



**SMP Linux Bring up on a MIPS32®
Coherent Processing System**

**Document Number: MD00721
Revision 01.01
February 7, 2013**

Unpublished rights (if any) reserved under the copyright laws of the United States of America and other countries.

This document contains information that is proprietary to MIPS Tech, LLC, a Wave Computing company ("MIPS") and MIPS' affiliates as applicable. Any copying, reproducing, modifying or use of this information (in whole or in part) that is not expressly permitted in writing by MIPS or MIPS' affiliates as applicable or an authorized third party is strictly prohibited. At a minimum, this information is protected under unfair competition and copyright laws. Violations thereof may result in criminal penalties and fines. Any document provided in source format (i.e., in a modifiable form such as in FrameMaker or Microsoft Word format) is subject to use and distribution restrictions that are independent of and supplemental to any and all confidentiality restrictions. UNDER NO CIRCUMSTANCES MAY A DOCUMENT PROVIDED IN SOURCE FORMAT BE DISTRIBUTED TO A THIRD PARTY IN SOURCE FORMAT WITHOUT THE EXPRESS WRITTEN PERMISSION OF MIPS (AND MIPS' AFFILIATES AS APPLICABLE) reserve the right to change the information contained in this document to improve function, design or otherwise.

MIPS and MIPS' affiliates do not assume any liability arising out of the application or use of this information, or of any error or omission in such information. Any warranties, whether express, statutory, implied or otherwise, including but not limited to the implied warranties of merchantability or fitness for a particular purpose, are excluded. Except as expressly provided in any written license agreement from MIPS or an authorized third party, the furnishing of this document does not give recipient any license to any intellectual property rights, including any patent rights, that cover the information in this document.

The information contained in this document shall not be exported, reexported, transferred, or released, directly or indirectly, in violation of the law of any country or international law, regulation, treaty, Executive Order, statute, amendments or supplements thereto. Should a conflict arise regarding the export, reexport, transfer, or release of the information contained in this document, the laws of the United States of America shall be the governing law.

The information contained in this document constitutes one or more of the following: commercial computer software, commercial computer software documentation or other commercial items. If the user of this information, or any related documentation of any kind, including related technical data or manuals, is an agency, department, or other entity of the United States government ("Government"), the use, duplication, reproduction, release, modification, disclosure, or transfer of this information, or any related documentation of any kind, is restricted in accordance with Federal Acquisition Regulation 12.212 for civilian agencies and Defense Federal Acquisition Regulation Supplement 227.7202 for military agencies. The use of this information by the Government is further restricted in accordance with the terms of the license agreement(s) and/or applicable contract terms and conditions covering this information from MIPS Technologies or an authorized third party.

MIPS, MIPS I, MIPS II, MIPS III, MIPS IV, MIPS V, MIPSr3, MIPS32, MIPS64, microMIPS32, microMIPS64, MIPS-3D, MIPS16, MIPS16e, MIPS-Based, MIPSsim, MIPSpro, MIPS-VERIFIED, Aptiv logo, microAptiv logo, interAptiv logo, microMIPS logo, MIPS Technologies logo, MIPS-VERIFIED logo, proAptiv logo, 4K, 4Kc, 4Km, 4Kp, 4KE, 4KEc, 4KEm, 4KEp, 4KS, 4KSc, 4KSd, M4K, M14K, 5K, 5Kc, 5Kf, 24K, 24Kc, 24Kf, 24KE, 24KEc, 24KEf, 34K, 34Kc, 34Kf, 74K, 74Kc, 74Kf, 1004K, 1004Kc, 1004Kf, 1074K, 1074Kc, 1074Kf, R3000, R4000, R5000, Aptiv, ASMACRO, Atlas, "At the core of the user experience.", BusBridge, Bus Navigator, CLAM, CorExtend, CoreFPGA, CoreLV, EC, FPGA View, FS2, FS2 FIRST SILICON SOLUTIONS logo, FS2 NAVIGATOR, HyperDebug, HyperJTAG, IASim, iFlowtrace, interAptiv, JALGO, Logic Navigator, Malta, MDMX, MED, MGB, microAptiv, microMIPS, Navigator, OCI, PDtrace, the Pipeline, proAptiv, Pro Series, SEAD-3, SmartMIPS, SOC-it, and YAMON are trademarks or registered trademarks of MIPS and MIPS' affiliates as applicable in the United States and other countries.

All other trademarks referred to herein are the property of their respective owners.

1	INTRODUCTION.....	4
2	YAMON AND LINUX BOOT UP FLOWS.....	5
2.1	YAMON BOOT UP FLOW.....	5
2.2	YAMON HIGH-LEVEL SET-UP FLOW FOR SMVP LINUX BOOT UP.....	11
2.3	LINUX HIGH-LEVEL SET-UP FLOW FOR SMVP LINUX BOOT UP.....	12
2.4	DETAILS OF CALLING FLOW IN SMP LINUX.....	13
3	Q & A.....	22
3.1	HOW DOES THE CODE DETERMINE WHICH CPU IS EXECUTING?.....	22
3.2	WHY IS CORE1 UNABLE TO ACKNOWLEDGE RESET?.....	22
3.3	WHAT DOES THE MESSAGE "UNABLE TO COMMUNICATE WITH CPU; PROCESSOR NEVER ACCESSED DMSEG" MEAN?.....	22
3.4	WHAT IS THE HIGH-LEVEL VIEW OF THE GCR, CPC, AND GIC ADDRESS MAPPING USED ON THE MALTA BOARD?.....	23
3.5	HOW IS CORE1 RELEASED FROM THE NAVIGATOR CONSOLE?.....	24
3.6	HOW DOES YAMON BIND TCs TO VPEs ON A COHERENT PROCESSING SYSTEM CORE?.....	24

1 Introduction

This application note describes how to boot up SMP Linux on a MIPS32® Coherent Processing System.

The application was developed using 1004K™ bit files executing on a Malta development platform.

- Linux kernel:
<http://www.linux-mips.org/pub/linux/mips/mti-stable/v2.6/linux-mti-2.6.35.9-2.tar.gz>
- Yamon source:
http://www.mips.com/secure-download/index.dot?product_name=/auth/yamon-src-02.21.tar.gz
- Bit files:
A00205-1004Kc-1_3c_0-2i-64ID-64TLB-noL2_ITU-REF00693 fl
A00206-1004Kc-1_3c_0-2i-64ID-64TLB-noL2_ITU-REF00693 fl
- Malta™ Development Board with 256 MB RAM

2 YAMON and Linux Boot Up Flows

2.1 YAMON Boot Up Flow

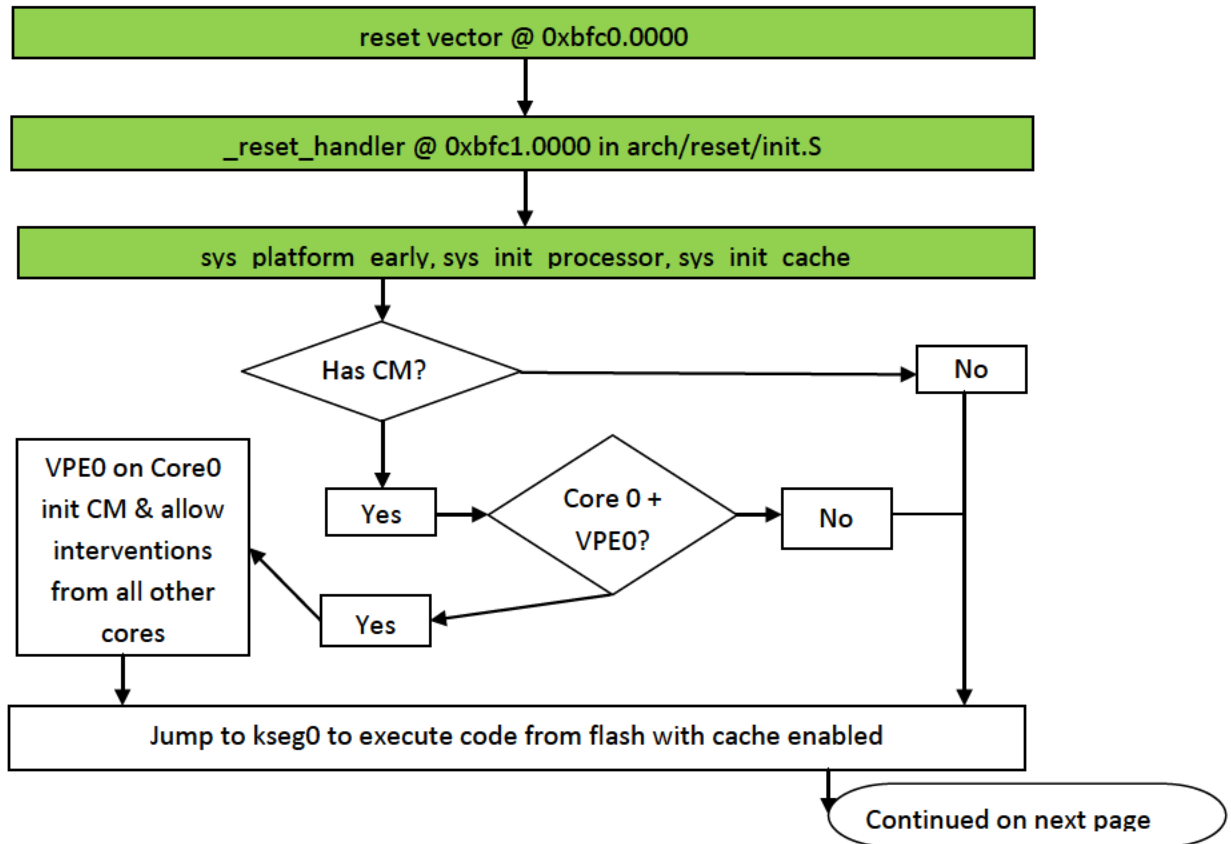
The Malta bit file is hardwired so that on cold reset, only Core0 in a multi-core cluster will power-up and begin executing code from the reset vector at address 0xbfc0.0000. Only VPE0 on Core0 will execute the code sequence in the flow chart and enter the shell. No other VPE/Core will execute YAMON code until VPE0 on Core0 powers-up the other cores in the `shell_cpu_init()` function.

YAMON is carefully coded such that only VPE0 on Core0 initializes the Coherence Manager (CM), Cluster Power Controller (CPC), and relocates code and data.

The `EBase.CPUNum` bit is used to identify which core is executing the code. In this case, VPE0 on Core0 has `EBase.CPUNum=0`. The CPU number is checked in the `sys_platform_early()` function. As a result, the value of `EBase.CPUNum` is stored in the V1 general-purpose register.

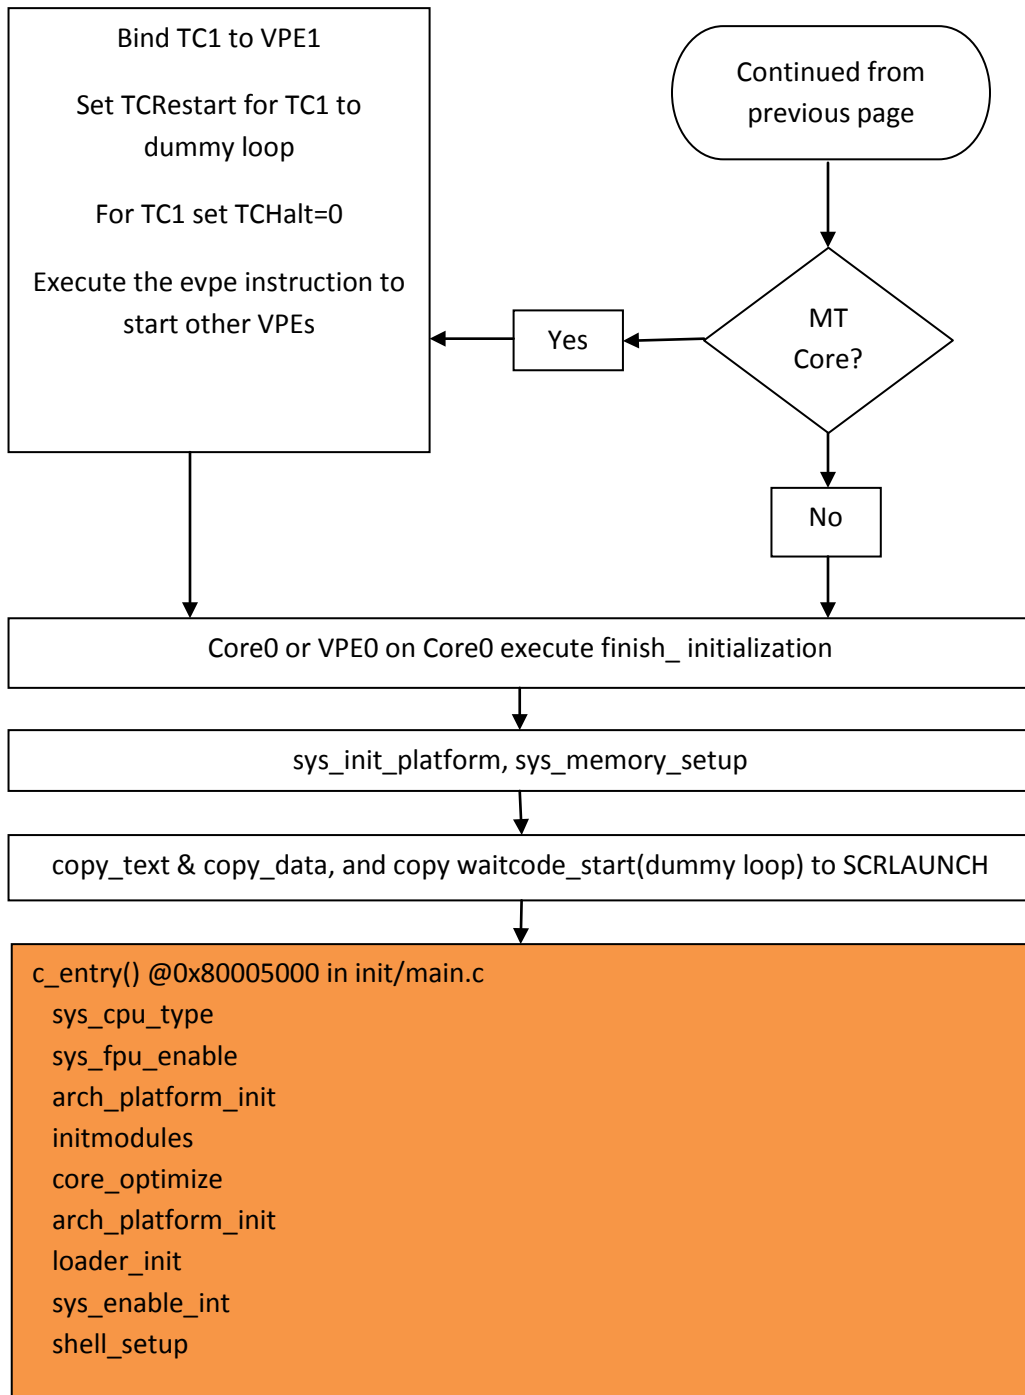
A flow chart of YAMON boot from the reset vector to the boot loader's shell is shown below. In the figure:

- A green background indicates that code is executed from flash with cache disabled.
- A white background indicates that code is executed from flash with cache enabled.
- An orange background indicates that code is executed from RAM with cache enabled.

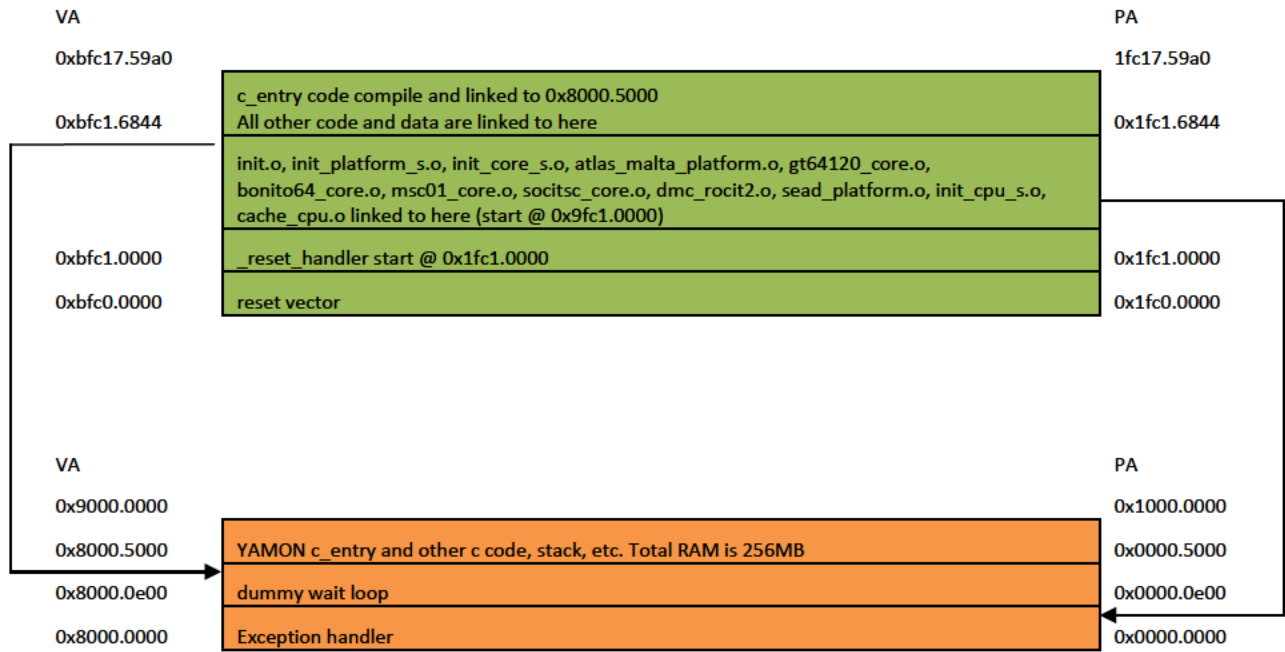


After CM has been initialized, YAMON initializes multi-threading (MT) by binding TC1 to VPE1.

VPE1 on Core0 does not begin executing code after the Coherent Processing system has been initialized by VPE0; VPE1 will be released later in `shell_cpu_init()`, which is required because the RAM has not been initialized, and the code has not been copied to the RAM. After multi-threading has been initialized, VPE0/TC0 will copy code and data to the RAM. For a non-MT core such as the 1074K™, the MT initialization mechanism is bypassed.



The high-level view of the memory map after the code relocation is shown below.



After the code relocation, YAMON will continue performing system-level initialization. Before entering the shell, VPE0 on Core0 will release other cores in the `gcmp_start_cores()` function. The calling flow is shown below. To release a core, it is powered-up by setting `CPC_CMD_REG.CMD=0x3`.

```
c_entry() in init/main.c
sys_cpu_type
sys_fpu_enable
arch platform init
initmodules
core_optimize
arch_platform_init
loader_init
sys enable int
shell_setup
```

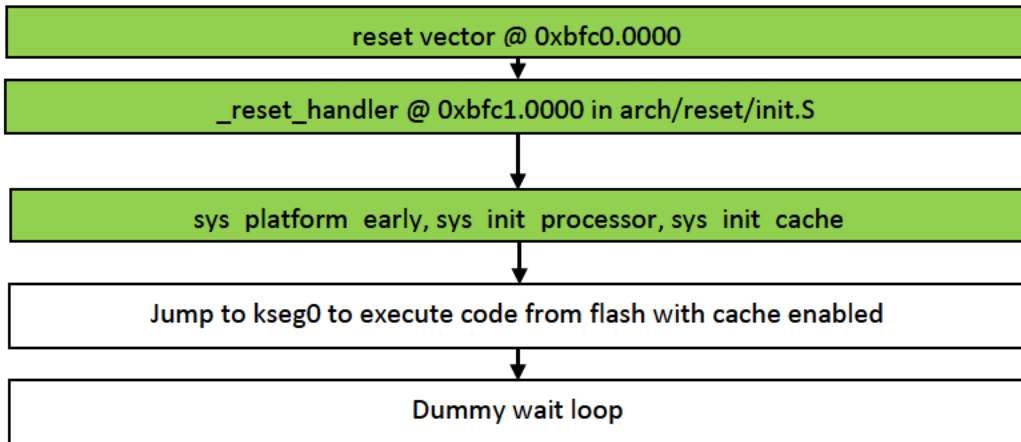
```
shell_setup(void) shell/shell_init.c
shell_arch
shell_help_init
shell( commands, command_count )
```

```
shell_arch(void) arch/shell/platform/shell_platform.c
shell_cksum_init
shell_register_cmd
:
:
shell_cpu_init()
```

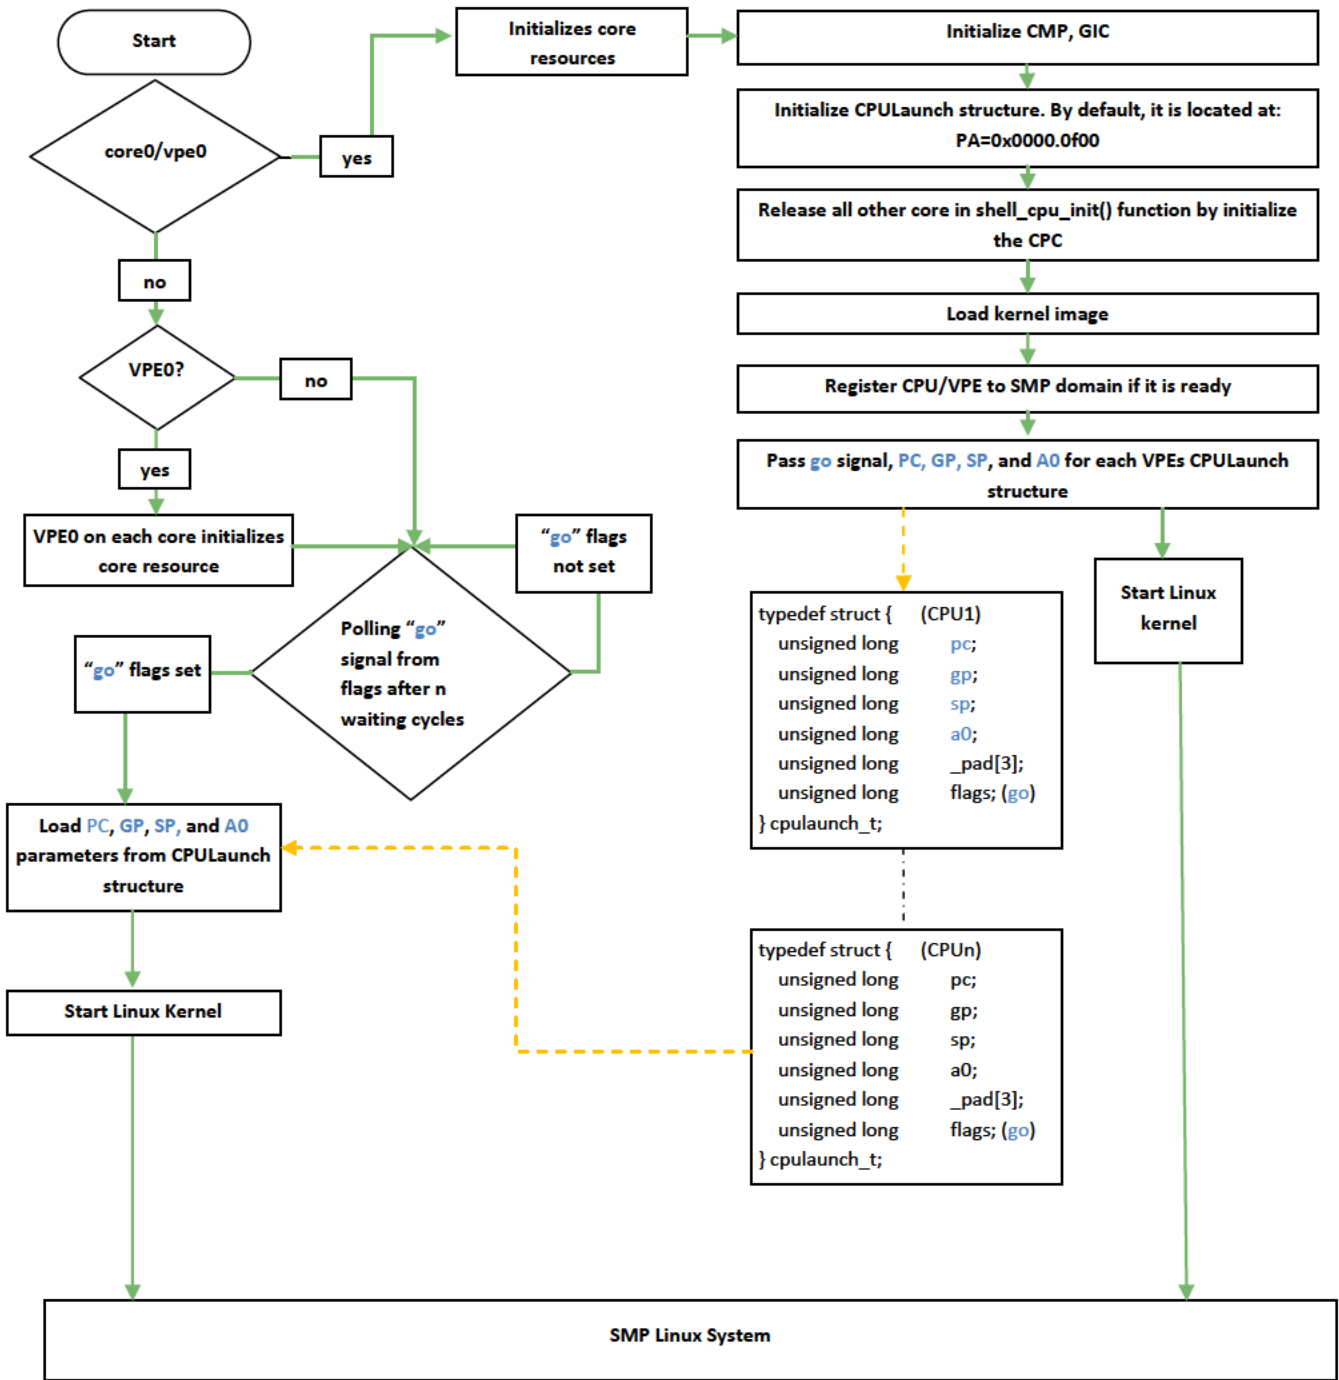
```
shell_cpu_init in shell/cpulaunch.c
release VPE1 by executing evpe instruction if MT is enabled
shell register cmd
:
:
gcmp_start_cores()
```

```
gcmp_start_cores in shell/cpulaunch.c
release other core by executing gcmp_requester(core_id)
Core released by program CPC_CMD_REG[CMD]=0x3
```

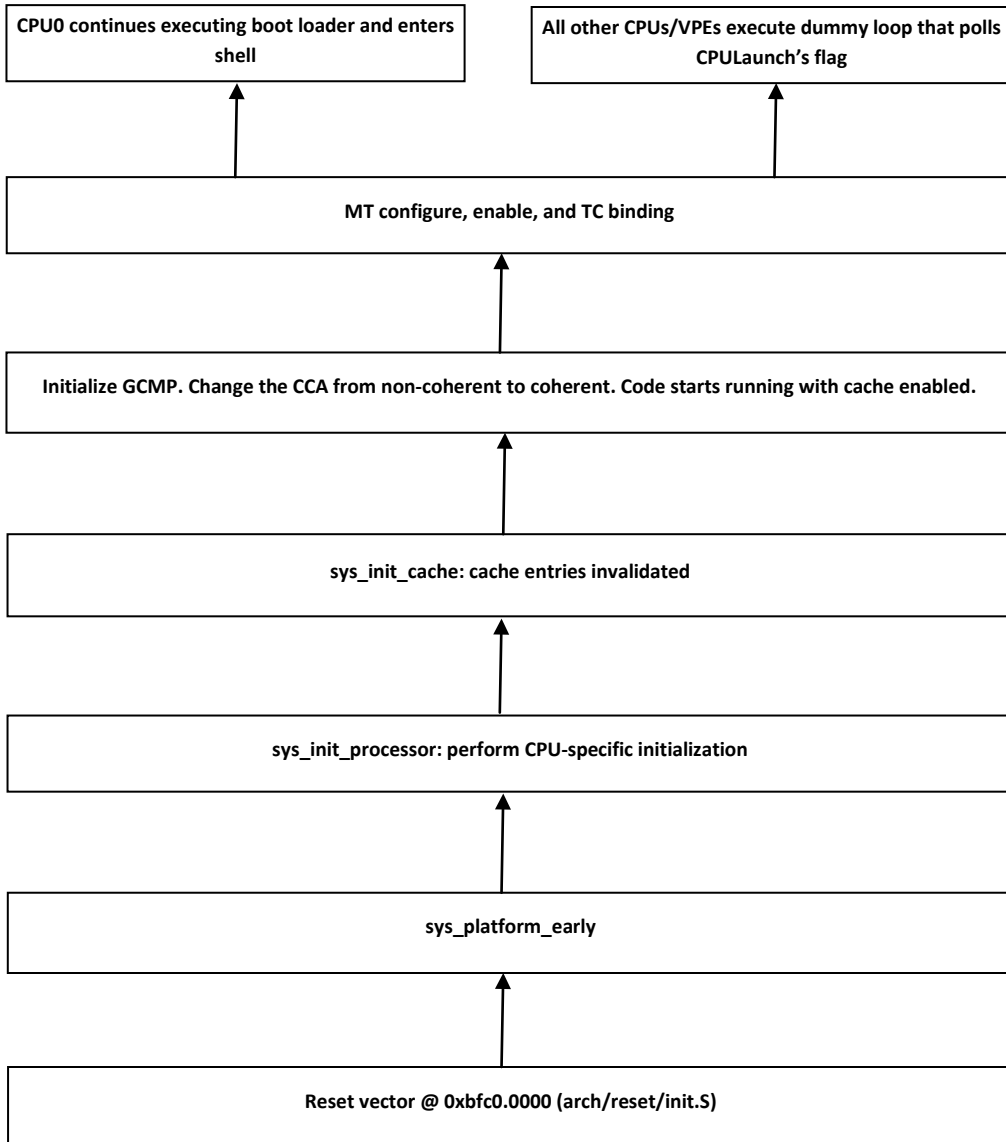

When all other cores in the multi-core cluster have powered-up, they will begin fetching code from the reset vector. The orange arrow shows how other cores/VPes (other than Core0) boot up and enter the dummy loop.



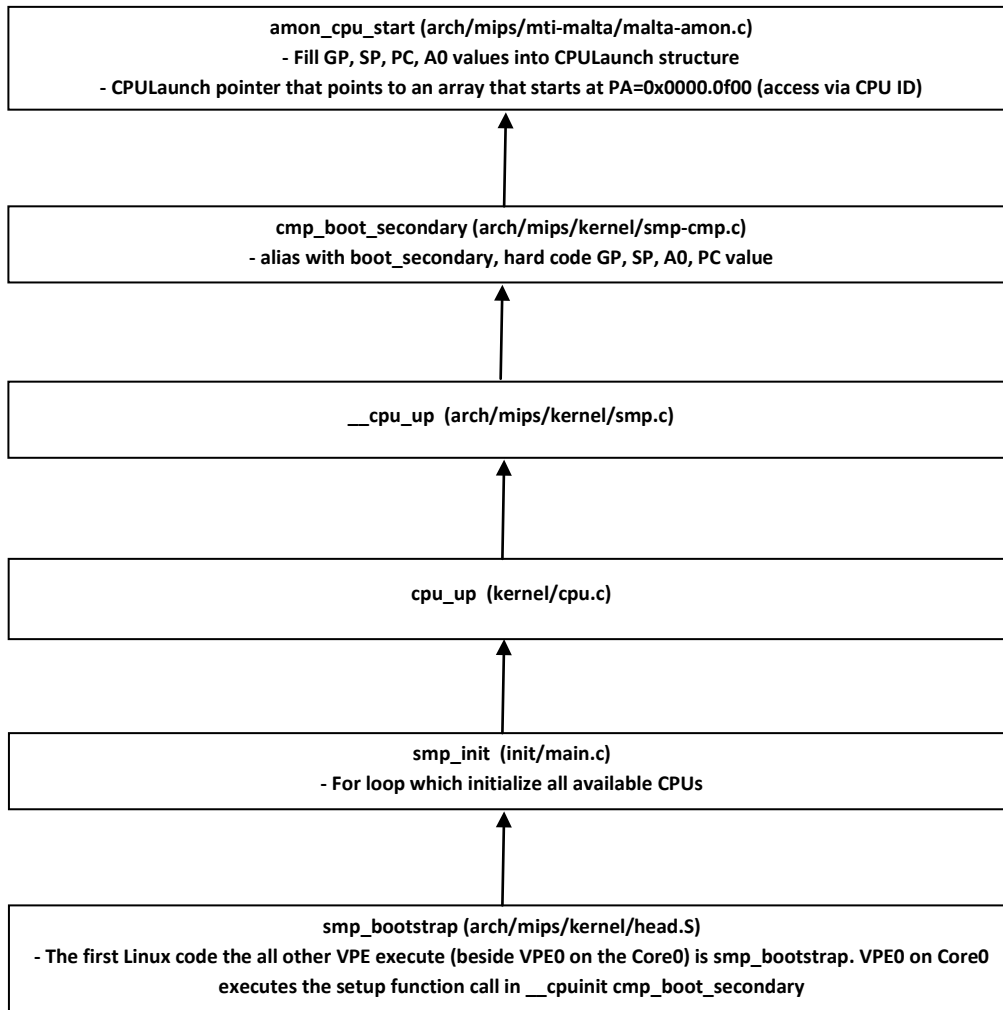
YAMON Bring up and Linux Handshake Flow Diagram



2.2 YAMON High-level Set-up Flow for SMVP Linux Boot Up



2.3 Linux High-level Set-up Flow for SMVP Linux Boot Up



2.4 Details of Calling Flow in SMP Linux

Below is the detailed SMP Linux calling flow. The orange background color indicates that the functions are executed by VPE0 on Core0 (or Core0 for non-MT cores). The blue background indicates that the call is executed only by other VPEs/CPU's.

Function name & calling flow	File where function is defined	Console output
<code>_kernel_entry ()</code>	<code>./arch/mips/kernel/head.S</code>	
<code>init_start_kernel ();</code>	<code>./init/main.c</code>	
<code>smp_setup_processor_id()</code>		
<code>lockdep_init();</code>		
<code>debug_objects_early_init();</code>		
<code>boot_init_stack_canary();</code>		
<code>cgroup_init_early();</code>		
<code>local_irq_disable();</code>		
<code>early_boot_irqs_off();</code>		
<code>early_init_irq_lock_class();</code>		
<code>lock_kernel();</code>		
<code>tick_init();</code>		
<code>boot_cpu_init();</code>		
<code>set_cpu_online(cpu, true);</code>		
<code>set_cpu_active(cpu, true);</code>		
<code>set_cpu_present(cpu, true);</code>		
<code>set_cpu_possible(cpu, true);</code>		
<code>page_address_init();</code>		
<code>setup_arch(&command_line);</code>	<code>arch/mips/kernel/setup.c</code>	
<code>cpu_probe();</code>	<code>arch/mips/kernel/cpu-probe.c</code>	
<code>cpu_probe_mips(c, cpu);</code>	<code>arch/mips/kernel/cpu-probe.c</code>	
<code>decode_configs();</code>	<code>arch/mips/kernel/cpu-probe.c</code>	
<code>spram_config()</code>		
<code>cpu_probe_vmbits(c);</code>		
<code>prom_init</code>	<code>arch/mips/mti-malta/malta-init.c</code>	Display "Linux" on LED and "LINUX started..." to console"
<code>prom_init_cmdline();</code>	<code>arch/mips/mti-malta/malta-cmdline.c</code>	
<code>prom_meminit();</code>	<code>arch/mips/mti-malta/malta-memory.c</code>	
<code>console_config();</code>	<code>arch/mips/mti-malta/malta-init.c</code>	
<code>gcmp_probe(GCMP_BASE_ADDR, GCMP_ADDRSPACE_SZ);</code>		
<code>register_smp_ops(&cmp_smp_ops);</code> or <code>register_smp_ops(&vsmp_smp_ops);</code>		

Function name & calling flow		File where function is defined	Console output
	setup_early_printk();	arch/mips/kernel/early_printk.c	
	register_console(&early_console);	./kernel/printk.c	Display "bootconsole [early0] enabled"
	cpu_report();	arch/mips/kernel/cpu-probe.c	Display ""CPU revision is:"
	check_bugs_early();		
	arch_mem_init(cmdline_p);	arch/mips/kernel/setup.c	Print Memory map "Determined physical RAM map"
	plat_mem_setup();		
	print_memory_map();		Print Memory map " memory: 00001000 @ 00000000 (reserved)"
	parse_early_param();		
	bootmem_init();	arch/mips/kernel/setup.c	Display "Wasting 44064 bytes for tracking 1377 unused pages"
	init_initrd();		
	init_bootmem_node();		
	reserve_bootmem(PFN_PHYS(mapstart), bootmap_size, BOOTMEM_DEFAULT);		
	finalize_initrd();		Display "Initrd not found or empty - disabling initrd"
	sparse_init();		
	paging_init();	./arch/mips/mm/init.c	
	pagetable_init();		
	kmap_init();		
	kmap_coherent_init();		
	free_area_init_nodes	./mm/page_alloc.c	Display "Zone PFN ranges:" "Movable zone start PFN for each node" "early node map[1] active PFN ranges"
	resource_init();		
	plat_smp_setup();	arch/mips/include/asm/smp-ops.h	
	cmp_smp_setup();	arch/mips/kernel/smp-cmp.c	Display "Detected 5 available secondary CPU(s)"
	mm_init_owner(&init_mm, &init_task);		
	setup_command_line(command_line);		
	setup_nr_cpu_ids();		
	setup_per_cpu_areas();	./mm/percpu.c	
	pcpu_embed_first_chunk	./mm/percpu.c	Display "PERCPU: Embedded 7 pages/cpu @81203000 s5888 r8192 d14592 u65536"
	pcpu_setup_first_chunk	./mm/percpu.c	
	pcpu_dump_alloc_info	./mm/percpu.c	Display "pcpu-alloc: s5888 r8192 d14592 u65536 alloc=16*4096" "pcpu-alloc: [0] 0 [0] 1 [0] 2 [0] 3 [0] 4 [0] 5"
	pcpu_free_alloc_info	./mm/percpu.c	
	free_bootmem	./mm/bootmem.c	
	smp_prepare_boot_cpu();	./arch/mips/kernel/smp.c	

Function name & calling flow	File where function is defined	Console output
set_cpu_possible(0, true);		
set_cpu_online(0, true);		
cpu_set(0, cpu_callin_map);		
build_all_zonelists();	./mm/page_alloc.c	
build_zonelists(pgdat);	./mm/page_alloc.c	Display "Built 1 zonelists in Zone order, mobility grouping on. Total pages: 65024"
build_zonelist_cache(pgdat);		
page_alloc_init();		
printk(KERN_NOTICE "Kernel command line: %s\n", boot_command_line);		Display "Kernel command line: init=/init ip=dhcp console=ttyS0,38400n8r"
parse_early_param();		
parse_args("Booting kernel", static_command_line, __start__param, __stop__param - start_param, &unknown_bootoption);		
pidhash_init();	./kernel/pid.c	
alloc_large_system_hash	./mm/page_alloc.c	Display "PID hash table entries: 1024 (order: 0, 4096 bytes)"
vfs_caches_init_early();	./fs/dcache.c	
dcache_init_early();	./fs/dcache.c	
alloc_large_system_hash	./mm/page_alloc.c	Display "Dentry cache hash table entries: 32768 (order: 5, 131072 bytes)"
INIT_HLIST_HEAD		
inode_init_early();	./fs/inode.c	
alloc_large_system_hash	./mm/page_alloc.c	Display "Inode-cache hash table entries: 16384 (order: 4, 65536 bytes)"
INIT_HLIST_HEAD		
sort_main_extable();	./kernel/extable.c	
sort_extable();	./lib/extable.c	
trap_init();	arch/mips/kernel/traps.c	
per_cpu_trap_init();	arch/mips/kernel/traps.c	
change_c0_status		
cpu_cache_init();	arch/mips/mm/cache.c	
r4k_cache_init();	arch/mips/mm/c-r4k.c	
probe_pcachel();		Display "Primary instruction cache 32kB, 4-way, VIPT, linesize 32 bytes." "Primary data cache 32kB, 4-way, PIPT, no aliases, linesize 32 bytes"
setup_scachel();		Display "MIPS secondary cache 512kB, 8-way, linesize 64 bytes."
r4k_blast_dcache_page_setup();		
r4k_blast_dcache_page_indexed_setup();		
r4k_blast_dcache_setup();		

Function name & calling flow		File where function is defined	Console output
	r4k_blast_icache_page_setup();		
	r4k_blast_icache_page_indexed_setup();		
	r4k_blast_icache_setup();		
	r4k_blast_scache_page_setup();		
	r4k_blast_scache_page_indexed_setup();		
	r4k_blast_scache_setup();		
	build_clear_page();		
	build_copy_page();		
	local_r4k_flush_cache_all(NULL);		
	coherency_setup();		
	tlb_init();		
	set_handler(0x180, &except_vec3_generic, 0x80);		
	set_except_vector(i, handle_reserved);		
	board_ejtag_handler_setup();		
	set_except_vector(23, handle_watch);		
	set_vi_handler(i, NULL);		
	parity_protection_init();	arch/mips/kernel/traps.c	Display "Writing ErrCtl register=00000000" "Readback ErrCtl register=00000000"
	board_be_init();		
	board_nmi_handler_setup();		
	signal_init();/signal32_init();		
	local_flush_icache_range(ebase, ebase + 0x400);		
	flush_tlb_handlers();		
	mm_init();	./init/main.c	
	page_cgroup_init_flatmem();		
	mem_init();	arch/mips/mm/init.c	Display "Memory: 252512k/255328k available (3412k kernel code, 2452k reserved, 851k data, 1376k init, 0k highmem)"
	calculate_totalram_pages, num_physpages, totalhigh_pagescodesize, datasize, initsize		
	kmem_cache_init();		
	pgtable_cache_init();		
	vmalloc_init();		
	sched_init();		
	preempt_disable();		
	rcu_init();	./kernel/rcupdate.c	
	__rcu_init	./kernel/rcutree.c	
	rcu_bootup_announce	./kernel/rcutree_plugin.h	Display "Hierarchical RCU implementation."
	__rcu_init_preempt		

Function name & calling flow	File where function is defined	Console output
open softirq(RCU_SOFTIRQ, rcu_process_callbacks);		
cpu_notifier(rcu_barrier_cpu_hotplug, 0);		
early_irq_init();		Display "NR_IRQS:256"
init_IRQ();	arch/mips/kernel/irq.c	
set_irq_noprobe(i)		
arch_init_irq();	arch/mips/mti-malta/malta-int.c	Display "CPU0: status register was 11002400" "CPU0: status register now 11002400" "CPU0: status register frc 11003c00"
init_i8259_irqs		
mips_cpu_irq_init();	arch/mips/kernel/irq_cpu.c	
init_msc_irqs		
set_vi_handler(MIPSCPU_INT_I8259A, malta_hw0_irqdispatch);		
set_vi_handler(MIPSCPU_INT_COREHI, corehi_irqdispatch);		
setup_irq(MIPS_CPU_IRQ_BASE+MIPSCPU_INT_I8259A, &i8259irq);		
setup_irq(MIPS_CPU_IRQ_BASE+MIPSCPU_INT_COREHI, &corehi_irqaction);		
fill_ipi_map();		
gic_init(GIC_BASE_ADDR, GIC_ADDRSPACE_SZ, gic_intr_map, ARRAY_SIZE(gic_intr_map), MIPS_GIC_IRQ_BASE);		
prio_tree_init();		
init_timers();		
hrtimers_init();		
softirq_init();		
timekeeping_init();		
time_init();	./arch/mips/kernel/time.c	Display "CPU frequency 24.99 MHz"
plat_time_init();	./arch/mips/mti-malta/malta-time.c	
estimate_cpu_frequency();		
mips_scroll_message();		
plat_perf_setup();		
mips_clockevent_init()		
init_mips_clocksource()		
profile_init();		
early_boot_irqs_on();		
local_irq_enable();		
set_gfp_allowed_mask(_GFP_BITS_MASK);		
kmem_cache_init_late();		
console_init();	drivers/char/tty_io.c	
con_initcall_start		
con_init()	drivers/char/vt.c	Display "Console: colour dummy device 80x25"

Function name & calling flow	File where function is defined	Console output
lockdep info();		
locking_selftest();		
page_cgroup init();		
enable_debug_pagealloc();		
kmemtrace init();		
kmemleak_init();		
debug_objects mem init();		
idr_init_cache();		
setup_per_cpu_pageset();		
numa_policy_init();		
if (late time init) late time init();		
sched_clock_init();		
calibrate_delay();	./init/calibrate.c	Display "Calibrating delay loop... 200.25 BogoMIPS (lpj=2093056)"
pidmap init();		
anon_vma_init();		
thread_info cache init();		
cred_init();		
fork_init(totalram pages);		
proc_caches_init();		
buffer init();		
key_init();		
security init();		
vfs_caches_init(totalram_pages);		
dcache init();		
inode_init();		
files init(mempages);		
mnt_init();	./fs/namespace.c	Display "Mount-cache hash table entries: 512"
bdev_cache init();		
chrdev_init();		
radix tree init();		
signals_init();		
page_writeback init();		
cgroup_init();		
cpuset init();		
taskstats_init_early();		
delayacct init();		

Function name & calling flow	File where function is defined	Console output
check_bugs();		
acpi_early_init();		
sfi_init_late();		
ftrace_init();		
rest_init();	./init/main.c	
rcu_scheduler_starting();		
kernel_thread(kernel_init, NULL, CLONE_FS CLONE_SIGHAND);		
kernel_init	./init/main.c	
smp_prepare_cpus(setup_max_cpus);		
cmp_prepare_cpus	arch/mips/kernel/smp-cmp.c	
mips_mt_set_cpuoptions		
do_pre_smp_initcalls();		
initcall_start		
start_boot_trace();		
smp_init();	./init/main.c	
cpu_up(cpu);	./kernel/cpu.c	
cpu_maps_update_begin();		
_cpu_up(cpu, 0);	./kernel/cpu.c	
cpu_hotplug_begin();		
raw_notifier_call_chain(&cpu_chain, CPU_UP_PREPARE mod, hcpu, -1, &nr_calls);		
__cpu_up(cpu);	arch/mips/kernel/smp.c	
fork_idle(cpu);		
boot_secondary(cpu, idle);	arch/mips/kernel/smp-cmp.c	
cmp_boot_secondary(cpu, idle);	arch/mips/kernel/smp-cmp.c	
amon_cpu_start(cpu, pc, sp, (unsigned long)gp, a0);	arch/mips/mti-malta/malta-amon.c	
update_struct_cpulaunch, each core start execute smp_bootstrap()		
smp_bootstrap()	arch/mips/kernel/head.S	
mips_ihb	arch/mips/kernel/entry.S	
setup_c0_status_sec	arch/mips/kernel/head.S	
smp_slave_setup	arch/mips/kernel/head.S	
start_secondary	arch/mips/kernel/smp.c	
cpu_probe()	arch/mips/kernel/cpu-probe.c	
cpu_report();		
per_cpu_trap_init();		
mips_clockevent_init();		

Function name & calling flow	File where function is defined	Console output
mp_ops->init_secondary();		
calibrate_delay();		
notify_cpu_starting(cpu);		
mp_ops->smp_finish();		
set_cpu_sibling_map(cpu);		
cpu_set(cpu, cpu_callin_map);		
synchronise_count_slave();	arch/mips/kernel/sync-r4k.c	
cpu_idle();		
cpu_isset(cpu,cpu_callin_map)		
udelay(100);		
cpu_set(cpu,cpu_online_map);		
set_cpu_active(cpu,true);		
raw_notifier_call_chain(&cpu_chain,CPU_ONLINE mod, hcpu);		
cpu_hotplug_done();		
cpu_maps_update_done();		
printk(KERN_INFO "Brought up %ld CPUs\n", (long)num_online_cpus());		Display "Brought up 6 CPUs"
smp_cpus_done(setup_max_cpus);	arch/mips/kernel/smp.c	
mp_ops->cpus_done();	arch/mips/kernel/smp-cmp.c	
cmp_cpus_done	arch/mips/kernel/smp-cmp.c	
synchronise_count_master();	arch/mips/kernel/sync-r4k.c	
sched_init_smp();		
do_basic_setup();		
numa_default_policy();		
kernel_thread(kthreadd, NULL, CLONE_FS CLONE_FILES);		
find_task_by_pid_ns(pid, &init_pid_ns);		
unlock_kernel();		
init_idle_bootup_task(current);		
preempt_enable_no_resched();		
schedule();		
preempt_disable();		
cpu_idle();		

NOTES:

1. The function `asmlinkage __cpuinit void start_secondary(void)` in `arch/mips/kernel/smp.c` is the first C code executed by all secondary CPUs.
2. The first function executed by the primary CPU is `__kernel_entry ()`. The first function executed by the secondary CPU is `smp_bootstrap ()`.
3. The functions `synchronise_count_master ()` and `synchronise_count_slave ()` are used to synchronize the clock source among primary and secondary CPUs. Two possible clock sources can be used: the traditional Count/Compare from each core or the global counter in GIC.

3 Q & A

3.1 How does the code determine which CPU is executing?

The code can check the *CPUNum* field in the CP0 *EBase* register. For a two-core/four-VPE configuration, VPE0 in Core0 has a value of 0. The other values are:

- VPE1 on Core0 has value of 1
- VPE0 on Core1 has value of 2
- VPE1 on Core1 has value of 3

On two-core two-VPE configurations:

- VPE0 on Core0 has value of 0
- VPE0 on Core1 has value of 1

3.2 Why is Core1 unable to acknowledge reset?

When using the Navigator Console to reset the target (EJTAGBOOT), the console will display “Unable to acknowledge reset”. This indicates that the CPC did not release Core1 in order that both VPE0 and VPE1 on Core1 cannot respond to the reset request from JTAG until the code issues a power-up command.



```
mips_1004k_keng.tcl - MIPS Navigator Console
File Console Edit Interp Prefs History Help Tools

(c0v0) 2 %reset
c0v0: Reset detected.
c0v0: Reset acknowledged.
c0v1: Reset detected.
c0v1: Reset acknowledged.
c0v0: Halted at 0xBFC00000.
0xBFC00000 10000005 b 0xbfc00018
c1v0: Reset detected.
c1v0: Unable to acknowledge reset.
c1v1: Reset detected.
c1v1: Unable to acknowledge reset.
c0v1: Halted at 0x80000E30.
(c0v0) 3 %
```

3.3 What does the message “Unable to communicate with CPU; Processor never accessed DMSEG” mean?

When using the System Navigator probe to read/write the CPC register, and the error “Unable to communicate with CPU; Processor never accessed DMSEG” appears on the console, most likely the CPC is not enabled. Set *GCR_CPC_BASE.CPC_EN* to 0x1 to enable the CPC unit.



```
mips_1004k_keng.tcl - MIPS Navigator Console
File Console Edit Interp Prefs History Help Tools

(c0v0) 9 %word 0xbbde2000
Unable to communicate with CPU; Processor never accessed DMSEG.
(c0v0) 10 %
(c0v0) 10 %
```


3.4 What is the high-level view of the GCR, CPC, and GIC address mapping used on the Malta board?

Below is the address mapping used by YAMON on the Malta board for the GCR, CPC, and GIC. The GCR is 64KB from the CPC, and the CPC is mapped next to the GIC.

	Start Addr	Size	Name	Note
Global Control Registers	0x1fbf.e000	8K	Global Debug Block	By default based address is hardwired to 0x1fbf.8000. Can be reprogramed by SW via GCR Base Register[GCR_BASE].
	0x1fbf.c000	8K	Core-Other Control Block	
	0x1fbf.a000	8K	Core-Local Control Block	
	0x1fbf.8000	8K	Global Control Block	
	0x1bde.8000	64K	Unused	
Cluster Power Controller	0x1bde.6000	8K	Unused	Based address is set by Cluster Power Controller Base Address Register[CPC_BaseAddress] bits.
	0x1bde.4000	8K	Core-Other Control Block	
	0x1bde.2000	8K	Core-Local Control Block	
	0x1bde.0000	8K	Global Control Block	
Global Interrupt Controller	0x1bdd.0000	64K	User-Mode Visible Section	Based address is set by Global Interrupt Controller Base Address Register [GIC_BaseAddress] bits.
	0x1bdc.c000	16K	VPE-Other Section	
	0x1bdc.8000	16K	VPE-Local Section	
	0x1bdc.0000	32K	Shared Section	

3.5 How is Core1 released from the Navigator Console?

To release Core1 from the System Navigator probe, enter the commands shown below in the Navigator Console.



The first command, `word 0xbfbf8088 0x1bde0001`, writes `0x1bde0001` to `0xbfbf8088` in KSEG1, which is the address of the *Cluster Power Controller Base Address Register (GCR_CPC_BASE)*. The value `0x1bde0001` is used to set the physical address of the CPC to `0x1bde 0000` and enable the CPC:

```
GCR_CPC_BASE[CPC_BaseAddress] = 0x1bde0
GCR_CPC_BASE [CPC_EN] = 0x1
```

The second command, `word 0xbbde2010 0x00010000`, writes `0x00010000` to the KSEG1 address `0xbbde 2010`, which is the address of *CPC_OTHER_REG*. The value of `0x00010000` is used to set *CORENUM* to `0x1` in that register.

```
CPC_OTHER_REG[CORENUM]= 0x1
```

This programs the CPU IDs of the other cores so that the *Core Other Addressing Register* can be used to power-up that core.

The third command powers up Core 1. The command `word 0xbbde4000 0x00000003` writes `0x0000 0003` to the KSEG1 address `0xbbde 4000`, which is the address of the Other Cores Cluster Power Controller's command register. The value of `0x00000003` writes the power-up command to that register and powers up the core.

```
CPC_CMD_REG[CMD]=0x3
```

3.6 How does YAMON bind TCs to VPEs on a Coherent Processing System core?

By default, YAMON binds TC0 to VPE0 and binds all other TCs to VPE1 on each core. Only TC1 on VPE1 runs the dummy polling loop. All other TCs remain inactive.