	Stack pointer register
Implicit	31	ra	Return address register

Table 14.7 16-Bit Instruction Special-Purpose Registers

Symbolic Name	Purpose
PC	Program counter. The PC-relative ADDIU can access this register as an operand.
HI	Contains high-order word of multiply or divide result.
LO	Contains low-order word of multiply or divide result.

14.3.3 32-Bit Category

14.3.3.1 New 32-bit instructions

The following table lists the 32-bit instructions introduced in the microMIPS ISA.

Table 14.8 32-bit Instructions introduced within microMIPS

Instruction	Major Opcode Name	Number of Register Fields	Immediate Field Size (bit)	Register Field Width (bit)	Total Size of Other Fields	Empty 0 Field Size (bit)	Minor Opcode Size (bit)	Comment
ADDIUPC	ADDIUPC	1	23	3		0	0	ADDIU PC-Relative
BEQZC	POOL32I	2:5 bit	16	5			0	Branch on Equal to Zero, No Delay Slot
BNEZC	POOL32I	2:5 bit	16	5			0	Branch on Not Equal to Zero, No Delay Slot
JALRS	POOL32A	2:5 bit	0	5			16	Jump and Link Register, Short Delay Slot
JALRS.HB	POOL32A	2:5 bit	0	5			16	Jump and Link Register with Hazard Barrier, Short Delay Slot
JALS	JALS32	0	26				0	Jump and Link, Short Delay Slot
JALX	JALX		26	5		0	5	Jump and Link Exchange
LWP	POOL32B	2:5 bit	12		5	0	4	Load Word Pair
LWXS	POOL32A	3:5 bit	0	5	0	1	10	Load Word Indexed, Scale
LWM32	POOL32B	1:5bit	12		5	0	4	Load Word Multiple
SWP	POOL32B	2:5 bit	12			0	4	Load Word Pair
SWM32	POOL32B	1:5bits	12		5	0	4	Store Word Multiple

References

This appendix lists other publications available from MIPS Technologies, Inc. that are referenced in this document. These documents may be included in the `$MIPS_PROJECT/doc` area of a typical M5100 soft or hard core release, or in some cases may be available on the MIPS web site <https://www.mips.com>.

1. MIPS32® M5100™ Processor Core Family Data Sheet
MIPS Document: MD00961
2. MIPS32® M5100™ Processor Core Family Integrator's Guide
MIPS Document: MD00963
3. MIPS32® M5100™ Processor Core Family Implementor's Guide
MIPS Document: MD00962
4. MIPS32® M5100™ Processor Core Family System Package & Simulation Flow User's Manual
MIPS Document: MD00965
5. *MIPS® Architecture For Programmers*, Volume I: Introduction to the MIPS32® Architecture
MIPS Document: MD0082
6. *MIPS® Architecture For Programmers*, Volume I: Introduction to the microMIPS32™ Architecture
MIPS Document: MD0741
7. *MIPS® Architecture For Programmers*, Volume II: The MIPS32® Instruction Set
MIPS Document: MD0086
8. *MIPS® Architecture For Programmers*, Volume II: The microMIPS32™ Instruction Set
MIPS Document: MD0582
9. *MIPS® Architecture For Programmers* Volume III: The MIPS32® and microMIPS32™ Privileged Resource Architecture
MIPS Document: MD00090
10. *MIPS® Architecture for Programmers* Volume IV-h: The MCU Application-Specific Extension to the MIPS32® Architectures
MIPS Document: MD00834
11. *MIPS® Architecture for Programmers* Volume IV-h: The MCU Application-Specific Extension to the microMIPS32™ Architectures
MIPS Document: MD00838
12. MIPS® EJTAG Specification
MIPS Document: MD00047

13. MIPS® cJTAG Adapter User's Manual
MIPS Document: MD00862
14. MIPS® Architecture for Programmers Volume IV-i: Virtualization Module of the MIPS32® Architecture
MIPS Document: MD00846
15. MIPS® Architecture for Programmers Volume IV-i: Virtualization Module of the microMIPS32™ Architecture
MIPS Document: MD00848
16. MIPS® Architecture Reference Manual Volume IV-e: The MIPS® DSP Module of the MIPS32® Architecture
MIPS Document: MD00372
17. MIPS® Architecture Reference Manual Volume IV-e: The MIPS® DSP Module of the microMIPS32® Architecture
MIPS Document: MD00762
18. Five Methods of Utilizing the MIPS® DSP Module
MIPS Document: MD00783
19. Efficient DSP Module Programming in C: Tips and Tricks
MIPS Document: MD00485
20. Accelerating DSP Filter Loops with MIPS® CorExtend® Instructions
MIPS Document: MD00303

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.

This document may refer to Architecture specifications (for example, instruction set descriptions and EJTAG register definitions), and change bars in these sections indicate changes since the previous version of the relevant Architecture document.

Revision	Date	Description
01.00	December 31, 2013	• Initial 1_0_0 release.
01.01	June 18, 2014	• Updated BP/BI bits in CP0 Config Register 3.
01.02	July 30, 2014	• Changes to timer interrupt input.
01.03	July 30, 2014	• Update IE bits in CP0 Config4 register.
01.04	March 15, 2015	• Update IE bits in CP0 Config4 register. See “Config4 Register (CP0 Register 16, Select 4)” on page 216.