



# **Programming the MIPS32® 34K™ Core Family**

**Document Number: MD00427**

**Revision 01.30**

**May 25, 2006**

**MIPS Technologies, Inc.  
1225 Charleston Road  
Mountain View, CA 94043-1353**

**Copyright © 2004-2006 MIPS Technologies Inc. All rights reserved.**

Copyright © 2004-2006 MIPS Technologies, Inc. All rights reserved.

Unpublished rights (if any) reserved under the copyright laws of the United States of America and other countries.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 Technologies, Inc. ("MIPS Technologies"). Any copying, reproducing, modifying or use of this information (in whole or in part) that is not expressly permitted in writing by MIPS Technologies 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 TECHNOLOGIES, INC.

MIPS Technologies reserves the right to change the information contained in this document to improve function, design or otherwise. MIPS Technologies does 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 Technologies 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, MIPS-3D, MIPS16, MIPS16e, MIPS32, MIPS64, MIPS-Based, MIPSsim, MIPSpro, MIPS Technologies logo, MIPS RISC CERTIFIED POWER logo, MIPS-VERIFIED, 4K, 4Kc, 4Km, 4Kp, 4KE, 4KEc, 4KEm, 4KEp, 4KS, 4KSc, 4KSd, M4K, 5K, 5Kc, 5Kf, 20K, 20Kc, 24K, 24Kc, 24Kf, 24KE, 24KEc, 24KEf, 25Kf, 34K, 34Kc, 34Kf, R3000, R4000, R5000, ASMACRO, Atlas, "At the core of the user experience.", BusBridge, CorExtend, CoreFPGA, CoreLV, EC, JALGO, Malta, MDMX, MGB, PDtrace, the Pipeline, Pro Series, QuickMIPS, SEAD, SEAD-2, SmartMIPS, SOC-it, and YAMON are trademarks or registered trademarks of MIPS Technologies, Inc. in the United States and other countries.

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

Template: nB1.00, Built with tags: 2B

Programming the MIPS32® 34K™ Core Family, Revision 01.30

**Copyright © 2004-2006 MIPS Technologies Inc. All rights reserved.**

# Table of Contents

<b>Chapter 1: Introduction .....</b>	<b>13</b>
1.1: Chapter summary .....	13
1.2: Typographical conventions .....	14
1.3: Finding information in this manual .....	14
1.4: Key features of the 34K™ core .....	14
1.5: Specification summary .....	15
1.6: Pipeline and implementation .....	16
<b>Chapter 2: The MIPS® MT ASE - Multithreading the RISC way .....</b>	<b>17</b>
2.1: What's a thread and its context? .....	18
2.2: Why multi-threading? .....	18
2.3: Different kinds of multi-threading: TCs and VPEs .....	19
2.3.1: How an MT CPU's hardware uses TCs and VPEs .....	19
2.3.2: CPU resources and registers shared between all threads .....	19
2.3.3: CPU resources and registers replicated per-TC .....	21
2.3.4: CPU resources and registers replicated per-VPE .....	22
2.4: When can't threads run? .....	22
2.5: Thread-scheduling decisions and the policy manager .....	24
2.6: Multithreading, exceptions and interrupts .....	24
2.6.1: Multithreading and interrupts .....	25
2.7: Multithreading, non-blocking loads and stores, and gating storage .....	26
2.7.1: Gating storage .....	26
2.8: MIPS® Multithreading ASE - new instructions .....	27
2.8.1: Yield, Yield Qualifiers and threads waiting for hardware events .....	30
2.8.2: All MT instructions in alphabetical order .....	30
2.9: Multithreading ASE - CP0 (privileged) registers .....	33
2.9.1: What CP0 registers are per-TC, per-VPE and per-CPU? .....	33
2.9.2: VPEControl .....	34
2.9.3: TCRestart, TCHalt and TCContextt .....	35
2.9.4: TCStatus .....	35
2.9.5: TCBind .....	37
2.9.6: MVPConf0-1 - read-only multithreading-specific configuration information .....	37
2.9.7: MVPControl Register - CPU-wide VPE control .....	38
2.9.8: VPEConf0-1 registers - initializable per VPE resource lists .....	39
2.9.9: YQMask register - enable yield "conditions" .....	39
2.9.10: VPEOpt register - reserve some cache "way" for use of one VPE .....	40
2.9.11: Shadow register configuration SRSCConf0-4 .....	40
2.9.12: Thread scheduling hints - TCSchedule, TCScheFBack, VPESchedule .....	40
<b>Chapter 3: How the 34K™ core implements multi-threading .....</b>	<b>41</b>
3.1: The 34K™ core pipeline and multithreading .....	41
3.1.1: Resource limits and consequences .....	43
3.1.2: Choosing instruction to issue .....	43
3.2: Thread scheduling in the 34K core .....	44
3.2.1: The Dispatch Scheduler .....	44
3.2.2: Policy manager interface .....	44
3.2.3: Policy managers available for the 34K™ core family .....	45

3.3: Inter-thread communication storage (ITC).....	47
3.3.1: Configuring ITC base address and cell repeat interval .....	48
3.4: The 34K™ core and interrupts .....	49
3.5: Synchronization: "ll" and "sc" instructions implementation .....	49
<b>Chapter 4: Initializing the 34K™ core - Multi-Threaded bootstrap issues.....</b>	<b>51</b>
4.1: Bootstrapping without worrying about multithreading.....	51
4.2: Configuring your choice of VPEs and TCs .....	52
4.2.1: Setting up a VPE for legacy software.....	53
4.2.2: Sharing and not sharing the TLB .....	53
4.3: Setting up a TC to run a thread .....	55
4.4: TCs recycled as Shadow registers .....	55
<b>Chapter 5: The MIPS32® DSP ASE .....</b>	<b>57</b>
5.1: Features provided by the MIPS® DSP ASE .....	57
5.2: The DSP ASE control register .....	58
5.2.1: DSP accumulators .....	59
5.3: Software detection of the DSP ASE .....	59
5.4: DSP instructions .....	60
5.4.1: Hints in instruction names .....	60
5.4.2: Arithmetic - 64-bit .....	61
5.4.3: Arithmetic - saturating and/or SIMD Types .....	61
5.4.4: Bit-shifts - saturating and/or SIMD types.....	61
5.4.5: Comparison and "conditional-move" operations on SIMD types.....	61
5.4.6: Conversions to and from SIMD types .....	62
5.4.7: Multiplication - SIMD types with result in GP register .....	62
5.4.8: Multiply Q15s from paired-half and accumulate .....	63
5.4.9: Load with register + register address .....	63
5.4.10: DSPControl register access .....	63
5.4.11: Accumulator access instructions.....	64
5.4.12: Dot products and building blocks for complex multiplication.....	64
5.4.13: Other DSP ASE instructions .....	65
5.5: Macros and typedefs for DSP instructions .....	66
5.6: Almost Alphabetically-ordered table of DSP ASE instructions .....	67
5.7: DSP ASE instruction timing .....	70
<b>Chapter 6: Memory map, caching, reads and writes and translation .....</b>	<b>71</b>
6.1: The memory map .....	71
6.2: Fixed mapping option .....	72
6.3: Reads, writes and synchronization.....	72
6.3.1: Read/write ordering and queues in the 34K core.....	72
6.3.2: The "sync" instruction in 34K.....	73
6.3.3: Write gathering and "write buffer flushing" in 34K.....	74
6.3.4: Parity error exception handling and the CacheErr register .....	74
6.3.5: ErrCtl register .....	75
6.3.6: Bus error exception .....	76
6.4: Caches .....	76
6.4.1: Cacheability options .....	76
6.4.2: Uncached accelerated writes .....	77
6.4.3: The cache instruction and software cache management.....	77
6.4.3.1: Read/write synchronization and the cache instruction .....	78
6.4.4: Cache management when writing instructions - the "synci" instruction .....	78

6.4.5: Cache management and multithreaded CPUs.....	78
6.4.6: Cache initialization and tag/data registers .....	79
6.4.7: Cache aliases.....	80
6.4.8: Cache locking.....	81
6.4.9: Cache control for Multithreading CPU.....	81
6.4.10: Cache management and multithreaded CPUs.....	81
6.4.11: Cache initialization and tag/data registers .....	82
6.5: Scratchpad memory/SPRAM.....	83
6.6: The TLB and translation .....	85
6.6.1: The TLB array .....	85
6.6.2: The TLB and the MIPS® MT ASE .....	86
6.6.3: Live translation and micro-TLBs.....	86
6.6.4: Reading and writing TLB entries: Index, Random and Wired .....	86
6.6.5: Reading and writing TLB entries - staging registers .....	86
6.6.6: TLB initialization and duplicate entries.....	88
<b>Chapter 7: Kernel-mode (OS) programming .....</b>	<b>89</b>
7.1: Hazard barrier instructions .....	89
7.2: MIPS32® Architecture Release 2 - enhanced interrupt system(s) .....	90
7.2.1: Traditional MIPS interrupt signalling and priority.....	91
7.2.2: VI mode - interrupt signalling and priority.....	92
7.2.3: External Interrupt Controller (EIC) mode.....	92
7.3: Shadow registers.....	93
7.3.1: Recycling multi-threading CPU's TCs as shadow sets .....	95
7.4: Saving Power .....	95
<b>Chapter 8: 34K™ core features for debug and profiling .....</b>	<b>97</b>
8.1: EJTAG on-chip debug unit .....	97
8.1.1: Debug communications through JTAG .....	98
8.1.2: Debug mode.....	98
8.1.3: The debug unit and multi-threading .....	99
8.1.4: Single-stepping .....	99
8.1.5: The “dseg” memory decode region .....	100
8.1.6: EJTAG CP0 registers, particularly Debug.....	102
8.1.7: The DCR (debug control) memory-mapped register.....	104
8.1.8: JTAG-accessible registers .....	105
8.1.9: EJTAG breakpoint registers .....	107
8.1.10: Understanding breakpoint conditions.....	109
8.1.11: Imprecise debug breaks.....	109
8.1.12: PC Sampling with EJTAG .....	110
8.2: PDtrace™ instruction trace facility.....	110
8.2.1: 34K core-specific fields in PDtrace™ JTAG-accessible registers.....	111
8.2.2: CP0 registers for the PDtrace™ logic .....	112
8.2.3: JTAG triggers and local control through TraceIBPC/TraceDBPC.....	113
8.2.4: UserTraceData reg.....	114
8.2.5: Summary of when trace happens .....	114
8.3: CP0 Watchpoints.....	115
8.4: Performance counters .....	115
8.4.1: Reading the event table. ....	117
<b>Chapter 9: Programming the 34K™ core in user mode .....</b>	<b>121</b>
9.1: The multiplier .....	121

9.2: User-mode accessible "Hardware registers" .....	122
9.3: Prefetching data .....	122
9.4: Using "synci" when writing instructions.....	123
9.5: Tuning software for the 34K family pipeline.....	123
9.5.1: Cache delays and mitigating their effect .....	124
9.5.2: Branch delay slot.....	124
9.5.3: Branch misprediction delays .....	125
9.5.4: Data dependency delays classified.....	125
9.6: Floating point instruction timing and data dependencies.....	128
9.6.1: FPU register dependency delays .....	128
9.6.2: Delays caused by "long-latency" instructions looping in the M1 stage.....	128
9.6.3: Delays on FPU load and store instructions .....	129
9.6.4: Delays when main pipeline waits for FPU to decide not to take an exception .....	129
9.6.5: Delays when main pipeline waits for FPU to accept an instruction.....	129
9.6.6: Delays on mfc1/mtc1 instructions .....	130
9.6.7: Delays caused by dependency on FPU status register fields .....	130
9.6.8: Slower operation in "MIPS I™" compatibility mode .....	130
<b>Appendix A: References .....</b>	<b>131</b>
<b>Appendix B: Glossary .....</b>	<b>133</b>
<b>Appendix C: CP0 register summary and reference.....</b>	<b>137</b>
C.0.0.1: Power-up state of CP0 registers .....	137
C.0.0.2: A note on unused fields in CP0 registers .....	137
C.1: CP0 registers by name .....	138
C.2: CP0 registers by number .....	138
C.3: CP0 registers by function .....	141
C.4: Miscellaneous CP0 register descriptions .....	142
C.4.1: Status register .....	142
C.4.2: Exception control: Cause and BadVAddr register .....	143
C.4.3: Count and Compare .....	145
C.4.4: PRId, Configuration and EBase registers.....	145
C.4.5: Configuring interrupts - The IntCtl and SRSCtl registers.....	150
C.4.6: TLB registers .....	151
C.4.7: Cache registers .....	151
C.4.8: EJTAG unit registers .....	151
C.4.9: Watchpoint registers.....	151
C.4.10: Performance counter registers .....	152
C.4.11: Parity/ECC control.....	152
C.4.12: Registers added for Multithreading.....	152
<b>Appendix D: MIPS® Architecture quick-reference sheet(s) .....</b>	<b>153</b>
D.1: General purpose register numbers and names .....	153
D.2: Floating point information.....	153
D.2.1: Data representation .....	153
D.2.2: Setting up the FPU and the FPU control registers .....	154
<b>Appendix E: CP0 Registers of the 34K Core.....</b>	<b>159</b>
E.1: CP0 Register Summary .....	159
E.2: CP0 Register Descriptions .....	161
E.2.1: Index Register (CP0 Register 0, Select 0).....	162

E.2.2: Random Register (CP0 Register 1, Select 0) .....	162
E.2.3: EntryLo0 and EntryLo1 Registers (CP0 Registers 2 and 3, Select 0) .....	163
E.2.4: Context Register (CP0 Register 4, Select 0) .....	164
E.2.5: PageMask Register (CP0 Register 5, Select 0).....	165
E.2.6: Wired Register (CP0 Register 6, Select 0) .....	166
E.2.7: HWREna Register (CP0 Register 7, Select 0).....	167
E.2.8: BadVAddr Register (CP0 Register 8, Select 0) .....	167
E.2.9: Count Register (CP0 Register 9, Select 0) .....	168
E.2.10: EntryHi Register (CP0 Register 10, Select 0).....	168
E.2.11: Compare Register (CP0 Register 11, Select 0).....	169
E.2.12: Status Register (CP0 Register 12, Select 0) .....	170
E.2.12.1: Operating Modes.....	170
E.2.12.2: Coprocessor Accessibility .....	171
E.2.13: IntCtl Register (CP0 Register 12, Select 1) .....	175
E.2.14: SRSCtl Register (CP0 Register 12, Select 2).....	176
E.2.15: SRSSMap Register (CP0 Register 12, Select 3) .....	178
E.2.16: Cause Register (CP0 Register 13, Select 0) .....	179
E.2.17: Exception Program Counter (CP0 Register 14, Select 0).....	182
E.2.18: Processor Identification (CP0 Register 15, Select 0).....	183
E.2.19: EBase Register (CP0 Register 15, Select 1) .....	184
E.2.20: Config Register (CP0 Register 16, Select 0) .....	185
E.2.21: Config1 Register (CP0 Register 16, Select 1) .....	187
E.2.22: Config2 Register (CP0 Register 16, Select 2) .....	189
E.2.23: Config3 Register (CP0 Register 16, Select 3) .....	190
E.2.24: Config7 Register (CP0 Register 16, Select 7) .....	192
E.2.25: LLAddr Register (CP0 Register 17, Select 0) .....	193
E.2.26: WatchLo Register (CP0 Register 18, Select 0-3) .....	193
E.2.27: WatchHi Register (CP0 Register 19, Select 0-3).....	194
E.2.28: Debug Register (CP0 Register 23, Select 0) .....	195
E.2.29: Trace Control Register (CP0 Register 23, Select 1).....	199
E.2.30: Trace Control2 Register (CP0 Register 23, Select 2).....	201
E.2.31: User Trace Data Register (CP0 Register 23, Select 3) .....	203
E.2.32: TracelBPC Register (CP0 Register 23, Select 4).....	204
E.2.33: TraceDBPC Register (CP0 Register 23, Select 5) .....	204
E.2.34: Debug Exception Program Counter Register (CP0 Register 24, Select 0).....	205
E.2.35: Performance Counter Register (CP0 Register 25, select 0-3).....	206
E.2.36: ErrCtl Register (CP0 Register 26, Select 0) .....	208
E.2.37: CacheErr Register (CP0 Register 27, Select 0) .....	210
E.2.38: TagLo Register (CP0 Register 28, Select 0,2,4) .....	213
E.2.39: DataLo Register (CP0 Register 28, Select 1,3).....	214
E.2.40: DataHi Register (CP0 Register 29, Select 1).....	215
E.2.41: ErrorEPC (CP0 Register 30, Select 0).....	215
E.2.42: DeSave Register (CP0 Register 31, Select 0).....	216

**Appendix F: Revision History..... 217**

# List of Figures

Figure 2-1: Fields in the VPEControl register.....	34
Figure 2-2: Fields in the TCStatus register .....	35
Figure 2-3: Fields in the TCBind register .....	37
Figure 2-4: Fields in the MVPConf0-1 registers .....	37
Figure 2-5: Fields in the MVPControl register .....	38
Figure 2-6: Fields in the VPEConf0-1 registers.....	39
Figure 2-7: Fields in the VPEOpt register .....	40
Figure 3-1: The 34K™ core pipeline .....	41
Figure 3-2: Fields in the TCSchedule and VPESchedule registers (WRR policy manager) .....	46
Figure 3-3: Field layout in an ITC cell status view.....	47
Figure 3-4: ITC configuration information.....	48
Figure 5-1: Fields in the DSPControl Register .....	58
Figure 6-1: Fields in the CacheErr register .....	74
Figure 6-2: Fields in the ErrCtl register .....	75
Figure 6-3: Fields in the encoding of a cache instruction.....	77
Figure 6-4: Fields in the ITagLo and DTagLo Registers .....	80
Figure 6-5: Fields in DTagLo/ITagLo when used for way-select RAM.....	80
Figure 6-6: Fields in the ITagLo and DTagLo Registers .....	82
Figure 6-7: Fields in DTagLo/ITagLo when used for way-select RAM.....	82
Figure 6-8: SPRAM (scratchpad RAM) configuration information in TagLo.....	84
Figure 6-9: Fields in a 34K™ core TLB entry.....	85
Figure 6-10: Fields in the EntryHi and PageMask registers.....	86
Figure 6-11: Fields in the EntryLo0-1 registers .....	87
Figure 7-1: Fields in the IntCtl register .....	91
Figure 7-2: Fields in the SRSCtl register (shadow register set control) .....	93
Figure 7-3: Fields in the SRSSMap register .....	94
Figure 7-4: Fields in the SRSSConf0 register .....	95
Figure 8-2: Exception cause bits in the debug register .....	103
Figure 8-1: Fields in the EJTAG CP0 Debug register .....	103
Figure 8-3: Debug register - exception-pending flags .....	104
Figure 8-4: Fields in the memory-mapped DCR (debug control) register .....	104
Figure 8-5: Fields in the JTAG-accessible ImpCode register.....	105
Figure 8-6: Fields in the JTAG-accessible EJTAG_CONTROL register .....	106
Figure 8-7: Fields in the IBS/DBS (EJTAG breakpoint status) registers .....	107
Figure 8-8: Fields in the hardware breakpoint control registers (IBCn, DBCn) .....	108
Figure 8-9: Fields in the TCBCONTROLA register .....	111
Figure 8-10: Fields in the TCBCONTROLB register .....	111
Figure 8-11: Fields in the TCBCONTROLC register .....	111
Figure 8-12: Fields in the TCBCONFIG register .....	112
Figure 8-13: Fields in the TraceControl and TraceControl2 registers .....	112
Figure 8-14: Fields in the TraceIBPC/TraceDBPC registers.....	114
Figure 8-15: Fields in the PerfCtl register .....	116
Figure 9-1: Overview of the FPU pipeline .....	128
Figure C-1: All Status register fields.....	142
Figure C-2: Fields in the Cause register .....	143
Figure C-3: 34K™ processor ID (PRId) register .....	146
Figure C-4: Fields in the Config register .....	146



Figure C-5: Fields in the Config1-2 registers .....	147
Figure C-6: Fields in the Config3 register .....	148
Figure C-7: Fields in the Config7 Register .....	149
Figure C-8: Fields in the EBase register .....	149
Figure C-9: Fields in the Context register .....	151
Figure C-10: Fields in the WatchLo/WatchHi registers .....	152
Figure D-1: How floating point numbers are stored in a register .....	154
Figure D-2: Fields in the FIR register .....	155
Figure D-3: Floating point control/status register and alternate views .....	156
Figure E-1: Index Register Format .....	162
Figure E-2: Random Register Format .....	163
Figure E-3: EntryLo0, EntryLo1 Register Format .....	163
Figure E-4: Context Register Format .....	164
Figure E-5: PageMask Register Format .....	165
Figure E-6: Wired and Random Entries in the TLB .....	166
Figure E-7: Wired Register Format .....	166
Figure E-8: HWREna Register Format .....	167
Figure E-9: BadVAddr Register Format .....	168
Figure E-10: Count Register Format .....	168
Figure E-11: EntryHi Register Format .....	169
Figure E-12: Compare Register Format .....	169
Figure E-13: Status Register Format .....	171
Figure E-14: IntCtl Register Format .....	175
Figure E-15: SRSCtl Register Format .....	176
Figure E-16: SRSTMap Register Format .....	178
Figure E-17: Cause Register Format .....	179
Figure E-18: EPC Register Format .....	183
Figure E-19: PRId Register Format .....	183
Figure E-20: EBase Register Format .....	184
Figure E-21: Config Register Format — Select 0 .....	185
Figure E-22: Config1 Register Format .....	187
Figure E-23: Config2 Register Format .....	189
Figure E-24: Config3 Register Format .....	191
Figure E-25: Config7 Register Format .....	192
Figure E-26: .....	193
Figure E-27: WatchLo Register Format .....	193
Figure E-28: WatchHi Register Format .....	194
Figure E-29: Debug Register Format .....	195
Figure E-30: TraceControl Register Format .....	199
Figure E-31: TraceControl2 Register Format .....	201
Figure E-32: User Trace Data Register Format .....	203
Figure E-33: TraceIBPC Register Format .....	204
Figure E-34: TraceDBPC Register Format .....	205
Figure E-35: DEPC Register Format .....	206
Figure E-36: Performance Counter Control Register .....	207
Figure E-37: Performance Counter Count Register .....	208
Figure E-38: ErrCtl Register .....	209
Figure E-39: CacheErr Register .....	210
Figure E-40: TagLo Register Format (ErrCtl[WST]=0, ErrCtl[SPR]=0) .....	213
Figure E-41: TagLo Register Format (ErrCtl[WST]=1, ErrCtl[SPR]=0) .....	213
Figure E-42: TagLo Register Format (ErrCtl[WST]=0, ErrCtl[SPR]=1) .....	213
Figure E-43: DataLo Register Format .....	215
Figure E-44: DataHi Register Format .....	215

Figure E-45: ErrorEPC Register Format ..... 216  
Figure E-46: DeSave Register Format ..... 216

# List of Tables

Table 2.1: MTTR/MFTR - "U" and "SEL" values .....	28
Table 2.2: MT instruction summary in alphabetical order .....	30
Table 2.3: MTTR/MFTR "assembler idioms" in alphabetical order .....	31
Table 2.4: CP0 registers required by MIPS® MT ASE .....	33
Table 2.5: Thread exception codes in VPEControl[EXCPT] .....	34
Table 2.6: TC summary state as expressed in per-TC register fields .....	37
Table 3.1: Dynamic priorities for finer resolution - group priority sequences .....	46
Table 3.2: ITC cell views and what they do.....	47
Table 5.1: Mask bits for instructions accessing the DSPControl register.....	63
Table 5.2: DSP instructions in alphabetical order .....	67
Table 6.1: Basic MIPS32® architecture memory map .....	71
Table 6.2: Fixed memory mapping.....	72
Table 6.3: Cache Code Values .....	77
Table 6.4: Operations on a cache line available with the cache instruction.....	79
Table 8.1: JTAG instructions for the EJTAG unit.....	98
Table 8.2: EJTAG debug memory region map ("dseg").....	101
Table 8.3: Performance counter events.....	118
Table 9.1: Hints for "pref" instructions.....	123
Table 9.3: Lazy producer → register delays .....	126
Table 9.2: Register → eager consumer delays.....	126
Table 9.4: Long-latency FP instructions.....	129
Table C.1: CP0 registers by name .....	138
Table C.2: Cross-referenced of CP0 registers by number .....	138
Table C.3: Exception Code values in Cause[ExcCode].....	144
Table C.4: Exception entry points .....	150
Table D.1: Conventional names of registers with usage mnemonics .....	153
Table D.2: FPU (co-processor 1) control registers.....	155
Table E.1: CP0 Registers .....	159
Table E.2: CP0 Register Field Types.....	161
Table E.3: Index Register Field Descriptions .....	162
Table E.4: Random Register Field Descriptions .....	163
Table E.5: EntryLo0, EntryLo1 Register Field Descriptions .....	163
Table E.6: Cache Coherency Attributes.....	164
Table E.8: PageMask Register Field Descriptions.....	165
Table E.9: Values for the Mask Field of the PageMask Register.....	165
Table E.7: Context Register Field Descriptions .....	165
Table E.11: HWREna Register Field Descriptions.....	167
Table E.10: Wired Register Field Descriptions .....	167
Table E.12: BadVAddr Register Field Description .....	168
Table E.13: Count Register Field Description .....	168
Table E.14: EntryHi Register Field Descriptions .....	169
Table E.15: Compare Register Field Description.....	169
Table E.16: Status Register Field Descriptions.....	171
Table E.17: IntCtl Register Field Descriptions .....	175
Table E.18: SRSCtl Register Field Descriptions .....	176
Table E.19: Sources for new SRSCtl[CSS] on an Exception or Interrupt .....	178
Table E.20: SRSMaP Register Field Descriptions .....	178

Table E.21: Cause Register Field Descriptions .....	179
Table E.22: Cause Register ExcCode Field.....	181
Table E.23: EPC Register Field Description .....	183
Table E.24: PRId Register Field Descriptions.....	183
Table E.26: Config Register Field Descriptions .....	185
Table E.25: EBase Register Field Descriptions .....	185
Table E.27: Cache Coherency Attributes.....	187
Table E.28: Config1 Register Field Descriptions .....	187
Table E.29: Config2 Register Field Descriptions .....	189
Table E.30: Config3 Register Field Descriptions .....	191
Table E.31: Config7 Register Field Descriptions .....	192
Table E.33: WatchHi Register Field Descriptions .....	194
Table E.32: WatchLo Register Field Descriptions.....	194
Table E.34: Debug Register Field Descriptions .....	196
Table E.35: TraceControl Register Field Descriptions .....	199
Table E.36: TraceControl2 Register Field Descriptions .....	201
Table E.38: TraceIBPC Register Field Descriptions .....	204
Table E.37: UserTraceData Register Field Descriptions .....	204
Table E.39: TraceDBPC Register Field Descriptions.....	205
Table E.40: BreakPoint Control Modes: IBPC and DBP .....	205
Table E.41: DEPC Register Formats .....	206
Table E.42: Performance Counter Register Selects .....	206
Table E.43: Performance Counter Control Register Field Descriptions.....	207
Table E.44: Performance Counter Count Register Field Descriptions.....	208
Table E.45: ErrCtl Register Field Descriptions.....	209
Table E.46: CacheErr Register Field Descriptions.....	211
Table E.47: TagLo Register Field Descriptions.....	213
Table E.48: DataLo Register Field Description .....	215
Table E.49: DataHi Register Field Description.....	215
Table E.50: ErrorEPC Register Field Description .....	216
Table E.51: DeSave Register Field Description.....	216

# Introduction

In this chapter you'll find:

- [Section 1.1, "Chapter summary"](#): what's in the chapters (hot links if you're reading online).
- [Section 1.2, "Typographical conventions"](#): a manual like this is made easier to read (though perhaps not made more beautiful) if we use typographical conventions so you can recognize machine registers, instructions and so on. Here's what they look like.
- [Section 1.4, "Key features of the 34K™ core"](#): a quick guide to the important features of the 34K core.
- [Section 1.5, "Specification summary"](#): a terse summary of facts and figures.

## 1.1 Chapter summary

- [Chapter 2, "The MIPS® MT ASE - Multithreading the RISC way" on page 17](#): about the MIPS Multi-Threading instruction set extension ("ASE").
- [Chapter 3, "How the 34K™ core implements multi-threading" on page 41](#): implementation options and more details.
- [Chapter 4, "Initializing the 34K™ core - Multi-Threaded bootstrap issues" on page 51](#): setting up the 34K core's multi-threading system.
- [Chapter 5, "The MIPS32® DSP ASE" on page 57](#): the instruction set extension for faster media algorithms.
- [Chapter 6, "Memory map, caching, reads and writes and translation" on page 71](#): all about memory accesses and translation.
- [Chapter 7, "Kernel-mode \(OS\) programming" on page 89](#): use of "hazard barriers", the advanced interrupt system, shadow registers and power management.
- [Chapter 8, "34K™ core features for debug and profiling" on page 97](#): EJTAG debug unit, watchpoints and performance counters.
- [Chapter 9, "Programming the 34K™ core in user mode" on page 121](#): on tuning code specifically for the 34K core family.

Appendices:

- [Appendix A, "References" on page 131](#): further reading.
- [Appendix B, "Glossary" on page 133](#): a glossary of terms which may be unfamiliar (particularly relating to multi-threading).

## Introduction

- [Appendix C, “CP0 register summary and reference” on page 137](#): functionally orientated index and guide to the 34K core’s “co-processor zero” registers and fields.
- [Appendix D, “MIPS® Architecture quick-reference sheet\(s\)” on page 153](#): handy guide to easily-forgotten data on MIPS.
- [Appendix E, “CP0 Registers of the 34K Core” on page 159](#): reference-manual long-format tables of the CP0 registers and fields (if you are printing this manual, you may like to bind this section separately).
- [Appendix F, “Revision History” on page 217](#): for this document.

## 1.2 Typographical conventions

CPU register names are in *oblique monospace*. *Co-processor zero* (CP0) registers fields are shown after the register name in brackets, so the interrupt enable bit in the *Status* register appears as *Status[IE]*. CP0 register numbers are denoted by *n.s*, where “n” is the register number (between 0-31) and “s” is the “select” field (0-7). If the select field is omitted, it’s zero. A select field of “x” denotes all eight potential select numbers.

The acronym *CP0* in the paragraph above is a word defined in [Chapter B, “Glossary” on page 133](#) and shows up in *italics* - but if you’re reading on-line it also shows up as *blue*, showing that it’s a link which you can click to get to the definition.

References to other manuals are collected together in [Appendix A, “References” on page 131](#) and look like this [\[MIPS32\]](#).

Instruction mnemonics and assembler code fragments are set in **bold monospace**, core interface signal names in *small italics*, and C or other programming language constructs in monospace.

To use register and field names in your program, you’ll need a C header file or something similar. It’s probably better and easier not to write your own: see [\[m32c0.h\]](#) and [\[mt.h\]](#).

## 1.3 Finding information in this manual

If you’re reading this manual on-screen, text shown in blue is a hot-link; click on the text to go to the section, figure or table referenced. The chapter index and lists of tables and figures at the start of the book is click-through too.

All the special *Co-processor zero* (CP0) registers are listed in [Appendix C, “CP0 register summary and reference” on page 137](#). That appendix has the registers listed by name, by number and by function. The by-number table has hot-links to other sections where each is mentioned - and for those reading on paper, all those links have page numbers. There’s an alternative description of all the CP0 registers and fields in the more formal style used in other MIPS Technologies manuals in [Appendix E, “CP0 Registers of the 34K Core” on page 159](#).

## 1.4 Key features of the 34K™ core

The 34K core is a 32-bit MIPS32 CPU with two novel instruction set extensions:

- *The MIPS® MT ASE*: The multithreading ASE (“application-specific extension” to the MIPS architecture). It’s a modest addition to the instruction set, but a profound change to the CPU, which can now run multiple threads

concurrently. The set of software-visible resources devoted to one thread are known as a *TC*. The 34K core allows for two multithreading models which are very different for software:

- Multiple *Virtual Processing Elements (VPEs)* in a CPU: each “VPE” has at least one TC together with its own copies of everything required to make it just like an independent MIPS CPU. Your 2-VPE (or more) system seems to software just like a 2-CPU “SMP” multiprocessor: indeed, it can run SMP software - software which knows nothing about MIPS MT - without requiring any CPU-related changes.
- Multiple concurrent threads running within one VPE, usable by software which knows about MIPS MT. These multiple threads are relatively cheap, because they’re equipped only with the resources necessary to run user-level programs (but they share a lot of OS-controlled resources.)

Much of this manual won’t make any sense until you get your head round multithreading, so unless you’re thoroughly familiar with it already you should acquaint yourself with [Chapter 2, “The MIPS® MT ASE - Multithreading the RISC way”](#) on page 17.

- *DSP ASE*: this is a lot of new computational instructions with a fixed-point math unit crafted to speed up popular signal-processing algorithms, which form a large part of the computational load for voice and imaging applications. Some of these functions are “SIMD” - they might, for example, do two math operations at once on two 16-bit values packed into one 32-bit register.

There’s a guide to the DSP ASE in [Chapter 5, “The MIPS32® DSP ASE”](#) on page 57 and the formal specification is [\[MIPSDSP\]](#).

## 1.5 Specification summary

The 34K core is provided as a synthesizable package, and customers have considerable freedom to customize it. But all 34K cores share these:

- *CPU architecture*: compliant to Release 2 of the MIPS32 Architecture [\[MIPS32\]](#).
- *Multi-threading*: as defined by the Multithreading extension to the MIPS32 architecture as specified by [\[MIPSMT\]](#).

The 34K core can be synthesized to be able to run five concurrent threads (5 *TCs*) in up to two “virtual processors” (2 *VPEs*).

It may be equipped with a bank of *Inter-Thread Communication storage (ITC)* locations, following the recommendations of [\[MIPSMT\]](#).

- *DSP-orientated instruction set*: it implements the DSP extension to the MIPS32 architecture, see [\[MIPSDSP\]](#).
- *MIPS16e™*: the 16-bit instruction set option for compact code, see [\[MIPS16e\]](#).
- *9-stage pipeline*<sup>1</sup>: a sophisticated branch prediction unit keeps the CPU efficient, even when it’s only running one thread.
- *Separate I- and D-caches*: 4-way set associative. The SoC designer may choose from 16, 32 or 64Kbytes size for each cache (and can even omit either cache). Parity checking in the cache is optional.

---

1. 11-stage for MIPS16e instructions

## Introduction

Caches are non-blocking, and both allow for hit-under-miss and miss-under-miss - the I-cache uses that to allow a cache-hitting thread to continue even though an I-cache refill is pending for some other thread.

The D-cache is write-back (memory regions may also be configured a write-through and a special "uncached accelerated" write mode). You can lock data into the caches.

- *OCP system interface*: industry-standard interconnect.

### SoC Builder's Optional features

Some features are provided only at the option of the SoC integrator, and may depend on separate licensed material from MIPS Technologies:

- *CorExtend™ user-defined instructions*: the 34K Pro Series™ core family allows you to add custom instructions as described in [\[CorExtend\]](#).
- *Floating point unit*: fitted to 34Kf™ cores, with 32 full 64-bit floating point registers.
- *Fixed mapping MMU*: reduces core size when a TLB is not required.
- *Instruction- or data-side "scratchpad" memory*: each can be up to 1Mbyte of high-performance on-chip memory, which can be dual-ported to the OCP interface for "push" I/O architectures.
- *EJTAG debug unit*: on-chip debug resources, summarized in [Section 8.1, "EJTAG on-chip debug unit"](#).
- *Power-management options*: summed up in [Section 7.4, "Saving Power"](#) below.
- *OCP L2 extensions*: to allow front-side L2 cache.

Refer to [\[34K\\_INT\]](#) for more details about the options.

## 1.6 Pipeline and implementation

In documents about MIPS Technologies other cores you'd have found a section here with the basic pipeline. With the 34K core that is hard to describe without knowing something about multi-threading (particularly) so we've moved it to [Section 3.1, "The 34K™ core pipeline and multithreading"](#) below.



## The MIPS® MT ASE - Multithreading the RISC way

We use “MT” for “multi-threading”. So what does a MIPS architecture CPU do to run multiple threads concurrently? That question is one about “architecture” - the corresponding “how does the 34K core run multiple threads?” question is about implementation, and is answered below in [Chapter 3, “How the 34K™ core implements multi-threading” on page 41](#).

In this chapter:

- [Section 2.1, “What’s a thread and its context?”](#): basic definitions.
- [Section 2.2, “Why multi-threading?”](#): motivation.
- [Section 2.3, “Different kinds of multi-threading: TCs and VPEs”](#): we offer two levels of multi-threading in one CPU.
- [Section 2.4, “When can’t threads run?”](#): and what they’re doing when stopped.
- [Section 2.5, “Thread-scheduling decisions and the policy manager”](#): what happens and what influence can you have.
- [Section 2.6, “Multithreading, exceptions and interrupts”](#): interrupts and other exceptions in the MIPS MT CPU.
- [Section 2.7, “Multithreading, non-blocking loads and stores, and gating storage”](#)
- [Section 2.8, “MIPS® Multithreading ASE - new instructions”](#)
- [Section 2.9, “Multithreading ASE - CP0 \(privileged\) registers”](#): understanding multi-threading in fine detail.

### Why multi-threading takes a lot of thinking about

Any form of concurrency makes your head hurt. Our brains are doubtless extremely parallel: we can talk on a cell-phone and drive with only a 50% increase in our chance of crashing. But our ability to reason correctly is distinctly sequential, and so far we have not bred a race of super-kids who can write explicitly parallel software.

Multi-tasking software has been successfully understood by dividing it into sequential chunks (“threads”, though a more precise definition follows) which communicate and synchronize with each other only in carefully controlled ways. You can then unleash a flock of threads and allow them to evolve separately. Programmers find it almost impossible to keep track of what every thread is doing at any one time - but with simple-enough rules about the interactions, the system will still work.

The multithreading CPU pushes thread concurrency down to the hardware level, so you should expect to find it somewhat mystifying from time to time. To really understand multi-threading and the 34K core you need to be able to switch between a software-orientated threads-eye-view (where threads are internally sequenced and other threads are

happening somewhere else) and a hardware engineers CPUs-eye-view (where everything happens in sequence along the pipeline). This is difficult, but we hope not impossible. This chapter takes the “thread” viewpoint, and the next chapter stays closer to the hardware.

## 2.1 What’s a thread and its context?

There are a couple of critical phrases and acronyms which it’s useful to define carefully before we start:

- *Thread*: a set of computer instructions read and activated in their programmed order.

Operating systems most often use the word “thread” specifically for application-software visible threads scheduled by the OS. But our wider definition means that any piece of software must have at least one thread. Skid buffer

By this definition something like an interrupt handler (which is not reached as a result of normal program flow) counts as a thread in its own right. This more general definition of “thread” seems to be a more logical starting point for describing multi-threading hardware.

- *Thread context*: you might want to consider the complete state of a running thread, enough so you could restart it successfully. But for our purposes we’re particularly interested in the part of the state which gets stored inside the CPU - what [MIPSM] calls the “thread context”. The thread context always (of course) includes the *Program Counter (PC)* and the general-purpose registers. There are some good justifications for narrowing our focus down to the state held in the CPU:
  1. We don’t need to encompass the thread’s data stored in memory, because we know how to share memory already (for OS-defined threads, for example);
  2. We don’t include state which is inherently inaccessible to this particular instruction stream - so kernel-only readable CP0 registers are invisible to a user-privilege thread;
  3. We don’t include state which is logically unnecessary, and just kept for efficiency - for example, cache contents, which generally make no difference to the underlying memory image.

With this definition, what is included in the thread context varies according to what sort of software is running. For a Linux interrupt handler on a conventional MIPS architecture CPU the CP0 registers are part of the thread state, but for a Linux application thread they’re not visible.

You could have found the definitions of *Thread* and *Thread context* in [Appendix B, “Glossary” on page 133](#) below. Any word or phrase in blue (or slightly faint in real black-and-white print) is probably explained. If you’re reading online and it’s blue, it will link to its definition: try it.

## 2.2 Why multi-threading?

Traditionally, a CPU only held one thread’s context (one PC, one set of registers). Operating systems providing multiple threads held all the state for the non-running threads in OS-specific data structures.

But MIPS MT CPUs are equipped with more than one PC and register set so they can hold more than one thread’s context.

There’s more than one reason why you might want to build a multithreading CPU. For MIPS MT the main motivation is to build a CPU which can continue to do useful work when some computation is held up for a period of a few to

some hundreds of CPU cycles - typical of cache misses and some other interactions in embedded systems. Such a hold-up is too short to allow an OS to borrow the CPU to do something else (the OS thread-switch overhead is itself probably 100 cycles or more). But in many workloads such hold-ups are frequent enough that the CPU spends half its time waiting for data.

A multithreading CPU can keep other threads making progress when one thread is held up. If (as is commonly the case these days) the real workload is already split into multiple threads, that can turn into extra application performance without modifying application code.

The extra thread state storage (mostly the register file) only represents a fraction of the gate count of a CPU, so this extra performance has cost only a small increment in area and complexity. That's why in 2005 everyone wants to do multithreading.

## 2.3 Different kinds of multi-threading: TCs and VPEs

In some ways the simplest thing to do is to replicate every software-visible piece of CPU state. Then your multithreading CPU will look pretty much like two CPUs which happen to share memory, creating a "virtual multiprocessor" (*VSMP*). That's what Intel's newer multithreading x86 processors do; you can drop a Linux kernel designed for a two-way multiprocessor onto such a CPU and it just works. It's an easy way to get a software market for a new technology.

But performance-critical embedded applications are those where the multithreading is an explicit part of the system design - we'll call it "explicit multithreading" or *EMT*. EMT is new, so we don't need to offer backward compatibility. An EMT application does not need the whole CPU replicated; it can manage with what is visible to user-level programs - the PC, GPRs and a little more.

The original and ingenious trick in the MIPS MT architecture is that you have a choice of either model, and can even do both in the same CPU at the same time. So a MIPS MT CPU has multiple *TCs* (the acronym started out as *Thread context*), but also provides for more than one *VPE* ("VPE" started out as a *Virtual Processing Element*.) A TC provides the minimum required to do explicit multithreading, while one or more TCs with their own VPE really look like an independent CPU, enough to provide a congenial home for software which doesn't really want to know about MIPS MT - perhaps even a non-MT-aware legacy operating system.

### 2.3.1 How an MT CPU's hardware uses TCs and VPEs

Each instruction being run by an MT CPU has a TC number. Whenever the instruction accesses some state - reads or writes a general-purpose register, for example - it uses its TC number to extend the register-number field which is already defined inside the instruction. An instruction sees a different set of registers depending on the TC number: it's very simple, and it just works.

It's not quite that simple on a MIPS architecture CPU, because of the TC/VPE trick mentioned above. So this instruction might be for TC #5 (it uses general purpose registers from the fifth bank) but VPE #1 (it gets most of its CPU registers from the first bank). Again, this should just work. What's more complicated, of course, is to get those CPU resources working which can't simply be reduced to registers. But that's not architecture, it's implementation, and described in [Chapter 3, "How the 34K™ core implements multi-threading" on page 41](#) below.

### 2.3.2 CPU resources and registers shared between all threads

Many of the CPU's resources are not replicated for MIPS MT, just used by whichever TC is identified by the instruction accessing the resource. They include:

- *Caches*: the cache's contents are just like memory (only faster) and unproblematic. On a CISC CPU the cache is usually completely invisible to running software, and there's no issue at all about multiple threads - but MIPS architecture CPUs generally need the OS to intervene in the caches at some points.

The MIPS MT ASE requires that the writeback and invalidate **cache** instructions used by real OS' when running are multi-threading safe. Cache manipulations may be independently mixed by two VPEs<sup>1</sup> without immediate harm; even if one VPE invalidates a cache entry from right under the feet of another one, everything should keep working - the consuming VPE will either get the old copy (which it was happy with) or cache-miss and pull in a new one (which should be just the same data).

However, arbitrary re-initialization of a cache already in use by another VPE will not be safe; writeback data could be lost. Programs running on separate VPEs would probably be well-advised to get cache initialization done by a thread running alone before other VPEs are enabled.

With a multithreading workload, cache performance could suffer; multiple threads will probably produce a larger and more diverse "working set" of active memory regions. However, a cache works well (or not) when optimizing repeated accesses over spans of code executing hundreds of thousands to millions of instructions. During that time which even a single-threaded workload will climb all over application and OS space. The 34K core's caches are already 4-way set associative, which should be enough to minimize misses caused by overlapping hot-spots of several concurrent threads. Our measurements to date back that up.

- *Main pipeline*: each of the 34K core's main pipeline stages just serve the TC associated with the current instruction. No problem.
- *The TLB (sometimes)*: the MIPS MT ASE allows the TLB entries to be shared between all VPEs, or partitioned between VPEs. The 34K core can be configured to do either (to share the TLB, set *MVPControl[STLB]* to 1.)

If the TLB is not shared, it is partitioned by hardware so each VPE sees its own independent array of entries.

When the TLB is shared, there's a problem of managing concurrent access by the two VPEs. It's up to OS software to control concurrent access by OS maintenance routines. But that still leaves the risk that one VPE's maintenance software will collide with another VPE's TLB refill exception handler: see [Section 4.2.2, "Sharing and not sharing the TLB"](#) for how that's avoided.

- *Basic configuration registers*: in a highly adaptable design like the 34K core the initialization software needs to know the full resource complement of the CPU, or it can't know how to share it between the VPEs.

The registers *MVPControl* and *MVPConf0-1* allow software to see what resources are provided CPU-wide, and these registers are not replicated per-VPE.

- *Performance counters*: since these are infrequently used, but it's valuable to have as many as possible available, the four registers are shared between both VPEs.

This is more implementation than architecture, but some software-invisible resources are also shared. Notably, the 34K core's "branch history table" (BHT) in the instruction fetch unit is shared. That seems quite wrong: the branch histories of different threads are certainly likely to be different. But the BHT was only statistically correct anyway; the branch history is only recorded in entries indexed by some modest number of low virtual address bits. Even in a conventional single-thread CPU, different branches could map onto the same entry and cause confusion (and thus

---

1. The CP0 registers used with the **cache** instruction are only replicated per-VPE, so EMT code must take care to avoid re-entry into cache management functions by other threads.

lower the prediction accuracy) - but there are enough different entries that this relatively rarely happens. Having multiple threads doesn't really make it much worse, and the BHT should continue to perform well in typical applications.

### 2.3.3 CPU resources and registers replicated per-TC

Some state needs to be independently kept for each TC, including:

- *Program counter and general purpose (integer) registers*: the TC's program counter can be seen and adjusted (when the TC is halted, otherwise it's a moving target) in *TCRestart*. The architecture does not define what you'll get if you read your own *TCRestart*; probably some "historical value".

Each TC, of course, has its own set of 32 general purpose registers. It also needs its own copies of the accumulator registers in the multiply-divide unit (*hi/lo*), and the extra accumulator registers and control register provided as part of the DSP ASE described in Chapter 5, "The MIPS32@ DSP ASE" on page 57.

- *Privilege state*: some TCs (sharing a VPE) may be in the kernel while others are running user-mode software. So each TC has its own copy of the user-mode/kernel-mode flags *Status[KSU]*. *TCStatus[TKSU]* provides a convenient per-TC view of the same flags. Each TC gets a copy of the *TCContext* register too: it has no hardware significance, but provides a useful scratch register for the OS to keep some key thread identifier.
- *Address space*: we don't want to insist that all TCs which share a VPE must execute in the same address space. Different address spaces in MIPS architecture CPUs are managed by only returning TLB translations for virtual addresses when they're presented together with the right "ASID" value, an arbitrary 8-bit token held in *EntryHi[ASID]* while the system runs.

So each TC also has its own copy of the *EntryHi[ASID]* field - the same field is accessible as *TCStatus[TASID]*.

- *Access to co-processors*: the 34K core's FPU - when fitted - is built with just one set of registers. That makes sense because the registers in the floating point unit already occupy a lot of logic space, and the 1-register-set FPU design is identical to that used in the 24K™ core family. But it means that the FPU can't be used by multiple concurrent threads.

Some other co-processors might have one set of data registers per TC, supporting arbitrary multi-threading.

In the MIPS architecture you can't use any co-processor unless you first turn on the corresponding *Status[CUx]* bit in the status register. MIPS MT uses that to provide a mechanism to share the co-processors, detailed in the notes to Figure 2-2 below. As part of that mechanism the *Status[CU3-1]* bits are also visible at *TCStatus[TCU3-1]*.

- *Which VPE we're using*: a TC must know which VPE it belongs to, or it can't get at the right copy of the per-VPE registers. The VPE affiliation is readable and writable in *TCBind[CurVPE]*. (Each VPE also has a distinct number readable at *EBase[CPUNum]*, to allow seamless use of multi-CPU software on multiple VPEs.)
- *TC halted*: think of this as "TC anesthetized" - it stops the TC from wriggling around when under surgery, or even just close inspection. It occupies its own 1-bit register *TCHalt* so it can be set and cleared atomically.

While this is set the TC is frozen: won't run, can't be picked by **fork**. The architecture abhors the idea of a halted thread being half-way through a synchronization access, and any pending load/store to *Gating Storage* will be rolled back when this bit is set. From a hardware point of view the gating storage access is aborted; but unless you do something special to stop it the access will be quietly retried once the OS is finished with its maintenance and clears *TCHalt*.

- *TC interrupt-exempt*: set *TCStatus[IXMT]* to mean this TC will never be picked to handle an interrupt exception (even if that means the interrupt is completely ignored).

- *Per-TC flags*: there are also bits to control the ability of **fork** to seize a “free” TC and make it run a new thread, and for other purposes. See the description of **fork** in Section 2.8, "MIPS® Multithreading ASE - new instructions" and the notes on Figure 2-2.
- *Debug state*: the single-step bit *Debug[SSt]* is replicated per-TC, for fine debugger control. The debugger is also given a control bit *Debug[OffLine]* which it can use to prevent TCs other than the one under debug from springing into life during single-step or when running a thread to the next breakpoint.

### 2.3.4 CPU resources and registers replicated per-VPE

We want a TC running alone in a VPE to be a MIPS32-compliant processor in its own right, so each VPE replicates all the CP0 registers required by release 2 of the MIPS32 specification (a few read-only registers are in fact shared between VPEs on the same CPU, but they're read-only, so who's to know?)

So what is replicated?

- *State related to exceptions*: MIPS architecture experts will recall that you enter exception mode by taking an exception, and remain in it until you either return with an **eret** or (more common in a complicated OS) you carefully clean up exception-dependent information and then manually clear *Status[EXL]*.

The MIPS MT architects determined that only one TC from a VPE is allowed to be in exception mode at any one time - when one TC takes the exception, its VPE siblings are suspended until the first TC clears *Status[EXL]*. To do otherwise would require a lot of extra replicated state, and would lead to some nasty concurrency hazards.

- *Interrupt system and interrupts*: interrupt signals to the chip are wired to VPEs separately (a reasonable strategy may be to wire all the VPEs in parallel to the same inputs, but that's an SoC designer's decision).

The interrupt management fields in the *Cause* and *Status* registers are all per-VPE.

- *Cache management registers*: all the cache operation staging registers are per-VPE. In fact, most of the CP0 registers are per-VPE.
- *The TLB (sometimes)*: on the 34K core the TLB may either be shared, or partitioned invisibly so that two VPEs each think they have their own dedicated chunk of the TLB<sup>1</sup>.
- *The EJTAG debug unit*: the physical unit may or may not be replicated, but the registers in its CP0 software interface (*DEPC*, *DESAVE* and *Debug*) are replicated per-VPE.

In debug mode all TCs other than the one running the debugger are suspended, regardless of VPE affiliation. Moreover, the TC in debug mode continues to run even if it is otherwise marked as halted, not-allocated etc. More details in Section 8.1.2, "Debug mode".

## 2.4 When can't threads run?

A CPU can be compliant to the MIPS MT ASE without being committed to any particular thread-scheduling algorithm - the decision as to which thread's instruction to pick next is implementation-dependent. But that level of abstraction is difficult, so let's make some working assumptions - which will, happily, turn out to be correct for the 34K core.

---

1. The amount of the TLB awarded to each VPE is configurable when your core is synthesized. Ask your hardware engineer.

Some implementations permit customizable hardware outside of the core to influence the CPU to favor one TC over another when deciding what instruction to run next; see [Section 2.5, "Thread-scheduling decisions and the policy manager"](#) below.

But before worrying about that, let's look at something simpler. A practical CPU might run instructions in turn ("round-robin") from each live thread. But what about that weasel word "live"? When can a thread *not* make progress? Well, it can be:

- *Waiting for memory data*: most often, to resolve a cache miss (for of the order of 50 cycles) - making use of this idle time is the first motivation for contemporary multithreading.

Or this might also be an uncached read of some device-register data (of the order of 100-500 cycles) - particularly relevant to embedded applications.

- *Blocked on read/write "gating storage"*: we envisage that multithreading applications are likely to use special memory locations where the wait-for-transfer is used as a deliberate way of matching the speed of the software to the arrival of data either from other threads, or some direct hardware source/sink. Waits of this kind may extend for thousands of cycles. So the MIPS MT *ASE* describes how some memory locations are accessed according to special rules which make them *Gating Storage*, and describes a particular application of gating storage to optional *ITC* locations. See [Section 3.3, "Inter-thread communication storage \(ITC\)"](#) for the facility provided by the 34K core.
- *Blocked on an "interrupt-like" external signal*: a thread which waits for a particular hardware signal is an obvious multithreading analogue of an interrupt handler, and likely to be useful. You'll see how the MIPS MT **yield** instruction can be used for that purpose.
- *Halted - closed for maintenance*: there are bound to be things the OS wants to do with TCs which can't be done while it's live, and each TC comes with a "Halt" button in the *TC Halt* register.
- *Not "allocated"*: the MT system includes the **fork** instruction, which provides a very lightweight way of starting a new thread - potentially, it's even usable from user-mode in a protected OS. An OS obviously can't simply relinquish control of thread scheduling, but it can arrange to provide a pool of "free" threads which **fork** can use - they're a bit like taxis waiting at a taxi-rank for customers. The TCs "at the rank" are prevented from running code by having their *TCStatus[A]* ("allocated") bit clear. If a system doesn't use **fork**, then it must take care to set the allocated bit explicitly on any TC which is to run.
- *Affiliated to an unactivated VPE*: that is, one with *VPEConf0[VPA]* zero.
- *Asleep after executing a wait instruction*: in which case it won't awake until its VPE gets an interrupt (it doesn't matter which TC runs the interrupt code, all TCs are woken from their sleep).
- *Suspended - temporarily inhibited to avoid some concurrency problem*: for example we'll see that a VPE becomes "single-threaded" while it is handling exceptions, so that implicitly suspends all the VPE's other TCs. OS software can achieve a similar effect using instructions such as **dmT** (stop all other threads with the same VPE affiliation) and **dvpe** (stop all other threads, even in different VPEs).
- *"Offlined" by a debugger*: using *Debug[OffLine]*, typically so the debugger can isolate another thread for test.

In this manual we'll try to consistently use the word *stopped* for a thread subject to any of the conditions above - and by analogy, we'll use the same adjective to describe the TC which is executing the thread. The opposite of "stopped" is *live*.

We'll distinguish a stopped thread as either:

## The MIPS® MT ASE - Multithreading the RISC way

- *Stalled*: waiting for a condition which could be experienced by a program on a single-threaded CPU - that includes waiting for data from a cache miss or an uncached read, OR:
- *Blocked*: waiting for something other than the above. That's some deliberate multithreading synchronization by **yield**, a gating storage read, or explicitly stopped as a result of software activity.

The blocked state is new with MIPS MT. The nearest thing that a thread on a non-MT MIPS CPU can come to "blocked" is when the CPU is asleep after executing **wait**.

For blocked threads we'll use *halted*, *suspended* and *asleep* in the specific senses above. The use of these terms is compatible with the formal specification [MIPSMT], though that uses *running* to instead of live. In the formal specification "running" means either live or waiting for a normal read/write.

Regardless of why a thread is stopped:

- *The CPU*: will be interested in issuing instructions from some other live thread. In a simple pipelined CPU, that may involve discarding some instructions from the stopped thread, if they've already entered the main pipeline.
- *The OS*: may be interested in taking control when a thread is blocked for a long time - the TC could be in principle given another thread which might be able to make more progress. The OS overhead in changing the TC to another thread - really the same job as a thread-switch on a conventional CPU - is likely to be more than 100 instructions so the OS should only do this when the thread is likely to remain stopped for many hundreds of cycles.

But it's important that the OS has the power to take a blocked thread and detach it from its TC cleanly, so it can be restarted. That motivates some of the key features of the architecture, including the details of *Gating Storage*, see Section 2.7.1, "Gating storage".

## 2.5 Thread-scheduling decisions and the policy manager

The MIPS MT architecture is agnostic about thread scheduling. The immediate choice of which thread to run next is made inside the core; in the absence of any directions to the contrary, this choice is required to be fair to TCs in the long run.

However, in MIPS Technologies cores we envisage a rather dumb in-core scheduler given long-term hints by a *Policy Manager (PM)* which, living outside the core, may be customized for specific applications.

In particular the *TCSchedule* and *VPESchedule* registers (if implemented at all) will typically be inside the policy manager block; so what they do is strictly implementation-dependent.

The way the in-core scheduler in the 34K core works is described in Section 3.2.1, "The Dispatch Scheduler", and the choice of policy managers available from MIPS Technologies is in Section 3.2.3, "Policy managers available for the 34K™ core family".

## 2.6 Multithreading, exceptions and interrupts

An exception in a single-threaded MIPS architecture CPU is usually quite disruptive in the pipeline, and is commonly implemented by discarding a lot of execution state (pipelines get flushed and instructions discarded). An exception on a MIPS MT machine happens within a thread context - and other threads (at least those on separate VPEs) expect to continue undisturbed. So you'd expect there to be some difficulties when we redefine exceptions on a multithreading machine.



There are two types of exceptions:

- Interrupts are “asynchronous” - they happen for reasons unconnected with any particular instruction and are discussed in [Section 2.6.1, "Multithreading and interrupts"](#) below.
- Synchronous exceptions, associated with a particular instruction. That’s what we’ll look at first.

Bear in mind that an OS is a program (a set of threads, in fact). It’s not characterized by the TC which happens to execute some part of it. The OS’ exception handlers are each separate threads in their own right, in the meaning given by our definition of *Thread*.

Synchronous exception handlers are run by the TC whose instruction caused the exception. The TC immediately ceases work on its thread and starts fetching instructions from the appropriate exception handler.

The MIPS MT ASE requires that once a TC enters exception state, all the other TCs within the same VPE are suspended. None of the other TC’s instructions may be executed until the VPE’s *Status[EXL]* bit is cleared<sup>1</sup> by the exception handler. The exception handler (a new thread, remember) runs with kernel privileges and has access to all the defined CP0 registers. Because only one TC can be in exception state, the exception-related CP0 registers need only be replicated per-VPE.

In your MIPS MT system an exception not only causes a hiccup to the thread which takes it, but also suspends unrelated threads in the same VPE. If your application needs to maximize concurrency, you should consider minimizing exceptions - you may be able to use a thread blocked on an ITC access or **yield** condition instead. And, of course, arrange that exception handlers (as soon as they can) save the state necessary that they can drop back out of exception mode.

### 2.6.1 Multithreading and interrupts

In the MIPS architecture interrupt management is by CP0 registers (in particular, *Cause* and *Status*). Those registers are replicated per-VPE, not per-TC; so interrupt masking and steering is managed per-VPE. Even interrupt "wiring" into the core is per-VPE.

Each interrupt input may be connected to just one VPE or to all of them: ask your hardware engineer. In some systems you may be able to redirect interrupts (outside the CPU) under software control. If you connect and unmask an interrupt on multiple VPEs, any number of them may take the interrupt exception - you probably don’t want that to happen, so either don’t connect or don’t enable some of them...

The interrupt exception may be taken by any available TC associated with the VPE.

The MIPS architecture already provides multiple ways to refuse an interrupt exception: an interrupt to any thread from this VPE can be prevented by exception mode, a global interrupt-enable flag which may be zero, and by per-interrupt mask bits: that is by *Status[EXL]*, *Status[IE]* and *Status[IM]*<sup>2</sup>. The MIPS MT architecture adds yet another reason not to take an interrupt. You can now set a new per-TC CP0 register field *TCStatus[IXMT]* to make the TC *Interrupt exempt*. That will prevent the particular TC from being used for an interrupt exception. It’s most obvious use is to permit some TC to run a thread which benefits from living in an interrupt-free universe.

- 
1. That may seem somewhat restrictive, but is necessary: critical exception handling state in the CP0 registers is not replicated per-TC, only per-VPE. And it’s not so bad as it looks, because it’s already good practice to minimize the amount of code which runs with *Status[EXL]* set.
  2. This list is not comprehensive.

## 2.7 Multithreading, non-blocking loads and stores, and gating storage

Most modern MIPS architecture cores implement non-blocking loads: that is, the core does not simply stop and wait for the load data to arrive. Instead, the register target of the load is marked and computation continues. If the data arrives before the program tries to use it, the data is sent directly to the register. But if some other instruction wants to read the register before the data arrives, the "consuming" instruction waits.

That means that a thread in a MIPS MT machine which does a "slow" load stops on the consuming instruction. When that happens the TC is still holding resources (e.g. the "fill buffer" in the CPU's bus interface unit which remembers the load, waits for the data, and associates it with the register).

If you are using long-delayed loads as a means of synchronizing your application, non-blocking loads are unwelcome: it would be preferable for the thread to stop on the load itself. So we provide a way to do that: memory locations used for synchronization can be mapped as *gating storage*.

### 2.7.1 Gating storage

The MIPS MT ASE provides for a kind of storage location whose behavior is adapted to loads which might be quite long-delayed, and which you may want to use for intentional thread synchronization. Such a location is called *Gating Storage*. A thread loading from a memory region marked as "gating" will block on the load itself. This is not the standard way of doing things: a thread which reads from a normal location which is slow to respond would run on until it attempted to use the data (that's a "non-blocking load").

It turns out to be useful to generalize this to writes as well as reads: even stores to gating storage locations block until the core gets an indication that everything went OK.

If a thread is blocked on a gated storage access and the OS decides that one of its valuable TCs has been hanging around too long, then the OS can take action. If the OS writes a 1 to the TC's *TC Halt* register any gating storage access will be aborted, with the *TC Restart* address set to re-execute the load/store. Once the TC is safely halted, the OS can decide to use the TC for something else. When the thread is eventually scheduled again, the load instruction will be re-executed. Meanwhile the CPU hardware can forget about it.

The core interface provided for gating storage locations also permits external logic to abort an uncompleted load or store. Perhaps it's better to describe this as "complete the operation with an exceptional condition". The thread doing the access gets an exception, with the restart address set so the load/store will be retried after the exception. The gating storage exception is synchronous, and you're guaranteed that the restart location captured in *EPC* will point to the load/store (or a preceding branch, if the load/store is in a branch delay slot). The exception can only happen if the thread is still waiting for the load/store, and the thread isn't otherwise prevented from running.

If required an OS can take control of all GS load/stores; set *VPEControl[GS!]* and all GS accesses trigger an exception.

Out on the gating storage interface, no external party can see whether a TC is waiting or not. All GS transactions involve delivering something which waits around until the other side responds (some software books call this kind of synchronization a *rendezvous*).

Gated storage provides the opportunity to provide *ITC* locations - a form of what some of you may have read about before as "full/empty storage". The ITC implementation which is optional in the 34K core is described in [Section 3.3, "Inter-thread communication storage \(ITC\)"](#) below.

## 2.8 MIPS® Multithreading ASE - new instructions

There are very few extra instructions:

- **fork rd,rs,rt**: fires up a thread on a free TC (if available, see below). *rs* points to the instruction where the new thread is to start, and the new thread's *rd* register gets the value from the existing thread's *rt*.

Some vital per-TC state is copied from the parent. That's whether you're in kernel or user mode, defined in *TCStatus[TKSU]*, (which is the same as *Status[KSU]*); and what address space you're part of, defined in *TCStatus[TASID]*, (which is the same as *EntryHi[ASID]*.)

When the thread has finished its job it should use **yield \$0** to free up the TC again.

**fork/yield** are the only MIPS MT instructions usable in user mode (they're also highly original, and are likely not to be extensively used in early MIPS MT architecture applications using substantial OS layers - they might be hidden inside the OS, but you won't see them for a while in Linux user code).

**fork** will only select a TC which is both "free" (*TCStatus[A]* is currently zero) and which is specifically marked as usable by fork because *TCStatus[DA]* is set.

**fork** may fail if a suitable TC isn't waiting at the "taxi-rank". In that case you get an exception ("Thread Overflow") which an OS may catch and fix up before restarting the application; that way the application remains unaware of the problem. This provides the illusion of an indefinite supply of TCs, in the same way that a virtual memory system provides an indefinite supply of memory - you'll hear this described as that "**fork** has been virtualized" or made *Virtualizable*.

**yield \$0** has a matching "Thread Underflow" exception, which occurs when you're about to reach a situation where all for-hire TCs are parked (because then the system might stop forever, with no threads running the code which might make another thread run...).

There's a lot more to say about yield, see the bullet below and [Section 2.8.1, "Yield, Yield Qualifiers and threads waiting for hardware events"](#).

- **mfttr rd,trno,u,sel,h** and **mttr rt,trno,u,sel,h**: are privileged (CP0) instructions ("move to/move from thread register") which provide read/write access to another TC's registers.

The other TC is identified by *VPEControl[TargTC]*. The **trno,u,sel** fields identify which register of that TC you are accessing. Their encoding is complicated: we'll present details in [Table 2.1](#) below, but here's a quick summary:

- When **u**==0, **trno** is a CP0 register and **sel** is the auxiliary 3-bit "select" field found in **mtc0/mfc0**;
- When **u**==1 and **sel**==0, **trno** is a general purpose register;
- When **u**==1 and **sel**>0, you get to access more exotic registers, as detailed in [Table 2.1](#) below.

The **h** value should be specified as 1 when you're obtaining the high half of a register which is double the size of a GPR. In other cases, omit it. However, this argument is *not* required for the multiply unit accumulators, where the low and high half have separate **trno** register numbers.

That's fairly confusing, and the details are presented again in [Table 2.1](#).

**Table 2.1 MTTR/MFTR - "U" and "SEL" values**

u	sel	trno	Other TC's register type
0	0-7	0-31	CP0 registers.
1	0	0-31	General-purpose integer registers
1	1	0	Multiply unit <i>lo</i> and <i>hi</i> respectively.
		1	
		4	Low and high half (respectively) of DSP accumulator 1
		5	
		8	Low and high half (respectively) of DSP accumulator 2
		9	
12	Low and high half (respectively) of DSP accumulator 3		
13			
		16	<i>DSPControl</i> register.
1	2	0-31	Floating point (CP1) registers
1	3	0-31	Floating point control registers, as usually accessible with <b>cfcl</b> / <b>ctcl</b> .
1	4	0-31	Co-processor 2 data and control register sets, respectively. Implementations are free to define large CP2 register sets; the MT ASE provides an extra 5-bit "rx" field to provide more bits for selecting the CP2 register, but the MT ASE does not define a standard assembler syntax to generate it.
1	5	0-31	

The hardware does nothing, inherently, to make sure that register changes as a result of other-thread activity are seen tidily; unless you are really sure that the other thread is currently leaving the register alone, it's safer to ensure that the other TC is halted (shut down for maintenance) before **mftr/mttr** will work reliably.

When disassembling binary code it is painful to have to hand-decode the **trno**, **u**, **sel**, **h** fields, so tool providers are recommended to support the alternative "idioms" described in Table 2.3 below, which are probably more memorable than binary numbers. Most tools will be symmetric, so you will be able to write the idioms too: but that doesn't necessarily mean you *should* write code with them. You will, I hope, use meaningfully-named C pre-processor constants for all the various fields in your assembly source code, so it may be kinder on those who come after you if you expect them to remember just the **mttr/mftr** mnemonics.

Note that access to the registers of a TC affiliated to a different VPE is available only when *VPEConf0[MVP]* is set - it's often used as a safety-catch. In some environments (where you're not meant to be able to get at the other VPE's state) you'll find you can't set *VPEConf0[MVP]*.

If you attempt to read a register number which is not valid on your CPU, you will get an all-ones (-1) value back.

- **dmt**: suspend all other threads affiliated to the same VPE.

Under the hood this atomically clears the *VPEControl[TE]* bit, returning the original value of *VPEControl* to an optional register argument<sup>1</sup>; so it is convenient to bracket a piece of code which needs to be single-threaded within the VPE by:

```
dmt rt
ehb    # need hazard barrier to be sure it took effect
...    # guaranteed to be the only live TC in this VPE
mtc0  rt, VPEControl
```

1. They fit in with the encoding already used for atomic update of a CP0 register by the disable-interrupts instruction **di** etc.

The “hazard barrier” should always be used when you change some CPO condition and need to know it’s taken effect when you run a subsequent instruction - see [Section 7.1, "Hazard barrier instructions"](#).

OS code which updates registers and resources which are only replicated per-VPE will typically need this kind of protection, unless already multithreading-protected by something higher-level.

The **emt** instruction atomically sets the *VPEControl[TE]* bit and returns the old value. It is relatively rarely used; it’s more robust to replace the whole original value of *VPEControl* with an **mtc0**, because then things still work if you inadvertently nest one single-threaded block within another.

- **dvpe**: suspend/un-suspend all other threads, even those in other VPEs. In many systems VPE independence is much prized, and then this instruction is likely to be restricted to initialization software. Under the hood it clears the *MVPCControl[EVP]* bit, returning the old value. Again, there’s an **evpe** instruction, but a single-threaded block is better terminated by restoring the whole *MVPCControl* register with an **mtc0**.
- **yield rd,rs**: a multi-purpose instruction, whose action depends on the value in *rs*. If and when it returns, *rd* is set to a bit-vector which shows the active inputs to **yield** - at least those enabled by the *YQMask* register. More in the section below.

So:

- When *rs* == 0: (also discussed under the bullet called “fork” at the start of [Section 2.8, "MIPS® Multithreading ASE - new instructions"](#)) terminate the thread and clear the *TCStatus[A]* bit, permitting re-allocation to another purpose by **fork**. If this was the only live TC with *TCStatus[DA]* set (that is, the last TC in the **fork** pool), you get a “thread underflow” exception.
- When *rs* == -1: polite pause while other threads get a chance to run. To be more precise, the thread will be stopped briefly while the yield indication is sent out to an external scheduling policy manager, if fitted (see [Section 2.5, "Thread-scheduling decisions and the policy manager"](#).) Such a policy manager may respond, in particular, to changes communicated by writing the *TCSchedule* and/or *VPESchedule* registers.

After this sort of **yield** this thread will not run again for long enough that the policy manager has time to respond. But the thread hasn’t been stopped and will normally run again soon, at the priority newly determined by the policy manager.

- When *rs* == -2: has no scheduling effect, purely done for the value delivered to *rd*. And a **yield -2** never produces a “yield scheduler” exception.
- when *rs* > 0: waits for one or more of a set of signals to be asserted; from up to 31 signals available on your CPU, it is sensitive only to those selected by a “1” bit in the *rs* value. That’s complicated, see [Section 2.8.1, "Yield, Yield Qualifiers and threads waiting for hardware events"](#) below.

But in particular, if the *rs* value includes a bit which is *not* set in *YQMask*, you get an “invalid qualifier” exception.

Software can ensure that *any yield* which would deschedule a thread (or any **yield -1** whose return status would be zero) produces a “yield scheduler” exception. A secure OS might do that because it wants to “scrub” the TC’s registers of any application data before the TC is returned to the free pool. To achieve this effect set *VPEControl[YSI]* (the “did any work” test depends on *TCStatus[DT]*.)

A **yield** instruction must not be in a branch delay slot.

## 2.8.1 Yield, Yield Qualifiers and threads waiting for hardware events

When the *rs* argument of the **yield rs** instruction is positive, the thread waits for a hardware condition; the thread will wait until the bitwise-and of *rs* and the hardware signal vector is non-zero. This is a cheap and efficient mechanism to get a thread to wait on the state of some input signal.

Cores in the 34K family may have up to 16 external hardware signals attached. Because the **yield** instruction is available to user (low-privilege) software, you might not want it to have sight of all your hardware signals. The CP0 register *YQMask* is a bit-field where a “1” bit marks an incoming signal as accessible to the **yield** instruction.

In any OS running more threads than TCs you might want to reclaim a TC blocked on such a **yield**. If you need to do that while continuing to monitor the condition, then you’ll probably want your system integrator to ensure that the yield condition is also available as an interrupt, so you can get the OS’ attention when the condition happens.

The OS can zero-out corresponding bits 0-15 of *YQMask* to prevent them being used - a **yield rs** which attempts to use one of those bits will result in an exception.

In the two-operand form **yield rd,rs** the *rd* register gets a result, which is a bit-map with a 1 for every active yield input which is enabled by *YQMask* (bits which are zeroed in *YQMask* may return any value, don’t rely on them). The single-register form **yield rs** is really **yield \$0,rs**.

## 2.8.2 All MT instructions in alphabetical order

That’s in [Table 2.2](#) - but not quite all. There are a large number of convenience mnemonics (“assembler idioms”) which are not separate instructions, but which map onto some variant of the access-other-TCs-register instructions **mttr** and **mftr**. Rather than fill the table with these, we’ve consigned them to [Table 2.3](#) below. If you’re looking up an unfamiliar instruction, please look in both tables.

**Table 2.2 MT instruction summary in alphabetical order**

<i>Instruction</i>	<i>Brief Description</i>
<a href="#">dmt rt</a>	Clear <i>VPEControl[TE]</i> , which suspends execution of all other TCs affiliated to the same VPE. The <i>rt</i> register receives the original value of <i>VPEControl</i> ; if you don’t specify <i>rt</i> it defaults to <i>\$0</i> .
<a href="#">dvpe rt</a>	Disable all multithreading, including any other TCs affiliated to other VPEs, leaving this thread running alone. Implemented as an atomic clear of the <i>MVPControl[EVP]</i> bit. If you specify a register <i>rt</i> it receives the previous contents of the <i>MVPControl</i> register.
<a href="#">emt</a>	The “enable” pairs of <b>dmt</b> / <b>dvpe</b> respectively.
<a href="#">evpe</a>	You may not need these instructions: when you’ve finished a section of code which must be single-threaded in some sense, it may be preferable to restore the whole <i>VPEControl</i> / <i>MVPControl</i> register from the value you got back when you ran the disable instruction, as suggested in the description of <b>dmt</b> in the running text above.
<a href="#">fork rd,rs,rt</a>	Find a TC and activate it, so it starts at <i>rs</i> . The new thread’s <i>rd</i> register will be set to the value provided in <i>rt</i> . Lots more details above.

Table 2.2 MT instruction summary in alphabetical order

<i>Instruction</i>	<i>Brief Description</i>
<code>mfr rd, trno, u, sel, h</code> <code>mttr rt, trno, u, sel, h</code>	<p>“Move from thread register” and “Move to thread register” - get/set the value of a register belonging to some other TC, using the general-purpose register <i>rd</i> as a sink, or <i>rt</i> as a source. For a per-VPE register, you will access the VPE affiliated to the target register - so to access a VPE first set up a TC affiliated to it.</p> <p>The other TC is identified by <i>VPEControl[TargTC]</i>, and the register you’re accessing is selected by all of <i>trno</i>, <i>u</i> and <i>sel</i> - as described above or in Section 2.1, “MTTR/MFTR - “U” and “SEL” values”.</p>
<code>yield rd, rs</code>	<p>A multi-purpose instruction, whose action depends on the value in <i>rs</i>. When <i>rs</i>==0, it terminates the thread and makes the TC available for a subsequent <b>fork</b>.</p> <p>When <i>rs</i>== -1, pauses while other threads run and any scheduling policy change filters through.</p> <p><b>yield</b> with <i>rs</i>== -2 is just done to poll yield inputs.</p> <p>When <i>rs</i>&gt;0, you wait for one of the yield input signals, but only one for which there’s a corresponding bit set in <i>rs</i>.</p>

Table 2.3 MTTR/MFTR "assembler idioms" in alphabetical order

<i>Write as</i>	<i>Equivalent</i>	<i>Description</i>
<code>cftc1 rd, ft</code>	<code>mfr rd, ft, 1, 3</code>	Get data from/send data to another TC’s floating-point (coprocessor 1, CPI) control register <i>ft</i> .
<code>cttc1 rt, ft</code>	<code>mttr rd, ft, 1, 3</code>	
<code>mftc0 rd, rt</code>	<code>mfr rd, tc0r, 0</code>	Read other TC’s CP0 register.
<code>mftc0 rd, rt, sel</code>	<code>mfr rd, tc0r, 0, sel</code>	
<code>mftc1 rd, ft</code>	<code>mfr rd, ft, 1, 2, 0</code>	Read low 32 bits from other TC’s floating point data register.
<code>mftdsp rd</code>	<code>mfr rd, 16, 1, 1</code>	Read other thread’s <i>DSPControl</i> register.
<code>mftgpr rd, rt</code>	<code>mfr rd, rt, 1, 0</code>	Read other thread’s general purpose register <i>rt</i> .
<code>mfthc1 rd, ft</code>	<code>mfr rd, ft, 1, 2, 1</code>	Read high 32 bits from other TC’s floating point data register.
<code>mfthi rd</code>	<code>mfr rd, 1, 1, 1</code>	Read the other TC’s <i>hi</i> multiply/divide unit register, which is the same as the first of...
<code>mfthi rd, ac0</code>	<code>mfr rd, 1, 1, 1</code>	Read the high half of one of the other TC’s <i>ac0-3</i> DSP accumulators.
<code>mfthi rd, ac1</code>	<code>mfr rd, 5, 1, 1</code>	
<code>mfthi rd, ac2</code>	<code>mfr rd, 9, 1, 1</code>	
<code>mfthi rd, ac3</code>	<code>mfr rd, 13, 1, 1</code>	
<code>mftlo rd</code>	<code>mfr rd, 0, 1, 1</code>	Read the other TC’s <i>lo</i> multiply/divide unit register, which is the same as the zeroth of...
<code>mftlo rd, ac0</code>	<code>mfr rd, 0, 1, 1</code>	Read the low half of the other TC’s <i>ac0-3</i> DSP accumulators.
<code>mftlo rd, ac1</code>	<code>mfr rd, 4, 1, 1</code>	
<code>mftlo rd, ac2</code>	<code>mfr rd, 8, 1, 1</code>	
<code>mftlo rd, ac3</code>	<code>mfr rd, 12, 1, 1</code>	
<code>mttc0 rt, rd</code>	<code>mttr rt, tc0r, 0</code>	Write other TC’s CP0 register.
<code>mttc0 rt, rd, sel</code>	<code>mttr rt, tc0r, 0, sel</code>	
<code>mttc1 rt, fd</code>	<code>mttr rt, fd, 1, 2, 0</code>	Write data from <i>rt</i> to the high half of the other TC’s floating point register <i>fd</i> .
<code>mttdsp rt</code>	<code>mttr rt, 16, 1, 1</code>	Write to the other TC’s <i>DSPControl</i> register.
<code>mttgpr rt, rd</code>	<code>mttr rt, rd, 1, 0</code>	Write to the other TC’s general purpose register <i>rd</i> .
<code>mtthc1 rt, fd</code>	<code>mttr rt, fd, 1, 2, 1</code>	Write data from <i>rt</i> to the low half of the other TC’s floating point register <i>fd</i> .

**Table 2.3 MTTR/MFTR "assembler idioms" in alphabetical order**

<i>Write as</i>	<i>Equivalent</i>	<i>Description</i>
<b>mtthi rt</b>	<b>mttr rt,1,1,1</b>	Write to the other TC's <i>hi</i> multiply unit register, which is the same as the zeroth of...
<b>mtthi rt,ac0</b>	<b>mttr rt,1,1,1</b>	Write to high part of the other TC's <i>ac0-3</i> accumulator.
<b>mtthi rt,ac1</b>	<b>mttr rt,5,1,1</b>	
<b>mtthi rt,ac2</b>	<b>mttr rt,9,1,1</b>	
<b>mtthi rt,ac3</b>	<b>mttr rt,13,1,1</b>	
<b>mttlo rt</b>	<b>mttr rt,0,1,1</b>	Write to the other TC's <i>lo</i> multiply unit register, which is the same as the zeroth of...
<b>mttlo rt,ac0</b>	<b>mttr rt,0,1,1</b>	Write to low part of the other TC's <i>ac0-3</i> accumulator.
<b>mttlo rt,ac1</b>	<b>mttr rt,4,1,1</b>	
<b>mttlo rt,ac2</b>	<b>mttr rt,8,1,1</b>	
<b>mttlo rt,ac3</b>	<b>mttr rt,12,1,1</b>	



## 2.9 Multithreading ASE - CP0 (privileged) registers

All the CP0 registers new to or affected by the MT ASE are in [Table 2.4](#).

**Table 2.4 CP0 registers required by MIPS® MT ASE**

<i>Register</i>	<i>Description</i>
<b>Per-TC registers</b>	
<i>TCStatus</i>	Per-TC run-time control/status fields. Includes alternate views of per-TC fields in old CP0 registers.
<i>TCBind</i>	VPE affiliation and own TC number of this TC.
<i>TCHalt</i>	Per-TC - write one to halt the TC for maintenance, zero to let it run again. No further description needed.
<i>TCRestart</i>	Per-TC - address of instruction the TC will run next. Unambiguous only when the TC is halted. Writing <i>TCRestart</i> (to control where the TC executes from next time it is made live) has side-effects; in particular it clears the link bit which associates a load-linked/store-conditional pair, see <a href="#">Section 3.5, "Synchronization: "ll" and "sc" instructions implementation"</a> .
<i>TCContext</i>	per-TC 32-bit read-write scratch register for OS use, no hardware-interpreted fields.
<b>Per-VPE registers</b>	
<i>VPEControl</i>	Per-VPE - status and control fields for exception and <b>mftr/mttr</b> instruction support.
<i>VPEConf0-1</i>	read-only status of VPE setup
<i>YQMask</i>	bitfield where "1" bits define valid select bits for a <b>yield</b> instruction - see <a href="#">Section 2.8.1, "Yield, Yield Qualifiers and threads waiting for hardware events"</a> .
<i>VPEopt</i>	Can be used to mark any cache "way" (a quarter) of the primary I- and D-caches as inaccessible to the owning VPE, to keep it clear for the other one.
<i>SRSCConf0-4</i>	Fixes which TC's GPRs are recycled as shadow register sets.
<b>Whole-CPU control and availability of MT resources</b>	
<i>MVPCControl</i>	writable register to determine how multiprocessing facilities work.
<i>MVPCConf0-1</i>	read-only summary of CPU MT resources
<b>Software hints and controls on thread scheduling</b>	
<i>TCSchedule</i>	Per-TC, writable register to influence thread scheduling. It's not really part of the core, and the description is for our sample thread scheduling policy manager.
<i>VPESchedule</i>	Per-VPE, writable register to influence scheduling
<i>TCSchefBack</i>	Optional read-only register providing statistical information about thread scheduling.
<b>New fields in old registers</b>	
<i>EBase[CPUNum]</i>	Identity of running VPE within CPU
<i>Config3[MT]</i>	set if this CPU implements the MIPS MT ASE.
<i>Debug[OffLine]</i>	a per-TC bit which a debugger can set to quiesce a TC while it debugs another thread (but without affecting any non-debug state).

### 2.9.1 What CP0 registers are per-TC, per-VPE and per-CPU?

At first sight the CP0 register map looks like quite a chaotic mixture of fields replicated per-TC, per-VPE or not replicated at all. But in fact the rules are fairly simple, and there are only a few special cases:


- All registers called "TCxx" are per-TC.

- All other CP0 registers (not called "TCxx") are per-VPE except for:
  - *MVPCControl*, *MVPCConf0-1* are not replicated, there's just one set on a CPU.
  - The performance counter count and control registers are per-CPU.
  - A handful of fields in pre-MT MIPS32-standard registers are replicated per-TC: they include those which are found in *TCStatus* as fields called "Txx", plus the debugger controls *Debug[OffLine]* (a "thread halt" control for debuggers) and *Debug[SSSt]* (single-step).

Like all other CP0 registers, many fields are not initialized by hardware when the CPU is reset. And - special for MT - CP0 registers other than those belonging to VPE #0 and TC #0 are *not initialized at all*. Unless you are confident that random contents in some particular register are safe, it's your responsibility to write registers to sensible values.

## 2.9.2 VPEControl

**Figure 2-1 Fields in the VPEControl register**



In *VPEControl*:

*YSI*, *GSI*: "intercept" bits for yield and *Gating Storage* operations.

By setting one or both of these bits 1, an OS can arrange to be notified (by an exception) if any thread would otherwise become blocked by a **yield** instruction, or on an access to gating storage. The exception will only happen if the TC's *TCStatus[DT]* bit is set, that is if the TC has run an instruction since it was last deallocated.

*YSI* affects any **yield** instruction which would block; but a **yield** which tests for a condition which is already true, or a **yield -2** will not be affected (a **yield -2** is just a poll - see [the bullet on "yield"](#))

*GSI* affects any gating storage access which will block the thread<sup>1</sup>.

*EXCPT*: encodes the cause of the last thread exception. This refines the information returned by *Cause[ExcCode]* - we don't have enough reserved values to encode all the thread exceptions separately. Like the old cause register field, *VPEControl[EXCPT]* is only valid so long as the TC remains in exception mode (recall that in exception mode only one TC within the VPE may run).

The possible values are:

**Table 2.5 Thread exception codes in VPEControl[EXCPT]**

0	Thread overflow on <b>fork</b> .
1	Thread underflow on <b>yield</b> .
2	Bad qualifier fed to <b>yield</b> .
3	Exception on gating storage operation
4	<b>yield</b> which would have blocked run while <i>VPEControl[YSI]</i> is set to 1.
5	Gating storage access which would have blocked attempted while <i>VPEControl[GSI]</i> is set to 1.

1. Gating storage operations are uncached, and may be slow; but *GSI* won't lead to an exception because the access is slow, only if the gating storage interface is told that the operation is blocked.

*TE*: "enable multithreading" - when clear, only one TC attached to the VPE is allowed to issue instructions. You normally set/unset this using the **dmt/emt** instructions.

*TargTC*: selects the TC number of the "other thread context" in the **mttr/mftr** instructions - TCs are numbered from 0 upward. There isn't room to encode the TC number in the instructions. Note that since the whole of *VPEControl* is a per-VPE register, TC-multithreaded software will need some kind of lock (perhaps **dmt/emt** brackets) around any code which uses **mttr/mftr**.

### 2.9.3 TCRestart, TCHalt and TCContextt

Three registers without internal fields:

- *TCRestart* holds the thread's "PC" - more accurately, when the TC is halted it holds the instruction address where the TC will restart. If a TC is to retry an instruction in a delay slot, *TCRestart* will point to the branch but the *TCStatus[TDS]* flag will be set.
- *TCHalt*: just write a 1 to the register, and the TC will halt and will be safe to inspect and reprogram. Write a zero to let it run again. *TCHalt* is for the use of MT-aware OS code.
- *TCContext* is a pure 32-bit software register, without any hardware effect. OS software finds it useful to have a per-TC register where it can write an ID or key pointer which identifies the thread.

### 2.9.4 TCStatus

**Figure 2-2 Fields in the TCStatus register**

31	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	0
TCU3-0	TMX	0	RNST	0	TDS	DT	0	TCEE	0	DA	0	A	TKSU	IXMT	0	TASID							

*TCU3-0*, *TMX*, *TCEE*, *TKSU*, *TASID*: These fields - most of those starting with a "T", in fact - provide convenient alternate views of some fields in legacy CP0 registers. They are fields which, with MIPS MT, need to be replicated per-TC. This is valuable because code can change them without the difficulty of doing a non-atomic read-modify-write sequence on one of the legacy registers (which would mean having to read-write many fields which are shared with any other TCs in the VPE.)

The fields listed are views of the *Status[CU0-3]*, *Status[MX]*, *Status[CEE]*, *Status[KSU]* and *EntryHi[ASID]* fields respectively, and to get a complete view of what any of them do you are recommended to look at the notes on each of those CP0 registers.

We'll deal with these alternate-view fields first:

*TCU3-0*: set a bit to enable this TC to run instructions for the corresponding co-processor. There are four bits but only two feasible co-processors: CP1 is the floating point unit (if fitted) and CP2 is available for custom use. CP3 is not available on cores in the 34K family, so *TCStatus[TCU3]* always reads zero.

The floating-point unit available as coprocessor 1 for 34K family cores has only one set of registers, so it may only be used by one TC. It is the OS' responsibility to make sure that's done. Custom or future coprocessors may replicate all their state per-TC (so they may be freely multithreaded) or provide fewer, perhaps just one, register set.

*TMX*: another view of *Status[MX]*, which is the enable bit for the instructions in the MIPS DSP ASE (in theory it plays the same role for the older but less-used MDMX, but that will never be found in a 34K family core.) It's visible here because it's a per-TC field.

*TCEE*: another view of the per-TC *Status[CEE]* enable bit implemented by a "CorExtend" (user-defined) instruction block which needs it, usually because it includes some registers which it may need a kernel to save across exceptions and context switches.

*TKSU*: a convenient view of the per-TC *Status[KSU]* bits which determines the privilege state of the CPU.

*TASID*: a convenient alternative view of the per-TC *EntryHi[ASID]* field.

*RNST*: (read-only) status, which can be used to find out why a blocked TC is blocked, with values meaning:

- 0 Not blocked.
- 1 Asleep after a **wait**.
- 2 Blocked on **yield** (that is, waiting for one or more of the *Yield Qualifier* signals to activate).
- 3 Waiting for gating storage load/store to complete.

*TDS*: qualifies the per-TC restart address *TCRestart*, indicating that the thread is stopped in a branch delay slot (in which case *TCRestart* will point to the branch.) An analogue of *Cause[BD]*.

*DT*: "dirty" bit - set whenever the thread being run by the TC makes progress. More precisely, set when any of this TC's instructions completes (though instructions in exception, error-exception or debug mode are not counted); it is also set if the TC is successfully started as a result of a **fork**.

This is for the use of an OS overseeing applications forking at user level; it can inspect its free-TC pool and discover which ones have done any work since last time it looked. This may be important, because a TC which has done work for one application<sup>1</sup> might have some of that application's data in its registers, and cannot be automatically allocated to a different application for fear of leaking data (applications are not supposed to see one another's data). The OS must make sure it scrubs all the TC's registers before that happens; this bit is part of the mechanism which lets it find out when it needs to scrub.

*DA*: "dynamically allocatable" - when clear, this TC may not be allocated as a result of a **fork**.

If as a result a **fork** can't find a TC to use, it takes a "thread overflow" exception.

If the only dynamically allocatable and live TC executes a **yield** (which might lead to the silence of the grave), it takes a "thread underflow" exception instead.

*A*: "activated" (sometimes, "allocated"). Can't run instructions without this bit, which is set by **fork** and cleared by **yield \$0**.

*IXMT*: set 1 to prevent this TC from handling interrupts.

### Summary TC status

The *TCHalt* and *TCStatus[A,DA]* fields interact as shown in [Figure 2.6](#) and may be best understood together. Note that the EJTAG debug *Debug[OffLine]* bit, if set, overrides all of them and prevents the TC from live its thread. *OffLine*, though, doesn't affect whether a TC may be selected by a **fork**.


---

1. In a Linux context "process" would be more precise than "application".

**Table 2.6 TC summary state as expressed in per-TC register fields**

Register bits			What happens to TC
TCHalt	TCStatus		
	[A]	[DA]	
1	X	X	TC is halted, fit for inspection and maintenance by software running on some other TC
0	0	0	TC is not running, nor may it be used by <b>fork</b> .
0	0	1	TC is "parked at the taxi-rank" ready to be used as a result of a <b>fork</b> instruction
0	1	X	TC is working through the instructions of some thread. Maybe it's not currently live, but that will be for a thread-specific reason.

## 2.9.5 TCBind

**Figure 2-3 Fields in the TCBind register**

*CurTC*: (read-only) returns the TC's own identity - the ID of the TC which ran the **mfc0** instruction, or which is the target of a **mftc**.


The zero field which occupies bits 18-20 below *CurTC* is reserved by the architectural definition. It can save a mask operation when you need to use *CurTC* as an index into an array of 4- or 8-byte objects.

*TBE*: set by hardware when a transaction causing a bus error is identified as resulting from a load issued by this TC: see [Section 6.3.6, "Bus error exception"](#) for details. It remains set until explicitly written to zero.

*CurVPE*: the ID number of the VPE affiliation of this TC. You write this field to change a TC's affiliation, but only when the "VPE configuration mode" safety-catch bit *MVPControl[VPC]* is set. In principle it's possible for a thread to set its own TC's affiliation, but that seems fraught with difficulty. This will most often be set by some supervisory thread using an **mttc** instruction).

## 2.9.6 MVPConf0-1 - read-only multithreading-specific configuration information

The *MVPConf0-1* registers present a read-only summary of the CPU's multithreading resources.

**Figure 2-4 Fields in the MVPConf0-1 registers**

*MVPConf0[M]*: a "continuation" bit - if zero, *MVPConf1* isn't implemented, and acts as if it was all-zero.

*MVPConf0[TLBS]*: 1 if it's possible to share the TLB. To do that you'd have to set *MVPControl[STLB]*, see [Figure 2-5](#).

*MVPConf0[GS]*: reads 1 if the CPU is able to support *Gating Storage* as described in [Section 2.7.1, "Gating storage"](#).

*MVPConf0[PCP]*: read-only bit. Reads 1 if it's possible to deny access to one or more primary cache "ways" to each VPE. This feature must be enabled in *MVPControl[CPA]* and the way inhibition programmed in *VPEopt*, as described in [Section 2.9.10, "VPEopt register - reserve some cache "way" for use of one VPE"](#).

*MVPCConf0[PTLBE]*: the number of TLB slots which may be provided to different VPEs according to initialization software. Will read zero - even on a CPU with a sharable TLB - if the TLB configuration has no option other than shared or split in some fixed way.

*MVPCConf0[TCA]*: reads 1 on 34K, because it is possible to dynamically assign TCs to a VPE (by writing *TCBind[CurVPE]*.) Other MIPS MT implementations may not let you do that.

*MVPCConf0[PVPE,PTC]*: how many separate VPEs/TCs respectively are available on this CPU (the field encodes “number of things minus one”, so that zero means “one VPE” (or TC).

*MVPCConf1[C1M]*: the floating point unit (co-processor 1) implements the MDMX™ extension to the instruction set, as described in [MDMX]. This will always be zero on CPUs in the 34K core family.


*MVPCConf1[C1F]*: co-processor 1 implements 64-bit floating point instructions as defined in [MIPS64].

*MVPCConf1[PCX,PCP2,PCP1]*: how many register set contexts are available for CorExtend™, co-processor 2 and co-processor 1 respectively.

## 2.9.7 MVPControl Register - CPU-wide VPE control

*MVPControl* is a read/write per-CPU control/status register.

Figure 2-5 Fields in the MVPControl register



*MVPControl[CPA]*: set 1 to enable the per-VPE *VPEOpt* registers (see Figure 2-7 and notes) to deny a VPE use of one or more ways of the primary caches. Check *MVPCConf0[PCP]* first, to see whether this feature is available.

*MVPControl[STLB]*: set to enable TLB sharing between the VPEs, see Section 4.2.2, "Sharing and not sharing the TLB".

*MVPControl[VPC]*: "configuration mode" - a heavy-duty safety catch. When this bit is set to "1", it becomes possible to write to configuration register fields which are read-only on conventional MIPS32-compliant CPUs.

This is obviously a fairly dangerous thing to do, and it's unlikely to be a good idea to change the configuration registers except when launching a VPE with software which believes it is re-initializing itself. In particular, make sure that no other VPE is running by executing a **dvpe; ehb** sequence - the **ehb** ("hazard barrier") instruction makes sure that subsequent instructions don't start until the **dvpe** has taken effect.

But with this bit set ("unsafe"), a MIPS MT CPU can be set up by MT-aware software to configure a VPE with its choice of CPU resources, then pass that VPE to legacy (non-MIPS-MT-aware) software with that choice of resources presented by the standard *ConfigNN* registers.

With this bit zero, the fields in the *ConfigNN* registers revert to read-only.

*MVPControl* is writable only if the "master VPE" safety catch *VPEConf0[MVP]* is set to 1.

*MVPControl[EVP]*: when clear, instructions will only be executed for the thread which was running when this bit was cleared - even TCs affiliated to other VPEs will not be run. This bit is usually manipulated with the **dvpe/mtc0** instructions.

## 2.9.8 VPEConf0-1 registers - initializable per VPE resource lists

These are per-VPE registers which are read to show what resources are available to software on the VPE. Some fields are writable, but that's only when the VPE-access safety catch *VPEConf0[MVP]* is already set. That means that setting your own *VPEConf0[MVP]* to zero is an irreversible abdication from inter-VPE power if other VPEs do the same.

**Figure 2-6 Fields in the VPEConf0-1 registers**

	31	30	29	28		21	20	19	18	17	16	15			2	1	0
VPEConf0	1	0		XTC	0	TCS	SCS	DCS	ICS				0			MVP	VPA
	31		28	27	20	19	18	17					10	9	8	7	0
VPEConf1			0		NCX		0						NCP2		0		NCP1

*VPEConf0[XTC]*: When only one TC in a VPE is running because the VPE is in exception mode or *VPEControl[TE]* is clear, this field identifies that one running TC. *XTC* can be written by **mttr** as part of cross-VPE initialization if you want to initialize a VPE so it starts with just one TC running alone. Anything might happen if you tried to write this field on a running VPE, so you're prevented from doing so - the field is not writable unless the target VPE's *VPEControl[VPA]* is zero.

Of course the initializing thread, running such a **mttr**, will need its own copy of *VPEConf0[MVP]* set to do cross-VPE access in the first place.

*VPEConf0[TCS,SCS,DCS,ICS]*: read-only bits which tell software which caches are shared with at least one other VPE. The separate bits are for tertiary, secondary, L1 D-cache and L1 I-cache respectively. There's no way for a 34K core to be fitted with un-shared caches, so a 34K core will have *DCS* and *ICS* set (and will have the other bits set if it has L2 or L3 cache).

*VPEConf0[MVP]*: "master virtual processor" - a safety catch bit, which must be set before software can touch registers in different VPEs (or in the TCs of different VPE affiliation).

It also controls write access to *MVPControl*.

*VPEConf0[VPA]*: Virtual Processor Activated. If zero, no TCs bound to this VPE will run.

*VPEConf1[NCX,NCP2,NCP1]*: number of CorExtend, coprocessor-2 and coprocessor-1/floating-point contexts available to this VPE. These fields are writable at configuration time to zero, one or the number of TCs affiliated to the VPE<sup>1</sup> and will be reflected in the VPE's view of *Config[UDI]* (for CorExtend) and *Config1[C2,FP]*. If a thread within the VPE is to run a legacy OS, you can use that to determine whether the legacy software sees UDI, CP2 and/or floating point capability.

## 2.9.9 YQMask register - enable yield "conditions"

*YQMask* is a bit-map where you write a "1" bit to make the corresponding yield condition usable for the **yield mask** instruction, as described in Section 2.8.1 "Yield, Yield Qualifiers and threads waiting for hardware events".

1. If this field was set to "number of TCs" but the number of TCs affiliated to the VPE subsequently changes (it can happen) the field will be automatically updated.

### 2.9.10 VPEOpt register - reserve some cache "way" for use of one VPE

Two OS programs running on separate VPEs of a MIPS MT CPU progress independently of each other, and the thread scheduling rules usually make sure that each gets a fair proportion of the CPU's attention. However, one unavoidable interaction is that threads on both VPEs are competing for the same cache resources.

The 34K core's primary caches are 4-way set associative, and this is usually enough to provide for the active "working set" of all loaded threads.


Occasionally a critical routine may need such a good response time that it is unacceptable for it to be dislodged from the cache by an unrelated thread. Where this affects a tiny piece of code, your best bet is probably to lock the code concerned into the cache, as described in [Section 6.4.8, "Cache locking"](#).

But if some legacy code consigned to the independent environment of one VPE is suffering because of competition for cache locations from an unrelated program on the other VPE, you may also have the choice of reserving some part of the cache for the use of a VPE: to check whether this facility is available on your core, test *MVPConf0[PCPJ]*.

This facility is large-scale, affecting a whole cache "way" - that's 25% of a cache, removing one out of the four cache locations available to store any particular cache-line sized piece of memory. It's implemented by getting a VPE to sacrifice the ability to load data into one or more cache ways, making it unusable for any of the VPE's threads.

**Caution:** *You almost certainly shouldn't do this.* This is a facility offered to dig systems out of a very particular kind of hole. Only use it after careful measurement has convinced you that you have a problem caused by competition for cache resources, and keep measuring to make sure you're getting the effect you need.

**Figure 2-7 Fields in the VPEOpt register**



But once you're sure: to enable this facility at all, you need to set *MVPControl[CPA]*. Then to renounce the use of one of the four cache ways in the I- or D-cache for anything you miss on in future, set the corresponding *IWXnn/DWXnn* bit in the *VPEOpt* register, as shown in [Figure 2-7](#). Since this is done on a per-VPE basis, it's possible (though odd) to completely deny yourself use of some part of the cache. After CPU reset these fields are cleared to zero, so if you don't need this facility, just ignore it.

### 2.9.11 Shadow register configuration SRSCnf0-4

A TC's registers can be borrowed and used as a "shadow set" for another TC. These registers control how this is done. It seemed simpler to combine their description with the rest of the shadow register system in [Section 7.3, "Shadow registers"](#).

### 2.9.12 Thread scheduling hints - TCSchedule, TCScheFBack, VPESchedule

The *TCSchedule*, *TCScheFBack*, and *VPESchedule* registers are inputs to wholly implementation-dependent logic, so their description is not in this chapter, but in [Section 3.2, "Thread scheduling in the 34K core"](#) below.



# How the 34K™ core implements multi-threading

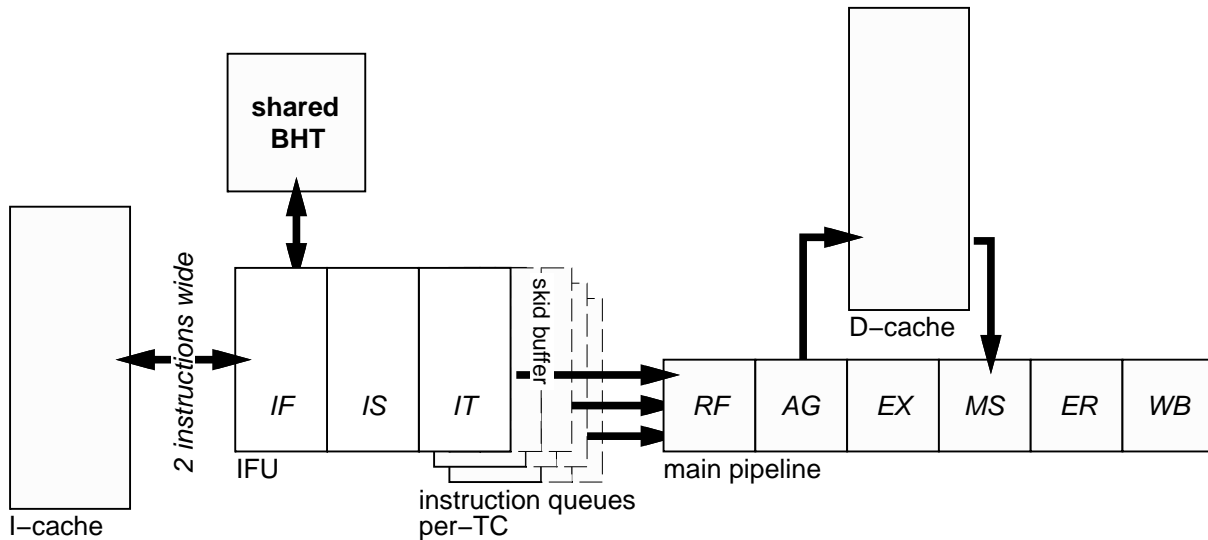
In this chapter:

- Section 3.1, "The 34K™ core pipeline and multithreading" how it all runs.
- Section 3.2, "Thread scheduling in the 34K core"
- Section 3.3, "Inter-thread communication storage (ITC)"
- Section 3.4, "The 34K™ core and interrupts"

## 3.1 The 34K™ core pipeline and multithreading

The 34K pipeline is shown in Figure 3-1. It inherits the 24K core's basic pipeline. You can find a simplified description of the 24K core pipeline in [PROG24K], which might be a useful introduction.

Figure 3-1 The 34K™ core pipeline



**Notes on the pipeline diagram Figure 3-1:**

The 34K core issues one instruction per clock and executes instructions for a particular thread strictly in order. We'll say an instruction is "fetched" when it's read from the I-cache in the IF stage, it's "issued" when it's sent to the RF stage and "executed" when it emerges from the ER stage without causing an exception.

- *Instruction fetch unit:* The instruction fetch unit ("IFU") occupies the three first stages and is decoupled from the main pipeline.

## How the 34K™ core implements multi-threading

Each TC in the processor has some of its own fetch unit state, and in particular each TC has a dedicated instruction queue which is kept filled whenever the TC is not stopped.

The IFU works a bit like a dog being taken for a walk. It rushes on ahead as long as the lead will stretch (the IFU, processing instructions two at a time, can rapidly get ahead). Even though you're in charge, your dog likes to go first - and so it is with the IFU. Like a dog, the IFU guesses where you want to go, strongly influenced by its observations of your habits. If you make an unexpected turn there is a brief hiatus while the dog comes back and gets up front again... but now we're anticipating the next bullet on "branch prediction". This kind of design is called a "decoupled" IFU.

Once instructions are issued from the IFU they are either completed in order or nullified (that is, essentially turned into a **nop** - such instructions continue to occupy a pipeline slot). When a thread stops for any reason (see [Section 2.4, "When can't threads run?"](#)) any of its instructions which have entered the main pipeline after the "stopping" victim will be nullified; the fetch unit holds the last couple of instructions issued in its *Skid buffer*, so it isn't necessarily going to have to go back and fetch the instructions from the I-cache all over again.

Even going back to the skid buffer is an avoidable overhead if the thread which stopped was the last runnable one. The hardware may detect this condition and decide to stall the main pipeline with the post-blockage instructions still in it when it knows there are no other runnable threads (in the hardware documents this is called "single-threaded mode").

- *Branch prediction*: The fetch unit has a couple of ways of predicting the branch target, allowing it to fill a TC's instruction buffer speculatively without waiting for the main pipeline to do calculations and report on branch conditions. It has:
  - *Target computation*: the fetch unit has logic which can compute the target of both PC-relative branches and long-displacement format **jal/j** instructions.
  - *A branch history table*: which is shared by all the TCs, is used to guess the direction of conditional branches. The table is indexed by the low address bits of the instruction's location, and keeps 2 bits of state for each slot. It's surprisingly effective, guessing right over 90% of the time. All branches (including the misnamed "branch likely" instructions) are treated the same.
  - *A return prediction stack*: a small stack on which the IFU pushes the return address of any subroutine call instruction. Subroutine return (i.e. **jr ra**) instructions pop the stack and guess that it delivers the correct target address.

When multiple TCs are running, only one of them may use this stack. A TC gets to use the stack whenever all other TCs are blocked for relatively long-term reasons, and gets to keep it (even though conditions change and other TCs become unblocked) until some other TC qualifies.

When the guess turns out to be wrong or the execution unit encounters an unpredictable computed branch the execution unit issues a *Redirect* and nullifies any of the TC's instructions in the pipe; the IFU has to discard all queued instructions for this TC, and start fetching again from the corrected address.

- *Main pipeline*: like the 24K core, the main pipeline is adjusted to provide something more than two clocks for accessing the L1 caches. It also shares the "skewed ALU" - load/store address calculation is done in the dedicated AG stage ahead of the EX stage where arithmetic and logical functions happen. The skewed ALU keeps the load-to-use delay down to just one clock.

There's no such thing as a free lunch; the downside is that a load/store instruction whose address generation depends on the immediately preceding instruction will have to wait for one clock. Compilers probably find it easier to move the address calculation back one place in the instruction stream, rather than to find yet another useful

instruction which can be moved between the load and use of the data. But code which follows a pointer chain is guaranteed to take at least three cycles per pointer.

### 3.1.1 Resource limits and consequences

The long pipeline, data interlocks, and the semi-autonomous IFU mean that the whole pipeline does not advance in lock-step as in the simplest MIPS architecture CPUs. Updates to internal states are not so easy to schedule at fixed times; instead they tend to wait in queues until a convenient moment. Most of the time, the convenient moment arrives quickly and there is no software-visible effect. But sometimes an unusual code sequence causes updates to be generated faster than they can be dealt with, the queue fills up and execution of the thread - perhaps the whole CPU pipeline - has to be stopped while the updates are done.

Queues which can fill up include:

- *Cache refills in flight*: there can be four or nine, at build-time option. In a single-threaded application you're unlikely to reach this limit unless you are using prefetch or otherwise deliberately optimizing loops. If a series of prefetches use all available resources, the next unrelated load-miss will stall the pipeline.

A hard-working multi-threading application might get there more often - hence the option to have nine entries in the "load data queue" in the 34K core.

- *Non-blocking loads to registers*: the 34K core has enough resources to have one load outstanding on each TC. They're used not just for non-blocking loads, but also for a TC blocked on gating storage.

That's enough so that compiled code is unlikely to reach this limit.

- *Lines evicted from the cache awaiting writeback (four)*: the 34K core's ability to write data will in almost all circumstances exceed the bandwidth available to memory: this queue will absorb short bursts without delaying anything. A long enough burst of writes will eventually slow to memory speed.
- *Register file write port (just one)*: that means that only one instruction can write a register value in each clock. For instructions which execute down the main CPU pipe this is not in the least problematic: they arrive at their register-write stage one at a time in sequence. But instructions with their own pipeline (multiply/divide operations, loads/stores, coprocessor operations) any result delivered to a general-purpose register must wait for a slot in which the main-pipeline instruction doesn't need to write a register. Typically, this happens very soon: but it depends on the instruction sequence.

### 3.1.2 Choosing instruction to issue

There's a critical piece of logic called the *Dispatch Scheduler* running in the IS/IT pipeline stages. It's job is to decide which TC's instruction to issue next, and that's the subject of the next section.

But if we look at the whole CPU, we see that the instructions in the main pipeline are not necessarily (nor even usually) all from the same TC. The hardware carries a TC number down the pipeline with each instruction, and that TC number is used to extend any register number defined by the instruction to read and write a register from the TC's own set.

This is fine, so long as no threads block. When a thread blocks, the fetch unit gets to know about it and will not issue any more instructions for that TC; but by that time some more instructions for that TC are likely to be in the main pipeline. These instructions are now doomed, but must be allowed to continue through the pipeline: otherwise instructions from unrelated, unblocked TCs could not make progress<sup>1</sup>. The doomed instructions are marked as "nullified": they will cause no exception, no load or store, and no register write-back.

## How the 34K™ core implements multi-threading

Meanwhile, the TC's own instruction queue is told of the blockage, and told where to restart after the thread is unblocked. The instruction at the head of the TC's instruction queue won't be the restart one (because one or more instructions were entered into the main pipeline and nullified). To avoid discarding the whole instruction queue and re-fetching all the instructions, the instruction queue includes a "skid buffer", which keeps a copy of a couple of instructions which have been issued but might still be nullified. So the TC which is stopped can back-up the skid buffer and wait to be running again.

## 3.2 Thread scheduling in the 34K core

In any multithreading CPU you have somehow to determine which instruction to run next.

The logic which does this job in the 34K core is called the *Dispatch Scheduler* ("DS"). On every cycle the DS selects an instruction from one of the per-TC instruction buffers and puts it into the main pipeline. Its decision is influenced by signals from the main pipeline but also by per-TC signals delivered from a piece of logic outside the core, the *Policy Manager* (PM). The simplest PMs just tie some interface signals to fixed levels; there are others which just feed back some bit-fields from the *TCSchedule* and *VPESchedule* registers. Customers can use a MIPS-supplied PM or create their own - for more details see [Section 3.2.3, "Policy managers available for the 34K™ core family"](#).

Instructions are fetched at the front of the IFU: so how do we choose the TC for which we'll fetch a pair of instructions for this clock? That's fairly simple. Fetch will rotate through each running TC which has room on its instruction queue (though there are minor tweaks in the hardware so an empty queue gets attention quickly).

### 3.2.1 The Dispatch Scheduler

The dispatch scheduler computes a priority for each TC. Where there are TCs with different priorities, it will do a cycle-by-cycle round-robin between the highest-priority TCs which are running.

The priority calculation includes the following bits, in something like this order:

- The MS bit represents "running": that means the priority mechanism automatically makes sure that we avoid selecting an instruction from a blocked TC (and we don't need any special purpose logic to do that).
- The priority includes the 2-bit per-TC priority which is supplied by the Policy Manager.
- The LS "priority" bits are for "round-robin" - again, the priority check is overloaded to implement the round-robin algorithm for otherwise-equal-priority TCs.

Along with the real TCs, the DS can have one or more *Relax* TC numbers; when a "relax" number wins the arbitration no instruction is issued, and perhaps some energy is saved. This feature is controlled from the external policy manager (see below) and in particular the *VPESchedule* register.

### 3.2.2 Policy manager interface

The interface is hardware, really. But if you are programming a core equipped with a custom PM, you probably need to know something about the hardware interface.

The *TCSchedule* and *VPESchedule* registers (if implemented at all) are inside the policy manager and their values may influence its behavior in any way the designer thinks fit.

- 
1. If there is only one running thread (so nothing else can happen if the stopped thread is evicted from the pipeline) the whole 34K pipeline may be stalled.

The policy manager supplies the core with:

- A 2-bit "group number" for each TC, mapping each onto one of four "scheduling groups".
- A 2-bit priority for every group. A TC then gets a priority from its group, which influences the internal dispatch scheduler as to which TC's instruction to schedule next.
- A per-TC "block" signal which when asserted freezes the TC completely. This is not used (that is, it's hard-wired deasserted) in MIPS Technologies' own PM designs.
- A set of "relax" signals corresponding to a bogus "relax" TC for each VPE; each has its own 2-bit priority and an enable.

The PM has access to many signals from the core. Per-TC information includes:

- VPE membership
- Instruction completion strobe.

(Signals below here are not used by any MIPS-designed PM):

- TC state: running, yielded, halted, suspended, waiting ITC, wait, used as SRS.
- TC running, as used by the dispatch scheduler. Note, though, that by the time the PM acts on signals like this it is always somewhat late; so it would be foolish to build hardware which attempted to respond to core signals without any "averaging".
- TC issue strobe, from DS.
- "TC has been forked". A 1-clock pulse asserted as a fork instruction completes.

Then there's some per-VPE information: the basic debug, exception and error-level bits. These are not used by any MIPS Technologies PM.

### 3.2.3 Policy managers available for the 34K™ core family

MIPS Technologies will ship the core with a couple of worked-example PMs, which themselves will be useful for many purposes. You can choose between:

#### **Equal priority (Basic round-robin)**

Just wire all the TC priority group values the same, and disable the do-nothing "relax" field in *VPESchedule*. Then you'll get simple round-robin influenced only by the heuristics used by the core to keep its pipeline full.

Many applications will work just fine with this simple mechanism. You are positively recommended to avoid using anything more complicated until you really understand why!

#### **Fixed priority**

Hard-wire TC groups and priorities as required, and disable "relax". TCs of equal priority will round-robin, but the scheduler will favor higher-priority TCs.

## How the 34K™ core implements multi-threading

The most likely arrangement is a two-level scheme offering higher priority for TCs to be used for threads which both (a) have real-time deadlines, and (b) can be trusted not to consume excess CPU cycles when they have no real work to do.

### The "Weighted Round Robin" (WRR) policy manager

It seems like a MIPS MT CPU with two running TCs can only be told to make them equal or to give one unconditional priority over the other. You might be interested in a system which - instead - would ensure that one of the TCs consistently got more cycles than the other, but that the less-favoured TC wouldn't be starved. You can do that, by feeding the CPU with a rapidly-changing set of priorities which average out to what you want.

The building block of this is a machine which runs through a sequence of states, allowing us to provide four distinct priority "groups": other things being equal, TCs in groups 0-3 get respectively 1/15, 2/15, 4/15 and 8/15 of the CPU. In our policy manager, we can now maintain a "priority group number" for each TC and have it turned into a cycle-by-cycle priority to achieve our goal.


We run a 15-cycle counter and in each cycle of 15 award different priorities to the groups as shown in [Table 3.1](#):

**Table 3.1 Dynamic priorities for finer resolution - group priority sequences**

	Priority in cycle (higher is better)														
Cycle	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Group0	0	1	0	1	0	1	0	3	0	1	0	1	0	1	0
Group1	1	0	1	3	1	0	2	1	2	0	1	3	1	0	2
Group2	2	3	2	0	2	3	1	2	1	3	2	0	2	3	1
Group3	3	2	3	2	3	2	3	0	3	2	3	2	3	2	3

The WRR PM uses dynamic priorities as shown above. You program it through the *TCSchedule* register, shown at [Figure 3-2](#).

**Figure 3-2 Fields in the TCSchedule and VPESchedule registers (WRR policy manager)**



*TCSchedule* and *VPESchedule* have lots of space for use by future (more sophisticated) policy managers. The fields defined are as follows:

*STP*: set 1 to prevent the associated TC running any instructions at all; for *VPESchedule[STP]* it disables the "relax" issue-nothing condition, which can be scheduled to save power.

*GRP*: determines which of the four priority groups this TC will be in, as described above; for *VPESchedule[GRP]* this is the scheduling group for the "relax" condition.

This policy manager does not define a *VPEScheFBack* register.

### TCScheFBack register

*TCScheFBack* counts up when any instruction is completed by this TC; it is an unsigned 32-bit value, which saturates at the maximum representable value. It is software's job to write it to zero or some other low value periodically.

### 3.3 Inter-thread communication storage (ITC)

ITC locations are magic memory locations used to provide low-level thread synchronization - which might be inter-thread (hence "ITC", from "Inter-Thread Communication storage") but could also be between customer-specific hardware and the software thread. Because ITC locations are places where threads wait for potentially long periods of time, they're accessed - always uncached - as *Gating Storage* - described in [Section 2.7.1, "Gating storage"](#) above. The ITC block is a piece of logic outside of the 34K core and connects through the gating storage interface. Because it's outside the core, SoC integrators are free to use the MIPS-supplied example logic in whole, in part, or to write their own. This section only describes the features of the sample ITC block supplied in the core package.

Each *ITC Cell* presents 32 bits of data. You should read/write these locations only as 32-bit data: partial-word loads and stores may misbehave. Be careful about compiled code too, to make sure optimization doesn't remove or alter any load or store operations. There are 16 different "views" of the same cell, all mapped to double-word boundaries for compatibility with 64-bit implementations, so each cell occupies 16×64-bits (128 bytes) of memory space. The different views have behaviors designed to support efficient implementations of popular synchronization operations, as listed in [Table 3.2](#). You can build your system with some or all of the ITC cells being FIFOs; to find out which cells are FIFOs look at the fields in the status view, described in [Figure 3-3](#).

**Table 3.2 ITC cell views and what they do**

Address within cell	Behavior
0	"bypass": load/store just read and write the data, without affecting the flags. If the cell is a FIFO, you write the newest entry and read the oldest (but <i>without</i> pushing the FIFO).
8	"status" view: read or write cell status as shown in <a href="#">Figure 3-3</a> .
16	"empty/full" synchronized view: the cell remembers whether anything has been written to it making it non-empty (and if it's not a FIFO, making it full at once). Loads from an empty cell block, as do stores to a full cell. A load from a full cell makes it non-full, and (eventually, if it's a FIFO) might empty it.
24	Empty/Full "try" (non-blocking) view. A load from an empty cell returns, but the data is always zero. A store to a full cell is quietly discarded, and the thread continues to run; but (more usefully) you can use an <b>sc</b> (store-conditional) instruction targeting this view and it will return 1 if the data was written, 0 if it was discarded.
32	"P/V" synchronized view: this implements a "P/V" counting semaphore. This synchronization trick was invented by Dijkstra - "p" and "v" are the "wait if zero, then count down" and "count up" functions respectively. A load from a zero cell blocks until a non-zero value appears. Otherwise the load returns the value and (atomically) decrements the stored value. Any store causes an atomic increment of the cell value, up to a maximum value of $2^{16}-1$ , at which it saturates. Stores never block. P/V operations do not modify the empty and full bits, which should both be cleared before an entry is used for P/V purposes. The P/V view of a FIFO location doesn't make sense, and the result of any such access is undefined.
40	P/V "try" (non-blocking) view. Same as the synchronized P/V view, except that a load does not block, even if the cell value is zero. Again, don't use this view for a FIFO cell
48-56	Reserved for future versions of the MIPS MT ASE.
64-120	Implementation-dependent views.

**Figure 3-3 Field layout in an ITC cell status view**

31	21	20	18	17	16	15	2	1	0
0	FIFO_PTR			FIFO	T	0	Full	Empty	

## How the 34K™ core implements multi-threading

The fields in the ITC cell status view mean:

*FIFO\_PTR*: the index of the oldest FIFO entry (on a read that's the next to be returned), but will read zero unless this is a FIFO cell. Write the ITC status view with *FIFO\_PTR* zero, *Full* = 0 and *Empty* = 1 to reset the FIFO.

*FIFO*: (read-only) 1 if this cell is a FIFO (that is, it has more than one word of storage.) In the ITC implementation distributed with the 34K family, all ITC FIFOs have four words of storage.

*T*: "trap" bit - causes any data access (i.e. any "empty/full" or "P/V" access) to this cell to result in an exception. Set by an OS which wants to keep track of reads and writes, perhaps because it's recycled a TC which was waiting here and wants to know when it might have been unblocked.


*Full/Empty*: described in Table 3.2. There are separate full and empty bits to allow ITC cells to quietly grow into FIFOs with multiple words of storage. Write *Empty* to 1 to reset the FIFO to a clean empty state.

### 3.3.1 Configuring ITC base address and cell repeat interval

The configuration information for the ITC space is held in two "tags" accessed by overloading the **cache Index Load Tag D** and **cache Index Store Tag D** instructions (it's much like the mechanism used for scratchpad RAM). Set the *ErrCtl[ITC]* bit to tell the instructions to access ITC space configuration "tags", and use addresses 0 or 8 in the **cache** instruction address field:

The ITC-configuration "tags" show up as in Figure 3-4.

**Figure 3-4 ITC configuration information**



What can you do with these?

*BaseAddress*, *AddrMask*: allow you to set the ITC starting physical address and region size, with at best a 1Kbyte resolution. Once this is set up and enabled, all accesses to this physical address range will go to ITC, and will no longer show up on the main system interface - so these locations will "overlay" anything else you expected to be there. Take care not to overlap any vital address.

The ITC cells can be put at any address whose alignment matches the total size of the ITC region (if you had 64 ITC cells at 256byte intervals you could place them at any 16Kbyte aligned address).

To do that set *AddrMask*:

<i>AddrMask</i>	<i>ITC region size</i>	<i>AddrMask</i>	<i>ITC region size</i>
0	= 1Kbyte	0xF	= 16Kbytes
1	= 2Kbytes	0x1F	= 32Kbytes
3	= 4Kbytes	0x3F	= 64Kbytes
7	= 8Kbytes	0x7F	= 128Kbytes

*ITC\_en*: set 1 to use ITC - it's zero from reset, making ITC invisible until you want it.



*NumEntries*: a read-only field which tells you how many 32-bit ITC cells are provided<sup>1</sup>.

*EntryGrain*: let's you control the cell spacing. Tightly spaced cells save on memory space, but widely spaced cells spread across a number of TLB pages, permitting different cells to be mapped to different processes. If you set the cell spacing very high, you'll limit the number of cells you can access in the usual ITC region.

When the *EntryGrain* field is zero, cells are packed at 128-byte intervals. Other values result in cells at intervals of  $128 \times 2^{\text{EntryGrain}}$  bytes, or:

<i>EntryGrain</i>	<i>ITC cell interval</i>	<i>EntryGrain</i>	<i>ITC cell interval</i>
0 =	128bytes	4 =	2Kbytes
1 =	256bytes	5 =	4Kbytes
2 =	512bytes	6 =	8Kbytes
3 =	1Kbyte	7 =	16Kbytes

To program these locations first set the *ErrCtl[ITC]* bit, which tells the cache instruction to access ITC information. Read the registers to find out how many ITC cells are available; then program your choice of cell interval and base address, with the region size set to match.

Don't forget to clear *ErrCtl[ITC]* afterwards, so that cache operations can continue as usual.

## 3.4 The 34K™ core and interrupts

As you may recall from [Section 2.6.1, "Multithreading and interrupts"](#), the interrupt system is replicated per-VPE; so the 34K core may have two interrupt systems. Interrupt inputs (including *Int0-5*, *NMI* and the EJTAG debug interrupt *DINT*) are presented separately for the two VPEs at the core interface.

Only the internally-generated timer and performance counter overflow interrupts are always local to the VPE (you can find out what interrupt number they use by looking in the *IntCtl* register shown as [Figure 7-1](#)).

In the 34K core any TC which is not interrupt-exempt may handle an interrupt. However, where there is a choice:

- An interrupt will be delivered to any thread which is asleep after a **wait** instruction (if there is one); otherwise:
- The interrupt will be delivered to any non-exempt, active thread which is not blocked waiting for a gating storage access; and only then:
- The interrupt will be delivered to an active-but-blocked thread.

See [Section 7.2, "MIPS32® Architecture Release 2 - enhanced interrupt system\(s\)"](#) for information about the interrupt signalling and handling options that the 34K core shares with other MIPS32 CPUs.

## 3.5 Synchronization: "ll" and "sc" instructions implementation

In coherent multi-processor or software multi-threaded systems, the **ll** and **sc** instructions work together to provide an RMW operation on a memory variable (with an arbitrary modification of the value) which succeeds only if it is guaranteed to have been atomic - that is, no other thread can have seen the value of the same variable between the read and the write. Moreover, **sc** returns a value which reports when atomicity could not be guaranteed, and the store

1. Earlier versions of this specification used a "logarithmic" code for number of entries.

## How the 34K™ core implements multi-threading

wasn't done; that allows software to build a retry loop to implement atomic operations. The risk of non-atomicity is detected by the cache snoop logic for cache-coherent multiprocessors, and by the intervention of an exception on software-scheduled uniprocessors.

The MIPS MT ASE requires that **ll/sc** also work between concurrent threads on an MT CPU. Each TC is equipped with a CP0 register called *LLAddr*, which remembers the physical address (at least to the enclosing 32-byte block) of the target location of any **ll/sc** sequence. The 34K core detects possible non-atomicity by checking every write made by any thread against the *LLAddr* of all other TCs.

The hardware keeps a single bit of state per TC called a "link bit" - the link bit is not directly visible to software. The link bit is most often zero, but is set by a **ll** instruction and then cleared by any condition threatening atomicity. It's cleared if:

- Some other TC's store is to the same block as our *LLAddr*;
- An **eret** instruction runs for this TC's VPE (which means there's been an exception, which could mean this TC has been rescheduled in the middle of its sequence);
- Some other software wrote this TC's *TCRestart* register to cause it to execute elsewhere. This is to catch conditions where OS software running on some other thread "reschedules" the TC: we don't want the link bit to survive such indignities.

Then the **sc** succeeds only if the link bit is still set when it executes.

In the MIPS MT ASE the **sc** instruction is also used to provide feedback from a store to an ITC location which might fail: see [Section 3.3, "Inter-thread communication storage \(ITC\)"](#).

## Initializing the 34K™ core - Multi-Threaded bootstrap issues

You are likely to deal with MIPS MT features at three stages as you boot the system:

- Boot the system, probably without dependence on the MIPS MT extension. It's good if first-level system bootstrap (which is bound to be awkward, system-dependent code) is not also sensitive to changes in the CPU feature set.

So the first-level bootstrap typically wants to make sure that any new multithreading behavior is suppressed until wanted: see [Section 4.1, "Bootstrapping without worrying about multithreading"](#).

- Set up VPEs and TCs.

Once you reach the point where you're running software which wants to exploit multithreading, you need to discover what resources the CPU has, and to set them up. That's described in [Section 4.2, "Configuring your choice of VPEs and TCs"](#).

Special care should be taken when you're initializing a VPE which is to run non-MT-aware "legacy" software - perhaps a whole legacy operating system: notes in [Section 4.2.1, "Setting up a VPE for legacy software"](#)

If you are running co-operative software on two VPEs and are able to make minor changes to the source code, it will usually be more efficient to share the TLB entries (the "legacy-ready" approach is really a hard-wired partition of the entries): see [Section 4.2.2, "Sharing and not sharing the TLB"](#).

- Thread initialization for explicit multi-threading, see [Section 4.3, "Setting up a TC to run a thread"](#).

### 4.1 Bootstrapping without worrying about multithreading

It's usually going to make sense to deal with the complexities of multi-threading only at the point in the system where you start to use the facilities. For many systems that means that initial bootstrap software (perhaps a boot monitor or reset-time diagnostic) will be better off ignoring multithreading.

Fortunately that's straightforward. A MIPS MT CPU comes out of reset with just TC #0 running, affiliated to VPE #0, and looking single-threaded. Moreover, the *VPEConf0[MVP]* bit is set, so the bootstrap software is all-powerful and can do whatever is required to set up the right VPEs and threads for the system.

Now bootstrap your computer. If the software needs to know it, it can read its own TC and VPE number from *TCBind*.

As always, bootstrap software is responsible for initializing CP0 registers; a register may only be skipped if you are certain that random contents in it will not disrupt your software.

## 4.2 Configuring your choice of VPEs and TCs

When you arrive at that software which wants to start an extra VPE or TC, you first need to discover what resources your CPU has. The pre-multithreading *Config* and *Config1-3* registers used to tell you everything; but in a MT CPU, those registers reflect just one VPE's resources, which in turn depend on what has been configured through *MVPControl* and *VPEConf0-1* (details in [Figure 2-5](#) and [Figure 2-6](#) and the notes to them) In fact, since in the 34K core initialization software has a fair amount of control over what resources are allocated to each VPE, some fields in the previously read-only *Config* registers are *writable*. However, that's only done where necessary: for example, since all VPEs share the caches, all VPEs can and do use read-only cache information from *Config1-2*.)

The total CPU resources are enumerated in *MVPConf0-1*, which you can see in detail in [Figure 2-4](#) .

### Getting a second VPE/TC into use

Suppose you have some software loaded into memory, but you want to start running it with TC #1 bound to VPE #1. Currently both the per-VPE and per-TC registers and other resources are in their post-reset state: the critical ones need software initialization before they can start.

Only a thread with *VPEConf0[MVP]* set can do this - fortunately VPE #0 will come out of reset with *MVP* set (if you already cleared it, you've resigned. Reset the CPU and try again!). Then set *MVPControl[VPC]*.

Set *VPEControl[TargTC]* to 1, the other TC's number, so you can write the other TCs registers with **mttr**.

You certainly don't want the other VPE running while you do this sort of thing, and you should clear your *VPEControl[EVP]* bit while you're working. You should probably use the **dvpe/mtc0** instructions as a "bracket", rather than manipulating the *EVP* bit directly. So your overall flow should be like this:

```

dvp t0
ehb                # execution hazard barrier, make sure dvp takes effect
...
set MVPControl[VPC]
(initialise VPE #1 and TC #1)
...
mtc0 t0, VPEControl # undo the dvp
    
```

OK, so now let's look at how you "initialize VPE #1 and TC #1".

From now on **mttr** will operate on the TC of your choice. You'll probably want to do quite a lot of set-up of both per-TC and per-VPE fields

So for TC#1:

- Set *TCHalt*. In fact TC#1 can't run anything yet, because you're still under **dvpe** control - but the MT specification allows CPUs to treat the "halted" state specially. Don't omit this.
- TC #1's VPE affiliation may not be set as you wish, so set *TCBind[CurVPE]* to 1 (the other VPE's number)
- A word of warning. This section lists all the "important" fields. When any MIPS CPU is powered up, only a rather small set of CP0 register fields are initialized. But when a MIPS MT CPU is powered up, only fields for VPE #0 and TC #0 are initialized at all. Your new VPE may have random garbage in any writable CP0 field. So if yours may be the first use of that VPE from power-on, iterate over all the CP0 registers setting all writable fields to "safe" values.

- Set *TCStatus*. *TCStatus[A]* will have to be set to 1 so the TC can run (this bit is the “allocated” bit for **fork**, and is required when setting up a thread manually.) All other bits can be zero, at least to start with - once you get something working, though, go back to the detailed description in for any other fields in [Section 2.9.4, "TCStatus"](#) on page 35.
- Set *TCRestart* to the program location where you want your new thread to start.

Now let's look at VPE#1's registers:

- Set *VPEControl* zero, and set *VPEConf0* to leave *VPEConf0[XTC]* = 1 (it should match your affiliated TC number, and must do so if you want to start the software “single-threaded”), *VPEConf0[MVP]* = 0, *VPEConf0[VPA]* = 1. For full details consult [Figure 2-1](#) and [Figure 2-6](#) and their notes.
- Set VPE #1/TC #1's *Status* register: If your intention is to have VPE #1's software mimic "coming out of reset" you might want to set its *Status[EXL]* set 1, so it starts in exception mode. Think about *Status[BEV]* - if set and your new thread takes an exception (which quite often happens with brand new code, due to one slip or another), then with *BEV* set it will use the ROM exception vectors, which are always shared with VPE #0 - and might not be what you wanted. On the other hand, if you do clear *Status[BEV]*, make sure you've set up code to catch any exceptions delivered at the non-ROM address. You might also set *EBase* to give VPE #1 different exception entry points from VPE #0 (unless you really want to share them).
- Set VPE #1's *VPEConf1* register. *VPEConf1[NCP2, NCP1, NCX]* determine whether your new VPE will be able to use coprocessors 2 and 1 (CP1 is the floating point unit) and the UDI instruction set, respectively. If the co-processor has only one bank of registers, you may well want to deny use of the co-processor to all but one of the VPEs.

At the end of our sequence you re-enable multithreading (by restoring the old value of *VPEControl*). Your last step is to use **mttz** to write TC #1's *TCHalt* to zero. Now TC#1 should start up and start running your code.

### 4.2.1 Setting up a VPE for legacy software

In general you can support only *one* piece of legacy software on a 34K family core. The VPEs see the same basic MIPS architecture memory map, and a few things are commonly shared - not least the exception entry points.

Your legacy software has to be told (by some means, beyond the scope of this manual) not to use all the physical memory in the system. Most likely the new MT-aware software will also need to use some virtual memory in the *kseg0/kseg1* regions, too.

The "legacy" VPE needs to be carefully set up to fool the old software into seeing and using a congenial MIPS32 CPU. That means:

1. Set up this VPE with just one TC;
2. You'll initialize all the relevant new MIPS MT registers and resources to keep the legacy software happy for its lifetime. Consult the full list of registers in [Section 2.9, "Multithreading ASE - CP0 \(privileged\) registers"](#).

### 4.2.2 Sharing and not sharing the TLB

It's not really visible to software whether there is really more than one TLB in any 34K core, but you can software configure it so that you get either:

## Initializing the 34K™ core - Multi-Threaded bootstrap issues

- *Hard partition*: each VPE appears to have its own TLB fully compatible with the MIPS32 architecture (the sizes of the VPE's TLBs are as set when the SoC was built - so while this will often be half-size each, it may not be); OR
- *Shared*: the VPEs share all the TLB's entries (up to a maximum of 64 entries). This is certainly a good choice if one of the VPEs doesn't use the TLB at all (which is not unusual in many legacy embedded systems.)

But it is also particularly likely to be a good choice if the VPEs are to run the *same* software and that software is under your control; for example, if you're using them to run a close approximation to a dual-CPU SMP Linux OS (a *VSMP* system.)

But to share the TLB you will need to make some modifications to the TLB maintenance code, as described below.

To partition the TLB just set *MVPControl[STLB]* zero, and set both VPE's *Config1[MMUSize]* fields to the appropriate size (the split is configurable by your SoC designer).

To share the TLB set *MVPControl[STLB]* to 1. It will probably be convenient to set *Config1[MMUSize]* to show the full array. A change to *STLB* should be made only by "unmapped" code, and with the TLB empty of valid entries.

You don't usually need to make any change to the critical TLB refill exception handler, so long as - as is usual - it relies on random replacement (that is, the update to the TLB is done with a **tlbwr** instruction.) The TLB CP0 registers used in a typical TLB miss handler include *Context*, *EntryHi*, *EntryLo0-1* and *PageMask*. All are replicated for each VPE.

The *Random* register is handled specially. **tlbwr** will not use a value for *Random* which coincides with the other VPE's value of the manually-set TLB index register *Index*.

Meanwhile, other kernel software may be doing software-driven updates to the TLB (mostly that means removing entries). The TLB maintenance software will need to run single-threaded, at least in part. There are three possible sources of unwanted concurrency, and software has to attend to two of them:

1. The other VPE may itself be performing some TLB maintenance. This can be fixed with a one-thread-at-a-time software semaphore, just like the ones you use in an SMP OS.
2. Another TC belonging to the same VPE may get a TLB-related exception. This can be fixed by bracketing critical parts of the TLB maintenance routine with a **dmt/mtc0** pair, disabling all TC-level parallelism while the operation is completed.
3. A TC belonging to the other VPE may get a TLB-related exception. But the hardware makes this OK. The only resources the two VPEs share are the TLB entries itself, and the only entry the other VPE will access is the one used by **tlbwr** as indexed by the *Random* register. The hardware will ensure that the *Random* value used by the other VPE will be different from the *Index* value you're using for your maintenance routine. So no software fix is needed.

For efficient use of TLB entries, maintenance software should return *Index* to an unused value (which represents no entry) as soon as it has finished - otherwise you're blocking *Random* from selecting some particular TLB slot. The recommended value is 0x8000.0000; the top bit of *Index* is writable on MIPS MT CPUs for this purpose.

There's another more subtle change. The TLB is used early in the pipeline to translate instruction addresses. A TLB-related exception ("TLB Invalid" for example) detected at this stage is not taken until and unless the instructions are scheduled into the main pipeline. By that time many instructions from other threads may have gone past, and perhaps one of them may have done a refill which dislodged the invalid TLB entry. So the TLB invalid exception handler

might find there's no translation entry in the TLB at all: your best bet, in that case, is probably to just return from the exception without doing anything, which will lead to a TLB miss exception and all should get fixed up.

### 4.3 Setting up a TC to run a thread

The easiest way to set a previously-unoccupied TC running a thread is to use the **fork** instruction.

To prepare to use **fork** you need to make sure there is at least one TC with the *TCStatus[DA]* bit set to 1 (which indicates it's available for fork), the bit *TCStatus[A]* zero (i.e. the TC is not already in use), and *TCHalt* zero.

However, there's nothing to prevent you from setting up a TC manually; just set the thread restart address *TCRestart* and set *TCStatus[A]*. You should almost always set *TCHalt* before doing manual adjustments on a TC, and clear it when you've finished. If the TC started to run (perhaps an interrupt routine) while being worked on it would be likely to lead to confusion.

### 4.4 TCs recycled as Shadow registers

The MIPS32 architecture permits CPUs to be configured such that a particular interrupt handler (or in some cases all exception handlers) should be invoked with a complete alternate set of general-purpose registers: a *Shadow register set*. That allows you to write a very low-overhead handler, because you don't have to save the interrupted thread's registers.

There are some applications where explicit multithreading will fix your problem better than shadow registers. But there are other cases where you really want shadow registers rather than multiple TCs, and the 34K core gives you the choice - you can close down a TC for thread business, and make its registers available for shadow set use. See [Section 7.3.1, "Recycling multi-threading CPU's TCs as shadow sets"](#).





## The MIPS32® DSP ASE

The MIPS DSP ASE is provided to accelerate a large range of DSP algorithms. You can get most programming information from this chapter. There's more detail in the formal DSP ASE specification [MIPSDSP], but expect to read through lots of material aimed at hardware implementors. You may also find [DSPWP] useful for tips and examples of converting DSP algorithms for the DSP ASE.

Different target applications generally need different data size and precision:

- *32-bit data*: audio (non-hand-held) decoding/encoding - a wide range of "hi-fi" standards for consumer audio or television sound.

Raw audio data (as found on CD) is 16-bit; but if you do your processing in 16 bits you lose precision beyond what is acceptable for hi-fi.

- *16-bit data*: digital voice for telephony. International telephony code/decode standards include G.723.1 (8Ksample/s, 5-6Kbit/s data rate, 37ms delay), G.729 (8Kbit/s, 15ms delay) and G.726 (16-40Kbit/s, computationally simpler and higher quality, good for carrying analogue modem tones). Application-specific filters are used for echo cancellation, noise cancellation, and channel equalization.

Also used for soft modems and much general "DSP" work (filters, correlation, convolution); lo-fi devices use 16 bits for audio.

- *8-bit data*: processing of printer images, JPEG (still) images and video data.

### 5.1 Features provided by the MIPS® DSP ASE

Those target applications can benefit from unconventional architecture features because they rely on:

- *Fixed-point fractional data types*: It is not yet economical (in terms of either chip size or power budget) to use floating point calculations in these contexts. DSP applications use fixed-point fractions. Such a fraction is just a signed integer, but understood to represent that integer divided by some power of two. A 32-bit fractional format where the implicit divisor is  $2^{16}$  (65536) would be referred to as a Q15.16 format; that's because there are 16 bits devoted to fractional precision and 15 bits to the whole number range (the highest bit does duty as a sign bit and isn't counted).

With this notation Q31.0 is a conventional signed integer, and Q0.31 is a fraction representing numbers between -1 and 1 (well, nearly 1). It turns out that Q0.31 is the most popular 32-bit format for DSP applications, since it won't overflow when multiplied (except in the corner case where  $-1 \times -1$  leads to the just-too-large value 1). Q0.31 is often abbreviated to Q31.

The DSP ASE provides support for Q31 and Q15 (signed 16-bit) fractions.

- *Saturating arithmetic*: It's not practicable to build in overflow checks to DSP algorithms - they need to be too fast. Clever algorithms may be built to be overflow-proof; but not all can be. Often the least worst thing to do

when a calculation overflows is to make the result the most positive or most negative representable value. Arithmetic which does that is called *saturating* - and quite a lot of operations in the DSP ASE saturate (in many cases there are saturating and non-saturating versions of what is otherwise the same instruction).

- *Multiplying fractions*: if you multiply two Q31 fractions by re-using a full-precision integer multiplier, then you'll get a 64-bit result which consists of a Q62 result with (in the very highest bit) a second copy of the sign bit. This is a bit peculiar, so it's more useful if you always do a left-shift-by-1 on this value, producing a Q63 format (a more natural way to use 64 bits). Q15 multiplies which generate a Q31 value have to do the shift-left too. That's what all the **mulq**... instructions do.
- *Rounding*: some fractional operations implicitly discard less significant bits. But you get a better approximation if you bump the truncated result by one when the discarded bits represent more than a half of the value of a 1 in the new LS position. That's what we mean by *rounding* in this chapter.
- *Multiply-accumulate sequences with choice of four accumulators*: (with fixed-point types, sometimes saturating).

The 34K already has quite a slick integer multiply-accumulate operation, but it's not so efficient when used for fractional and saturating operations.

The sequences are made more usable by having four 64-bit result/accumulator registers - (the old MIPS multiply divide unit has just one, accessible as the *hi/lo* registers). The new *ac0* is the old *hi/lo*, for backward compatibility.

- *Benefit from "SIMD" operations*.: Many DSP calculations are a good match for "Single Instruction Multiple Data" or *vector* operations, where the same arithmetic operation is applied in parallel to several sets of operands.

In the MIPS DSP ASE, some operations are SIMD type - two 16-bit operations or four 8-bit operations are carried out in parallel on operands packed into a single 32-bit general-purpose register. Instructions operating on vectors can be recognized because the name includes **.ph** (paired-half, usually signed, often fractional) or **.qb** (quad-byte, always unsigned, only occasionally fractional).

The DSP ASE hardware involves an extensive re-work of the normal integer multiply/divide unit. As mentioned above it has four 64-bit accumulators (not just one) and a new control register, described immediately below.

## 5.2 The DSP ASE control register

This is a part of the user-mode programming model for the DSP ASE, and is a 32-bit value read and written with the **rddsp/wrdsp** instructions. It holds state information for some DSP sequences.

**Figure 5-1 Fields in the DSPControl Register**

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

In Figure 5-1:

*ccond*: condition bits set by compare instructions (there have to be four to report on compares between vector types). "Compare" operations on scalars or vectors of length two only touch the lower-numbered bits. *DSPControl* bits 31:28 are used for more *ccond* bits in 64-bit machines.

*ouflag*: one of these bits may be set when a result overflows (whether or not the result is saturated depends on the instruction - the flag is set in either case). The "ou" stands for "overflow/underflow" - "underflow" is used here for a value which is negative but with excessive absolute value.

Any overflowed/underflowed result produced by any DSP ASE instruction sets a *ouflag* bit, *except* for **addsc**/**addwc** and **shilo**/**shilov**.

The 6 bits are set according to the destination of the operation which overflowed, and the kind of operation it was:

<i>Bit No</i>	<i>Overflowed destination/instruction</i>
16-19	Destination register is a multiply unit accumulator: separate bits are respectively for accumulators 0-3.
20	Add/subtract.
21	Multiplication of some kind.
22	Shift left or conversion to smaller type
23	Accumulator shift-then-extract

*EFI*: set by any of the accumulator-to-register bitfield extract instructions **extp**, **extpv**, **extpdp**, or **extpdpv**. It's set to 1 if and only if the instruction finds there are insufficient bits to extract. That is, if *DSPControl[pos]* - which is supposed to mark the highest-numbered bit of the field we're extracting - is less than the size value specified by the instruction.

c: Carry bit for 32-bit add/carry instructions **addsc** and **addwc**.

*scount*, *pos*: Fields for use by "variable" bitfield insert and extract instructions, such as **insv** (the normal MIPS32 **ins/ext** instructions have the field size and position hard-coded in the instruction).

*scount* specifies the size of the bit field to be inserted, while *pos* specifies the insert position.

**Caution**: in all inserts (following the lead of the standard MIPS32 insert/extract instructions) *pos* is set to the lowest bit number in the field. But in the DSP ASE extract-from-accumulator instructions (**extp**, **extpv**, **extpdp** and **extpdpv**), *pos* identifies the *highest*-numbered bit in the field.

The latter two ("dp") instructions post-decrement *pos* (by the bitfield length *size*), to help software which is unpacking a series of bitfields from a dense data structure.

The **mthlip** instruction will increment the *pos* value by 32 after copying the value of *lo* to *hi*.

### 5.2.1 DSP accumulators

Whereas a standard MIPS32 architecture CPU has just one 64-bit multiply unit accumulator (accessible as *hi/lo*), the DSP ASE provides three 64-bit accumulators. Instructions accessing the extra accumulators specify a 2-bit field as 0-3 (0 selects the original accumulator).

## 5.3 Software detection of the DSP ASE

You can find out if your core supports the DSP ASE by testing the *Config3[DDSP]* bit (see [Table C-6](#)).

Then you need to enable the instruction set by setting *Status[MX]* (or its alternate view *TCStatus[TMX]*) to 1, for any TC which will execute code in the MIPS DSP ASE.

## 5.4 DSP instructions

The DSP instruction set is nothing like the regular and orthogonal MIPS32 instruction set. It's a collection of special-case instructions, in many cases aimed at the known hot-spots of important algorithms.

We'll summarize the instructions under headings, but then list all of them in [Section 5.2, "DSP instructions in alphabetical order"](#), an alphabetically-ordered list which provides a terse but usually-sufficient description of what each instruction does.

### 5.4.1 Hints in instruction names

An instruction's name may have some suffixes which are often informative:

- q**: generally means it treats operands as fractions (which isn't important for adds and subtracts, but is important for multiplications and convert operations);
- \_s**: usually means the full-precision result is saturated to the size of the destination; **\_sa** is used for instructions which saturate intermediate results before accumulating; and **r**: denotes rounding (see above);
- .w**, **.ph**, **.qb**: suggest the operation is dealing with 32-bit, paired-half or quad-byte values respectively. Where there are two of these (as in **macq\_s.w.ph1**) the first one suggests the type of the result, and the second the type of the operand(s).
- v**: (in a shift instruction) suggests that the shift amount is defined in a register, rather than being encoded in a field of the instruction.

To help you get your arms around this collection of instructions we'll group them by likely usage - guided by the type of the result performed, with an eye to the application. The multiplication instructions are more tricky: most of them have multiple uses. We've sorted them by the most obvious use (likely also the most common). The classification we've chosen divides them into:

- [Arithmetic - 64-bit](#)
- [Arithmetic - saturating and/or SIMD Types](#)
- [Bit-shifts - saturating and/or SIMD types](#)
- [Comparison and "conditional-move" operations on SIMD types - includes \*\*pick\*\* instructions.](#)
- [Conversions to and from SIMD types](#)
- [Multiplication - SIMD types with result in GP register](#)
- [Multiply Q15s from paired-half and accumulate](#)
- [Load with register+register address](#)
- [DSPControl register access](#)
- [Accumulator access instructions](#)
- [Dot products and building blocks for complex multiplication - includes full-word \(Q31\) multiply-accumulate](#)
- [Other DSP ASE instructions - everything else...](#)

## 5.4.2 Arithmetic - 64-bit

**addsc/addwc** generate and use a carry bit, for efficient 64-bit add.

## 5.4.3 Arithmetic - saturating and/or SIMD Types

- *32-bit signed saturating arithmetic:* **addq\_s.w**, **subq\_s.w** and **absq\_s.w**.
- *Paired-half and quad-byte SIMD arithmetic:* perform the same operation simultaneously on both 16-bit halves or all four 8-bit bytes of a 32-bit register. The "q" in the instruction mnemonic for the PH operations here is cosmetic: Q15 and signed 16-bit integer add/subtract operations are bit-identical - Q15 only behaves very differently when converted or multiplied.

The paired half operations are: **addq.ph/addq\_s.ph**, **subq.ph/subq\_s.ph** and **absq\_s.ph**.

The quad-byte operations (all unsigned) are: **addu.qb/addu\_s.qb**, **subu.qb/subu\_s.qb**.

- *Sum of quad-byte vector:* **raddu.w.qb** does an unsigned sum of the four bytes found in a register, zero extends the result and delivers it as a 32-bit value.

## 5.4.4 Bit-shifts - saturating and/or SIMD types

All shifts can either have a shift amount encoded in the instruction, or - indicated by a trailing "v" in the instruction name - provided as a register operand. PH and 32-bit shifts have optional forms which saturate the result.

- *32-bit signed shifts:* include a saturating version of shift left, **shll\_s.w**; and an auto-rounded shift right (just the "arithmetic", sign-propagating form): **shra\_r.w**. Recall from above that rounding can be imagined as pre-adding a half to the least significant surviving bit.
- *Paired-half and quad-byte SIMD shifts:* **shll.ph/shllv.ph/shll\_s.ph/shllv\_s** are as above. For PH only there's a shift-right-arithmetic instruction ("arithmetic" means it propagates the sign bit downward) **shra.ph**, which has a variant which rounds the result **shra\_r.ph**.

The quad-byte shifts are unsigned and don't round or saturate: **shll.qb/shllv.qb**, **shrl.qb/shrlv.qb**.

## 5.4.5 Comparison and "conditional-move" operations on SIMD types

The "cmp" operations simultaneously compare and set flags for two or four values packed in a vector (with equality, less-than and less-than-or-equal tests). For PH that's **cmp.eq.ph**, **cmp.lt.ph** and **cmp.le.ph**. The result is left in the two LS bits of *DSPControl[ccond]*.

For quad-byte values **cmpu.eq.qb**, **cmpu.lt.qb** and **cmpu.le.qb** simultaneously compare and set flags for four bytes in *DSPControl[ccond]* - the flag relating to the bytes found in the low-order bits of the source register is in the lowest-numbered bit (and so on). There's an alternative set of instructions **cmpgu.eq.qb**, **cmpgu.lt.qb** and **cmpgu.le.qb** which leave the 4-bit result in a specified general-purpose register.

**pick.ph** uses the two LS bits of *DSPControl[ccond]* (usually the outcome of a paired-half compare instruction, see above) to determine whether corresponding halves of the result should come from the first or second source register. Among other things, this can implement a paired-half conditional move. You can reverse the order of your conditional inputs to do a move dependent on the complementary condition, too.

**pick.qb** does the same for QB types, this time using four bits of *DSPControl[ccond]*.

## 5.4.6 Conversions to and from SIMD types

Conversion operations from larger to smaller fractional types have names which start "**precrq...**" for "precision reduction, fractional". Conversion operations from smaller to larger have names which start "**prece...**" for "precision expansion".

- *Form vector from high/low parts of two other paired-half values:* **packr1.ph** makes a paired-half vector from two half vectors, swapping the position of each sub-vector. It can be used to acquire a properly formed sub-vector from a non-aligned data stream.
- *One Q15 from a paired-half to a Q31 value:* **preceq.w.phl/preceq.w.phr** select respectively the "left" (high bit numbered) or "right" (low bit numbered) Q15 value from a paired-half register, and load it into the result register as a Q31 (that is, it's put in the high 16 bits and the low 15 bits are zeroed).
- *Two bytes from a quad-byte to paired-half:* **precequ.ph.qbl/precequ.ph.qbr** picks two bytes from either the "left" (high bit numbered) or "right" (low bit numbered) halves of a quad-byte value, and unpacks to a pair of Q15 fractions.

**precequ.ph.qbla** does the same, except that it picks two "alternate" bytes from bits 31-24 and 15-8, while **precequ.ph.qbra** picks bytes from bits 23-16 and 7-0.

Similar instructions without the **q** - **preceu.ph.qbl**, **preceu.ph.qbr**, **preceu.ph.qbla** and **preceu.ph.qbra** - work on the same register fields, but treat the quantities as integers, so the 16-bit results get their low bits set.

- *2×Q31 to a paired-half:* both operands and result are assumed to be signed fractions, so **precrq.ph.w** just takes the high halves of the two source operands and packs them into a paired-half; **precrq.rs.ph.w** rounds and saturates the results to Q15.
- *2×paired-half to quad-byte:* you need two source registers to provide four paired-half values, of course. This is a fractional operation, so it's the low bits of the 16-bit fractions which are discarded.

**precrq.qb.ph** treats the paired-half operands as unsigned fractions, retaining just the 8 high bits of each 16-bit component.

**precrqu.s.qb.ph** treats the paired-half operands as Q15 signed fractions and both rounds and saturates the result (in particular, a negative Q15 fraction produces a zero byte, since zero is the lowest representable quantity).

- *Replicate immediate or register value to paired-half:* in **repl.ph** the value to be replicated is a 10-bit signed immediate value (that's in the range  $-512 \leq x \leq 511$ ) which is sign-extended to 16 bits, whereas in **replv.ph** the value - assumed to be already a Q15 value - is in a register.
- *Replicate single value to quad-byte:* there's both a register-to-register form **replv.qb** and an immediate form **repl.qb**.

## 5.4.7 Multiplication - SIMD types with result in GP register

When a multiply's destination is a general-purpose register, the operation is still done in the multiply unit, and you should expect it to overwrite the *hi/lo* registers (otherwise known as *ac0*.)

- *8-bit×16-bit 2-way SIMD multiplication:* **muleu.s.ph.qbl/muleu.s.ph.qbr** picks the "left" (high bit numbered) or "right" (low bit numbered) pair of byte values from one source register and a pair of 16-bit values

from the other. Two unsigned integer multiplications are done at once, the results unsigned-saturated and delivered to the two 16-bit halves of the destination.

The asymmetric use of the source operands is not a bit like a Q15 operation. But 8×16 multiplies are heavily used in imaging and video processing (JPEG image encode/decode, for example).

- *Paired-half SIMD multiplication:* **mulq\_rs.ph** multiplies two Q15s at once and delivers it to a paired-half value in a general-purpose register, with rounding and saturation.
- *Multiply half-PH operands to a Q31 result:* **muleq\_s.w.phl/muleq\_s.w.phr** pick the "left"/"right" Q15 value respectively from each operand, multiply and store a Q31 value.

"Precision-doubling" multiplications like this *can* overflow, but only in the extreme case where you multiply -1×-1, and can't represent 1 exactly.

### 5.4.8 Multiply Q15s from paired-half and accumulate

**maq\_s.w.phl/maq\_s.w.phr** picks either the left/high or right/low Q15 value from each operand, multiplies them to Q31 and accumulates to a Q32.31 result. The multiply is saturated only when it's -1×-1.

**maq\_sa.w.phl/maq\_sa.w.phr** differ in that the final result is saturated to a Q31 value held in the low half of the accumulator (required by some ITU voice encoding standards).

### 5.4.9 Load with register + register address

Previously available only for floating point data<sup>1</sup>: **lwx** for 32-bit loads, **lhx** for 16-bit loads (sign-extended) and **lbux** for 8-bit loads, zero-extended.

### 5.4.10 DSPControl register access

**wrdsp rs,mask** sets *DSPControl* fields, but only those fields which are enabled by a 1 bit in the 6-bit mask.

**rddsp** reads *DSPControl* into a GPR; but again it takes a mask field. Bitfields in the GPR corresponding to *DSPControl* fields which are not enabled will be set all-zero.

The mask bits tie up with fields like this:

**Table 5.1 Mask bits for instructions accessing the DSPControl register**

<i>Mask Bit</i>	<i>DSPControl field</i>
0	pos
1	scount
2	c
3	ouflag
4	ccond
5	EFI

1. Well, an integer instruction is also included in the MIPS SmartMIPS™ ASE.

### 5.4.11 Accumulator access instructions

- *Historical instructions which now access new accumulators:* the familiar **mfhi/mflo/mthi/mtlo** instructions now take an optional extra accumulator-number parameter.
- *Shift and move to general register:* **extr.w/extr\_r.w/extr\_rs.w** gets a 32-bit field from an accumulator (starting at bit 0 up to 31) and puts the value in a general purpose register. At your option you can specify rounding and signed 32-bit saturation.

**extrv.w/extrv\_r.w/extrv\_rs.w** do the same but specify the field's starting bit number with a register.

- *Extract bitfield from accumulator:* **extp/extpv** takes a bitfield (up to 32 bits) from an accumulator and moves it to a GPR. The length of the field can be an immediate value or from a register. The position of the field is determined by *DSPControl[pos]*, which holds the bit number of the most significant bit.

**extpdp/extpdpv** do the same, but also auto-decrement *DSPControl[pos]* to the bit-number just below the field you extracted.

- *Accumulator rearrangement:* **shilo/shilov** has a *signed* shift value between -32 and +31, where positive numbers shift right, and negative ones shift left. The "v" version, as usual, takes the shift value from a register. The right shift is a "logical" type so the result is zero extended.
- *Fill accumulator pushing low half to high:* **mthlip** moves the low half of the accumulator to the high half, then writes the GPR value in the low half. Generally used to bring 32 more bits from a bitstream into the accumulator for parsing by the various **ext...** instructions.

### 5.4.12 Dot products and building blocks for complex multiplication

In 2-dimensional vector math (or in any doubled-up step of a multiply-accumulate sequence which has been optimized for 2-way SIMD) you're often interested in the *dot product* of two vectors:

$$v[0]*w[0] + v[1]*w[1]$$

In many cases you take the dot product of a series of vectors and add it up, too.

Some algorithms use complex numbers, represented by 2D vectors. Complex numbers use *i* to stand for "the square root of -1", and a vector  $[a, b]$  is interpreted as  $a + ib$  (mathematicians leave out the multiply sign and use single-letter variables, habits which would not be appreciated in C programming!) Complex multiplication just follows the rules of multiplying out sums, remembering that  $i*i = -1$ , so:

$$(a + ib)*(c + id) = (a*c - b*d) + i(a*d + b*c)$$

Or in vector format:

$$[a, b] * [c, d] = [a*c - b*d, a*d + b*c]$$

The first element of the result (the "real component") is like a dot product but with a subtraction, and the second (the "imaginary component") is like a dot product but with the vectors crossed.



- *Q15 dot product from paired-half, and accumulate:* **dpag\_s.w.ph** does a SIMD multiply of the Q15 halves of the operands, then adds the results and saturates to form a Q31 fraction, which is accumulated into a Q32.31 fraction in the accumulator.

**dpsq\_s.w.ph** does the same but subtracts the dot product from the accumulator.

For the imaginary component of a complex multiply, first swap the Q15 numbers in one of the register operands with a **rot** (bit-rotate) instruction.

For the real component of a complex Q15 multiply, you have the difference-of-products instruction **mulsaq\_s.w.ph**, which parallel-multiplies both Q15 halves of the PH operands, then computes the difference of the two results and leaves it in an accumulator in Q32.31 format (beware: this does not accumulate the result).

- *16-bit integer dot-product from paired-half, and accumulate:* **dpau.h.qb1/dpau.h.qbr** picks two QB values from each source register, parallel-multiplies the corresponding pairs to integer 16-bit values, adds them together and then adds the whole lot into an accumulator. **dpsu.h.qb1/dpsu.h.qbr** do the same sum-of-products, but the result is then subtracted from the accumulator. In both cases, note this is integer (not fractional) arithmetic.
- *Q31 saturated multiply-accumulate:* is the nearest thing you can get to a dot-product for Q31 values. **dpag\_sa.1.w** does a Q31 multiplication and saturates to produce a Q63 result, which is added to the accumulator and saturated again. **dpsq\_sa.1.w** does the same, except that the multiply result is subtracted from the accumulator (again, useful for the real component of a complex number).

#### 5.4.13 Other DSP ASE instructions

- *Branch on DSPControl field:* **bposge32** branches if *DSPControl[pos] ≥ 32*.

Typically the test is for "is it time to load another 32 bits of data from the bitstream yet?".

- *Circular buffer index update:* **modsub** takes an operand which packs both a maximum index value and an index step, and uses it to decrement a "buffer index" by the step value, but arranging to step from zero to the provided maximum.
- *Bitfield insert with variable size/position:* **insv** is a bit-insert instruction. It acts like the MIPS32 standard instruction **ins** except that the position and size of the inserted field are specified not as immediates inside the instruction, but are obtained from *DSPControl[pos]* (which should be set to the lowest numbered bit of the field you want) and *DSPControl[scount]* respectively.
- *Bit-order reversal:* **bitrev** reverses the bits in the low 16 bits of the register. The high half of the destination is zero.

The bit-reverse operation is a computationally crucial step in buffer management for FFT algorithms, and a 16-bit operation supports up to a 32K-point FFT, which is much more than enough. A full 32-bit reversal would be expensive and slow.

## 5.5 Macros and typedefs for DSP instructions

It's useful to be able to use fragments of C code to describe what some instructions do. To do that, we need to be able to refer to fractional types, saturation and vectors. Here are the definitions we're using<sup>1</sup>:

```
typedef long long int64;
typedef int int32;

/* accumulator type */
typedef signed long long q32_31;

typedef signed int q31;

#define MAX31 0x7FFFFFFF
#define MIN31 -(1<<31)
#define SAT31(x) (x > MAX31 ? MAX31: x < MIN31 ? MIN31: x)

typedef signed short q15;
#define MAX15 0x7FFF
#define MIN15 -(1<<15)
#define SAT15(x) (x > MAX15 ? MAX15: x < MIN15 ? MIN15: x)

typedef unsigned char u8;
#define MAXUBYTE 255
#define SATUBYTE(x) (x > MAXUBYTE ? MAXUBYTE: x < 0 ? 0: x)

/* fields in the vector types are specified by relative bit
   position, but C definitions are in memory order, so these
   definitions need to be endianness-dependent */

#ifdef BIG_ENDIAN
typedef struct{
    q15 h1, h0;
} ph;

typedef struct{
    u8 b3, b2, b1, b0;
} qb;
#else
typedef struct{
    q15 h0, h1;
} ph;

typedef struct{
    u8 b0, b1, b2, b3;
} qb;
#endif
```

---

1. This page needs more work, and I hope it will be improved in a future version of the manual.

## 5.6 Almost Alphabetically-ordered table of DSP ASE instructions

Table 5.2 DSP instructions in alphabetical order

<i>Instruction</i>	<i>Description</i>
<b>absq_s.w rd,rt</b>	Q31/signed integer absolute value with saturation
<b>addq.ph rd,rs,rt</b> <b>addq_s.ph rd,rs,rt</b>	2×SIMD Q15 addition, without and with saturation of the result
<b>addq_s.w rd,rs,rt</b>	Q31/signed integer addition with saturation
<b>addsc rd,rs,rt</b> <b>addwc rd,rs,rt</b>	Add setting carry, then add with carry. The carry bit is kept in <i>DSPControl[c]</i> . So to add the 64-bit values in registers <i>yhi/ylo, zhi/zlo</i> to produce a 64-bit value in <i>xhi/xlo</i> , just do: addsc xlo, ylo, zlo; addwc xhi, yhi, zhi
<b>addu.qb rd,rs,rt</b> <b>addu_s.qb rd,rs,rt</b>	4×SIMD QBYTE addition, without and with SATUBYTE saturation.
<b>bitrev rd,rt</b>	Delivers the bit-reversal of the low 16 bits of the input (result has high half zero).
<b>bposge32 offset</b>	Branch if <i>DSPControl[pos]&gt;=32</i> . Like most branch instruction, it has a 16-bit "PC-relative" target encoding.
<b>cmp.eq.ph rs,rt</b> <b>cmp.le.ph rs,rt</b> <b>cmp.lt.ph rs,rt</b>	Signed compare of both halves of two paired-half ("PH") values. Results are written into <i>DSPControl[cond1-0]</i> for high and low halves respectively (1 for true, 0 for false). A signed compare works for both Q15 or signed 16-bit values.
<b>cmpgu.eq.qb rd,rs,rt</b> <b>cmpgu.le.qb rd,rs,rt</b> <b>cmpgu.lt.qb rd,rs,rt</b>	Unsigned simultaneous compare of all four bytes in quad-byte values. The four result bits are written into the four LS bits of general register <i>rd</i> .
<b>cmpu.eq.qb rs,rt</b> <b>cmpu.le.qb rs,rt</b> <b>cmpu.lt.qb rs,rt</b>	Unsigned simultaneous compare of all four bytes in quad-byte values. The four result bits are written into register <i>DSPControl[cond3-0]</i> .
<b>dpaq_s.w.ph ac,rs,rt</b>	"Dot product and accumulate", with Q31 saturation of each multiply result: ph rs,rt; ac += SAT31(rs.h0*rt.h0 + rs.h1*rt.h1); The accumulator is effectively used as a Q32.31 fraction.
<b>dpaq_sa.l.w ac,rs,rt</b>	Q31 saturated multiply-accumulate
<b>dpau.h.qb1</b>  <b>dpau.h.qbr</b>	qb rs, rt; ac += rs.b3*rt.b3 + rs.b2*rt.b2; Dot-product and accumulate of quad-byte values ("1" for left, because these are the higher bit-numbered bytes in the 32-bit register). Not a fractional computation, just unsigned 8-bit integers. Then for the lower bit-numbered bytes: qb rs, rt; ac += rs.b1*rt.b1 + rs.b0*rt.b0;
<b>dpsq_s.w.ph ac,rs,rt</b>	Paired-half fractional "dot product and subtract from accumulator" ph rs, rt; q32_31 ac; ac -= SAT31(rs.h1*rt.h1 + rs.h0*rt.h0);
<b>dpsq_sa.l.w ac,rs,rt</b>	Q31 saturated fractional-multiply, then subtract from accumulator: q31 rs, rt; q32_31 ac; ac -= SAT31(rs*rt);
<b>dpsu.h.qb1 ac,rs,rt</b>  <b>dpsu.h.qbr ac,rs,rt</b>	QB format dot-product and subtract from accumulator. This is an integer (not fractional) multiplication and comes in "left" and "right" (higher/lower-bit numbered pair) versions: qb rs,rt; ac -= rs.b3*rt.b3 + rs.b2*rt.b2; qb rs,rt; ac -= rs.b1*rt.b1 + rs.b0*rt.b0;

Table 5.2 DSP instructions in alphabetical order

<i>Instruction</i>	<i>Description</i>
<b>extp</b> <i>rt,ac,size</i> <b>extpdp</b> <i>rt,ac,size</i> <b>extpdpv</b> <i>rt,ac,rs</i> <b>extpv</b> <i>rt,ac,rs</i>	Extract bitfield from an accumulator to register. The length of the field (number of bits) can be an immediate constant or can be provided by a second source register (in the <b>v</b> variants). The field position, though, comes from <i>DSPControl[pos]</i> , which marks the highest-numbered bit of the field (note that the MIPS32 standard bitfield extract instructions specify the <i>lowest</i> bit number in the field). In the <b>dp</b> variants like <b>extpdp/extpdpv</b> , <i>DSPControl[pos]</i> is auto-decremented by the length of the field extracted, which is useful when unpacking the accumulator into a series of fields.
<b>extr.w</b> <i>rt,ac,shift</i> <b>extr_r.w</b> <i>rt,ac,shift</i> <b>extr_rs.w</b> <i>rt,ac,shift</i> <b>extrv.w</b> <i>rt,ac,rs</i> <b>extrv_r.w</b> <i>rt,ac,rs</i> <b>extrv_rs.w</b> <i>rt,ac,rs</i>	Extracts a bit field from an accumulator into a general purpose register. The LS bit of the extracted field can start anywhere from bit zero to 31 of the accumulator: <code>int64 ac; unsigned int rt;</code> <code>rt = (ac &gt;&gt; shift) &amp; 0xFFFFFFFF;</code> At option you can specify rounding ( <b>_r</b> names): <code>int64 ac; unsigned int rt;</code> <code>rt = ((ac + 1&lt;&lt;(shift-1)) &gt;&gt; shift) &amp; 0xFFFFFFFF;</code> and signed 32-bit saturation of the result ( <b>_s/_rs</b> names). The <b>extrv...</b> variants specify the shift amount (still limited to 31 positions) with a register.
<b>extr_s.h</b> <i>rt,ac,shift</i> <b>extrv_s.h</b> <i>rt,ac,rs</i>	Obtain a right-shifted value from an accumulator and form a signed 16-bit saturated result.
<b>insv</b> <i>rt,rs</i>	The bitfield insert in the standard MIPS32 instruction set is <b>ins</b> <i>rt,rs,pos,size</i> , and the position and size must be constants (encoded as immediates in the instruction itself). This instruction permits the position and size to be calculated by the program, and then supplied as <i>DSPControl[pos]</i> and <i>DSPControl[scount]</i> respectively. In this case <i>DSPControl[pos]</i> must be set to the <i>lowest</i> numbered bit in the field to be inserted: yes, that's different from the <b>extp...</b> instructions.
<b>lbux</b> <i>rd,index(base)</i> <b>lhx</b> <i>rd,index(base)</i> <b>lwx</b> <i>rd,index(base)</i>	Load operations with register+register address formation. <b>lbux</b> is a load byte and zero extend, <b>lhx</b> loads half-word and sign-extends, and <b>lwx</b> loads a whole word. The full address must be naturally aligned for the data type.
<b>maq_s.w.phl</b> <i>ac,rs,rt</i> <b>maq_s.w.phr</b> <i>ac,rs,rt</i> <b>maq_sa.w.phl</b> <i>ac,rs,rt</i> <b>maq_sa.w.phr</b> <i>ac,rs,rt</i>	Non-SIMD Q15 multiply-accumulate, with operands coming from either the "left" (higher bit number) or "right" (lower bit number) half of each of the operand registers. In all versions the Q15 multiplication is saturated to a Q31 results. The "_sa" variants saturates the add result in the accumulator to a Q31, too.
<b>mfhi</b> <i>rd,ac</i> <b>mflo</b> <i>rd,ac</i>	Legacy instruction, which now works on new accumulators (if you provide a second nonzero argument). Copies high/low half (respectively) of accumulator to <i>rd</i> .
<b>modsub</b> <i>rd,rs,rt</i>	Circular buffer index update. <i>rt</i> packs both the decrement amount (low 8 bits) and the highest index (high 24 bits), then this instruction calculates: <code>rd = (rs == 0) ? ((unsigned) rt &gt;&gt; 8) : rs - (rt &amp; 0xFF);</code>
<b>mthi</b> <i>rs,ac</i>	Legacy instruction working on new accumulators. Moves data from <i>rd</i> to the high half of an accumulator.
<b>mthlip</b> <i>rs,ac</i>	Moves the low half of the accumulator to the high half, then writes the GPR value in the low half.
<b>mtlo</b> <i>rs,ac</i>	Legacy instruction working on new accumulators. Moves data from <i>rd</i> to the low half of an accumulator.
<b>muleq_s.w.phl</b> <i>rd,rs,rt</i> <b>muleq_s.w.phr</b> <i>rd,rs,rt</i>	Multiply selected Q15 values from "left"/"right" (higher/lower numbered bits) of <i>rd/rs</i> to a Q31 result in a general purpose register, Q31-saturating. Like all multiplies which target general purpose registers, it may well use the multiply unit and overwrite <i>hi/lo</i> , also known as <i>ac0</i> .

Table 5.2 DSP instructions in alphabetical order

<i>Instruction</i>	<i>Description</i>
<b>muleu_s.ph.qbl rd,rs,rt</b> <b>muleu_s.ph.qbr rd,rs,rt</b>	A 2xSIMD 16-bitx8-bit multiplication. <b>muleu_s.ph.qbl</b> does something like: $rd = ((LL\_B(rs) * LEFT\_H(rt)) \ll 16) \mid ((LR\_B(rs) * RIGHT\_H(rt));$ <p>Note that the multiplications are unsigned integer multiplications, and each half of the result is unsigned-16-bit-saturated. The asymmetric source operands are quite unusual, and note this is not a fractional computation. <b>muleu_s.ph.qbr</b> is the same but picks the RL and RR (low bit numbered) byte values from <i>rs</i>.</p>
<b>mulq_rs.ph rd,rs,rt</b>	2xSIMD Q15 multiplication to two Q15 results. Result in general purpose register, <i>hi/lo</i> or <i>ac0</i> may be overwritten.
<b>multsaq_s.w.ph ac,rs,rt</b>	$ac += (LEFT\_H(rs) * LEFT\_H(rt)) - (RIGHT\_H(rs) * RIGHT\_H(rt));$ The multiplications are done to Q31 values, saturated if they overflow (which is only possible when $-1 \neq -1$ makes $+1$ ). The accumulator is really a Q32.31 value, so is unlikely to overflow; no overflow check is done on the accumulation.
<b>packrl.ph rd,rs,rt</b>	pack a “right” and “left” half from different registers, ie $rd = ((rs \& 0xFFFF) \ll 16) \mid (rt \gg 16) \& 0xFFFF);$
<b>pick.ph rd,rs,rt</b>	Like a 2-way SIMD conditional move: <code>ph rd,rs,rt;</code> <code>rd.l = DSPControl[ccond1] ? rs.l: rt.l;</code> <code>rd.r = DSPControl[ccond0] ? rs.r: rt.r;</code>
<b>pick.qb rd,rs,rt</b>	Kind of a 4-way SIMD conditional move: <code>qb rd,rs,rt;</code> <code>rd.ll = DSPControl[ccond3] ? rs.ll: rt.ll;</code> <code>rd.lr = DSPControl[ccond2] ? rs.lr: rt.lr;</code> <code>rd.rl = DSPControl[ccond1] ? rs.rl: rt.rl;</code> <code>rd.rr = DSPControl[ccond0] ? rs.rr: rt.rr;</code>
<b>preceq.w.phl rd,rt</b> <b>preceq.w.phr rd,rt</b>	Convert a Q15 value (either left/high or right/low half of <i>rt</i> ) to a Q31 value in <i>rd</i> .
<b>precequ.ph.qbl rd,rt</b> <b>precequ.ph.qbla rd,rt</b> <b>precequ.ph.qbr rd,rt</b> <b>precequ.ph.qbra rd,rt</b>	Simultaneously convert two unsigned 8-bit fractions from <i>rt</i> to Q15 and load into the two halves of <i>rd</i> . <b>precequ.ph.qbl</b> uses <i>rt.ll/rt.lr</i> ; <b>precequ.ph.qbla</b> uses <i>rt.ll/rt.rl</i> ; <b>precequ.ph.qbr</b> uses <i>rt.rl/rt.rr</i> ; and <b>precequ.ph.qbra</b> uses <i>rt.lr/rt.rr</i> .
<b>preceu.ph.qbl rd,rt</b> <b>preceu.ph.qbla rd,rt</b> <b>preceu.ph.qbr rd,rt</b> <b>preceu.ph.qbra rd,rt</b>	Zero-extend two unsigned byte values from <i>rt</i> to unsigned 16-bit and load into the two halves of <i>rd</i> . <b>preceu.ph.qbl</b> uses <i>rt.ll/rt.lr</i> ; <b>preceu.ph.qbla</b> uses <i>rt.ll/rt.rl</i> ; <b>prececu.ph.qbr</b> uses <i>rt.rl/rt.rr</i> ; and <b>prececu.ph.qbra</b> uses <i>rt.lr/rt.rr</i> .
<b>precrq.ph.w rd,rs,rt</b> <b>precrq_rs.ph.w rd,rs,rt</b>	<b>precrq.ph.w</b> makes a paired-Q15 value by taking the MS bits of the Q31 values in <i>rs</i> and <i>rt</i> , like this: $rd = (rs \& 0xFFFF0000) \mid ((rt \gg 16) \& 0xFFFF);$ <b>precrq_rs.ph.w</b> is the same, but rounds and Q15-saturates both half-results.
<b>precrq.qb.ph rd,rs,rt</b>  <b>precrqu_s.qb.ph rd,rs,rt</b>	Form a quad-byte value from two paired-halves. We use the upper 8 bits of each half-word value, as if we were converting an unsigned 16-bit fraction to an unsigned 8-bit fraction. In C: $rd = (rs \& 0xFF000000) \mid (rs \ll 8 \& 0xFF0000) \mid (rt \gg 16 \& 0xFF00) \mid (rt \gg 8 \& 0xFF);$ Does the same, but each conversion is rounded and saturated to an unsigned byte. Note in particular that a negative Q15 quantity yields a zero byte, since zero is the smallest representable value.
<b>raddu.w.qb rd,rs</b>	Set <i>rd</i> to the unsigned 32-bit integer sum of the four unsigned bytes in <i>rs</i> .

Table 5.2 DSP instructions in alphabetical order

<i>Instruction</i>	<i>Description</i>
<b>rddsp</b> <i>rt,mask</i>	Read the contents of the <i>DSPControl</i> register into <i>rt</i> , but zeroing out any fields for which the appropriate mask bit is zeroed, see <a href="#">Figure 5.1</a> above.
<b>repl.ph</b> <i>rd,imm</i> <b>replv.ph</b> <i>rd,rt</i>	Replicate the same signed value into the two halves of a PH value in <i>rd</i> ; the value is either provided as an immediate whose range is limited between -512 and +511 ( <b>repl.ph</b> ) or from the <i>rt</i> register ( <b>replv.ph</b> ).
<b>repl.qb</b> <i>rd,imm</i> <b>replv.qb</b> <i>rd,rt</i>	Replicate the same 8-bit value into all four parts of a QB value in <i>rd</i> ; the value can come from an immediate constant, or the <i>rt</i> register of the <b>replv.qb</b> instruction.
<b>shilo</b> <i>ac,shift</i> <b>shilov</b> <i>ac,rs</i>	Do a right or left shift (use a negative value for a left shift) of a 64-bit accumulator. The right shift is "logical", bringing in zeroes into the high bits. <b>shilo</b> takes a constant shift amount, while <b>shilov</b> get the shift amount from <i>rs</i> . The shift amount may be no more than 31 right or 32 left.
<b>shll.ph</b> <i>rd, rt, sa</i> <b>shllv.ph</b> <i>rd, rt, rs</i> <b>shll_s.ph</b> <i>rd, rt, sa</i> <b>shllv_s.ph</b> <i>rd, rt, rs</i>	2×SIMD (paired-half) shift left. The " <b>v</b> " versions take the shift amount from a register, and the " <b>_s</b> " versions saturate the result to a signed 16-bit range.
<b>shll.qb</b> <i>rd, rt, sa</i> <b>shllv.qb</b> <i>rd, rt, rs</i>	4×SIMD quad-byte shift left, with shift-amount-in-register and saturating (to an unsigned 8-bit result) versions.
<b>shll_s.w</b> <i>rd, rt, sa</i> <b>shllv_s.w</b> <i>rd, rt, rs</i>	Signed 32-bit shift left with saturation, with shift-amount-in-register <b>shllv_s</b> option.
<b>shra.ph</b> <i>rd, rt, sa</i> <b>shra_r.ph</b> <i>rd, rt, sa</i> <b>shrav.ph</b> <i>rd, rt, rs</i> <b>shrav_r.ph</b> <i>rd, rt, rs</i>	2×SIMD paired-half shift-right arithmetic ("arithmetic" because the vacated high bits of the value are replaced by copies of the input bit 16, the sign bit) - thus performing a correct division by a power of two of a signed number. As usual the <b>shra_v</b> variant has the shift amount specified in a register. The <b>_r</b> versions round the result first (see <a href="#">the bullet on rounding</a> above).
<b>shra_r.w</b> <i>rd, rt, sa</i> <b>shrav_r.w</b> <i>rd, rt, rs</i>	32-bit signed/arithmetic shift right with rounding, see <a href="#">the bullet on rounding</a> .
<b>shrl.qb</b> <i>rd, rt, sa</i> <b>shrlv.qb</b> <i>rd, rt, rs</i>	4×SIMD shift right logical ("logical" means that the vacated high bits are filled with zero, appropriate since the byte quantities in a quad-byte are usually treated as unsigned.)
<b>subq.ph</b> <i>rd,rs,rt</i> <b>subq_s.ph</b> <i>rd,rs,rt</i>	2×SIMD subtraction. <b>subq_s.ph</b> saturates its results to a signed 16-bit range.
<b>subq_s.w</b> <i>rd,rs,rt</i>	32-bit saturating subtraction.
<b>subu.qb</b> <i>rd,rs,rt</i> <b>subu_s.qb</b> <i>rd,rs,rt</i>	4×SIMD quad-byte subtraction. Since quad-bytes are treated as unsigned, the saturating variant <b>subu_s.qb</b> works to an unsigned byte range.
<b>wrdsp</b> <i>rt,mask</i>	Write the <i>DSPControl</i> register with data from <i>rt</i> , but leaving unchanged any fields for which the appropriate mask bit is zeroed, see <a href="#">Figure 5.1</a> above.

## 5.7 DSP ASE instruction timing

Most DSP ASE operations are pipelined, and instructions can often be issued at the maximum CPU rate, but getting results back into the general-purpose register file takes a few clocks. The timings are generally fairly similar to those for the standard multiply instructions, and are listed - together with delays for the standard instruction set - in [Section 9.5.4, "Data dependency delays classified"](#).

## Memory map, caching, reads and writes and translation

In this chapter:

- Section 6.1, "The memory map": basic memory map of the system.
- Section 6.3, "Reads, writes and synchronization"
- Section 6.4, "Caches"
- Section 6.5, "Scratchpad memory/SPRAM": optional on-chip, high-speed memory (particularly useful when dual-ported to the OCP interface).
- Section 6.6, "The TLB and translation": how translation is done and supporting CP0 registers.

### 6.1 The memory map

A 34K core system can be configured with either a TLB (virtual memory translation unit) or a fixed memory mapping, or even with one VPE using the TLB and one with fixed mapping.

A TLB-equipped VPE sees the memory map described by the [MIPS32] architecture, which will be familiar to anyone who has used a 32-bit MIPS architecture CPU and is summarized in Figure 6.1. The TLB gives you access to a full 32-bits physical address on the system interface. More information about the TLB in Section 6.6, "The TLB and translation".

**Table 6.1 Basic MIPS32® architecture memory map**

<i>Segment Name</i>	<i>Virtual range</i>	<i>What happens to accesses here?</i>
kuseg	0x0000.0000-0x7FFF.FFFF	The only region accessible to user-privilege programs. Mapped by TLB entries.
kseg0	0x8000.0000-0x9FFF.FFFF	a fixed-mapping window onto physical addresses 0x0000.0000-0x1FFF.FFFF. Almost invariably cacheable - but in fact other choices are available, and are selected by <i>Config[K0]</i> , see Section C-4, "Fields in the Config register". Accessible only to kernel-privilege programs.
kseg1	0xA000.0000-0xBFFF.FFFF	a fixed-mapping window onto the same physical address range 0x0000.0000-0x1FFF.FFFF as "kseg0" - but accesses here are uncached. Accessible only to kernel-privilege programs.
kseg2 sseg	0xC000.0000-0xDFFF.FFFF	Mapped through TLB, accessible with supervisor or kernel privilege (hence the alternate name).
kseg3	0xE000.0000-0xFFFF.FFFF	Mapped through TLB, accessible only with kernel privileges.

## 6.2 Fixed mapping option

To save chip area for applications not needing a full TLB, threads in one or both VPEs can use a simple fixed mapping (“FMT”) memory translator, which plays the same role. You can find out whether a VPE has fixed mappings by reading the CP0 field *Config[MT]* (see [Figure C-4](#) and descriptions). With the fixed mapping option, virtual address ranges are hard-wired to particularly physical address windows, and cacheability options are set through CP0 register fields as summarized in [Table 6.2](#):

**Table 6.2 Fixed memory mapping**

Segment Name	Virtual range	Physical range	Cacheability bits from
kuseg	0x0000.0000-0x7FFF.FFFF	0x4000.0000-0xBFFF.FFFF	<i>Config[KU]</i>
kseg0	0x8000.0000-0x9FFF.FFFF	0x0000.0000-0x1FFF.FFFF	<i>Config[K0]</i>
kseg1	0xA000.0000-0xBFFF.FFFF	0x0000.0000-0x1FFF.FFFF	(uncached)
kseg2/3	0xC000.0000-0xFFFF.FFFF	0xC000.0000-0xFFFF.FFFF	<i>Config[K23]</i>

Note that even in fixed-mapping mode, the cache parity error status bit *Status[ERL]* still has the effect (required by the MIPS32 architecture) of usurping the normal mapping of “kuseg”; addresses in that range are used unmapped as physical addresses, and all accesses are uncached, until *Status[ERL]* is cleared again.

## 6.3 Reads, writes and synchronization

The MIPS architecture permits implementations a fair amount of freedom as to the order in which loads and stores appear at the CPU interface. Most of the time anything goes, so long as the software behaves correctly.

### 6.3.1 Read/write ordering and queues in the 34K core

To understand the timing of loads and stores (and sometimes instruction fetches), we need to say a little more about the internal construction of the 34K core. In order to maximize performance:

- *Loads are non-blocking*: execution continues “through” a load instruction, and only stops when the program tries to use the GPR value it just loaded.
- *Writes are “posted”*: a write from the core is put aside (the hardware stores both address and data) until the CPU can get access to the system interface and send it off.
- *Cache refills are completed “opportunistically”*: the CPU may still be running on from a non-blocking load or prefetch when data arrives back from the cache. The data required to make good a miss can be forwarded to the relevant GP register, so the returning data is not urgently needed in the cache. The data waits until a convenient moment before it gets put into the cache line.

All of these are implemented with “queues”, called the LDQ, WBB and FSB respectively. All the queues handle data first-come, first served. They need to be *snooped* - a subsequent store to a location with a load pending had better not be allowed to go ahead until the load is complete, for example. So each queue entry is tagged with the address of the data it contains.

An LDQ entry is required for every load that misses in the cache. Moreover, an LDQ entry must be available for any load - even if it will hit in the cache, the logic requires that the LDQ entry is available if needed. This queue allows the pipeline to keep running even though there are outstanding loads. When the load data is finally returned



from the system, the LDQ and the main core logic act together to write this data into the correct GPR (which will then wake the GPR's TC, if it was blocked on an attempt to use this register).

The WBB (Write Back Buffer) queue holds data waiting to be sent out over the system interface, either from D-cache writebacks or uncached/write-through store instructions.

FSB (Fill Store buffer) queue entries are used to hold data that is waiting to be written into the D-cache. D-cache writes are completed opportunistically in a gap between CPU-side cache accesses for loads and stores. An FSB entry gets used during a cache miss (when it holds the refill data), or a write which hits in the cache (when it holds the data the CPU wrote). Loads and stores snoop the FSB so that accesses to lines "in flight" can be dealt with correctly.

All this has a number of consequences which may be visible to software:

- *Number of non-blocking loads which may be pending:* the CPU has either four or nine LDQ entries according to configuration (but it's always at least one per TC.) That limits the number of outstanding loads. As mentioned above, you can't start a load - even one which will in fact hit in the cache - unless you have a free LDQ entry.
- *Hit-under-miss:* the D-cache continues to supply data on a hit, even though there are outstanding misses with data in flight. FSB entries remember the in-flight data. So it is quite normal for a read which hits in the cache to be "completed" - in the sense that the data reaches a register - before a previous read which missed.
- *Write-under-miss:* the CPU pipeline continues and can generate stores even though a read is pending, so long as WBB slots are available. The 34K core's "OCP" interface is non-blocking too (reads consist of separate address and data phases, and writes are permitted between them), so this behavior can often be visible to the system.
- *Miss under miss:* the 34K core can continue to run until the pending read operations exhaust FSB or LDQ entries.
- *Core interface ordering:* at the core interface, read operations may be split into an address phase and a later data phase, with other bus operations in between.

The 34K core - as is permitted by [MIPS32] - makes only limited promises about the order in which reads and writes happen at the system interface. In particular, uncached or write-through writes may be overtaken by cache line reads triggered by a load/store cache miss *later* in sequence. However, uncached reads and writes are always presented in their pipeline sequence (program sequence inside a thread). Use a **sync** instruction where required, as described in the next section.

### 6.3.2 The "sync" instruction in 34K

If you want to be sure that some other agent in the system sees a pair of transactions to uncached memory in the order of the instructions that caused them, you should put a **sync** instruction between the instructions. Other MIPS32/64-compliant CPUs may reorder loads and stores even more; portable code should use **sync**<sup>1</sup>.

But sometimes it's useful to know more precisely what **sync** does on a particular core. On 34K **sync**:

- Stalls until all loads, stores, refills are completed and all write buffers are empty (that is until the LDQ, FSB and WBB are empty);
- If the *Config7[ES]* bit is set, it will cause a synchronizing transaction on the OCP system interface<sup>2</sup>.

1. Note that **sync** is described as only working on "uncached pages or cacheable pages marked as coherent". But **sync** also acts as a synchronization barrier to the effects produced by routine cache-manipulation instructions - hit-writeback and hit-invalidate.

### 6.3.3 Write gathering and “write buffer flushing” in 34K

We mentioned above that uncached writes to the system are performed somewhat lazily, the write being held in the WBB queue until a convenient moment. That can have two system-visible effects:

- Writes can happen later than you think. Your write will happen before the next uncached read or write, but that’s all you know. To make sure that a write has gone out on the OCP bus you can use a **sync** (as above): but that meaning of **sync** is CPU-dependent, so that code is non-portable. Also, your write might still be posted somewhere in a system controller, unless you know your system is built to prevent it. Sometimes it’s better to code a dummy uncached read from a nearby location (which will “flush out” buffered writes on pretty much any system).
- Uncached writes to locations in the same “cache line”-sized chunk of memory may be gathered - stored together in the WBB, and then dealt with by a single “wider” OCP write than the one you originally coded. Sometimes, this is what you want. When it isn’t, put a **sync** between your successive writes.

### 6.3.4 Parity error exception handling and the CacheErr register

The 34K core does not check parity on data (or control fields) from the external interface - so this section really is just about parity protection in the cache. It’s a build-time option, selected by your system integrator, whether to include parity bits in the cache and parity check logic.

At a system level, a cache parity exception is usually fatal - though recovery might be possible sometimes, when it is useful to know that the restart address is in *ErrorEPC* and you can return from the exception with an **eret**.

But mainly, diagnostic-code authors will probably find the *CacheErr* register’s extra information useful.

**Figure 6-1 Fields in the CacheErr register**

31	30	29	28	27	26	25	24	23	22	21	20	19	16	15	0
ER	EC	ED	ET	ES	EE	EB	EF	SP	EW	Way		0		Index	

*ER*: was the error on an I-fetch (0) or on data (1)?

*EC*: in L1 cache (0) or higher-level cache (1)?

*ED,ET*: 1 for error in data field/tag field respectively.

*ES*: always zero on the 34K core - used to indicate an error detected during an external "snoop" for cache-coherent CPUs.

*EE*: always 0 on the 34K core (it would be 1 if this error came from a parity error in data at the system interface, but there’s no system parity on the core).

*EB*: 1 if data and instruction-fetch error reported on same instruction, which is unrecoverable. If so, the rest of the register reports on the instruction-fetch error.

*EF*: unrecoverable (fatal) error (other than the *EB* type above). Some parity errors can be fixed by invalidating the cache line and relying on good data from memory. But if this bit is set, all is lost... It’s one of the following:

1. Line being displaced from cache ("victim") has a tag parity error, so we don’t know whether to write it back, or whether the writeback location (which needs a correct tag) would be correct.
2. This will be a read with the signal *OC\_MReqInfo[3]* set. Handling of this transaction is system dependent, but a typical system controller will flush any external write buffers and complete all pending transactions before telling the CPU that the transaction is completed. Ask your system integrator how it works in your SoC.

2. The victim's tag indicates it has been written by the CPU since it was obtained from memory (the line is "dirty" and needs a write-back), but it has a data parity error.
3. Writeback store miss and *CacheErr[EW]* error.
4. At least one more cache parity error happened concurrently with or after this one, but before we reached the relative safety of the cache parity error exception handler.

*SP*: error affecting a scratchpad RAM access, see [Section 6.5, "Scratchpad memory/SPRAM"](#) below.

*EW*: parity error on the "dirty" (cache modified) or way-selection bits. That's not recoverable.

*Way*: the way-number of the cache entry where the error occurred.

*Index*: the cache index (within the cache way) of the entry where the error occurred... except that the low bits are not meaningful. The index is aligned as if a byte address, which is good because that's what Index-type **cache** instructions need. It identifies the failing doubleword for a data error, or just the failing line for a tag error.

The index-type **cache** instruction will need an "index" with the way bits glued on top of this cache-entry field; you know how to put that together, because the shape of the cache is defined in the *Config1-2* registers as shown in [Figure C-5](#).

### 6.3.5 ErrCtl register

This register controls parity protection of the L1 caches (if it was configured in your core in the first place) and provides for software testing of the whole cache array, including the otherwise-inaccessible way-selection RAM.

**Figure 6-2 Fields in the ErrCtl register**

31	30	29	28	27	26	25	24	23	19	18	13	12	4	3	0
PE	PO	WST	SPR	PCO	ITC	LBE	WABE	0		PCI		PI			PD

In summary: running software should set this register to 0x8000.0000 to enable cache parity checking, and to zero otherwise. Other uses are more obscure, but the fields are as follows:

*PE*: 1 to enable cache parity checking. Hard-wired to zero if parity isn't implemented.

*PO*: (parity overwrite) - set 1 to set the parity bit regardless of parity computation, which is only for diagnostic/test purposes.

After setting this bit you can use **cache IndexStoreTag** to set the cache data parity to the value currently in *ErrCtl[PI]* (for I-cache) or *ErrCtl[PD]* (for D-cache), while the tag parity is forcefully set from *TagLo[P]*.

*WST*: test mode for **cache IndexLoadTag/cache IndexStoreTag** instructions, which then read/write the cache's internal "way-selection RAM" instead of the cache tags.

*SPR*: when set, index-type **cache** instructions work on the scratchpad/SPRAM, if fitted - see [Section 6.5, "Scratchpad memory/SPRAM"](#).

*PCO/PCI*: precode override and data. Used for diagnostic/test of the I-cache feature which partially decodes instructions at cache refill time. Not properly documented here.

*PI/PD*: parity bits being read/written to caches (I- and D-cache respectively).

*ITC*: set to make **cache IndexLoadTag/cache IndexStoreTag** operate on the control/configuration "tags" for ITC storage locations - see [Section 3.3.1, "Configuring ITC base address and cell repeat interval"](#).

## Memory map, caching, reads and writes and translation

*LBE*, *WABE*: field indicating whether a bus error (the last one, if there's been more than one) was triggered by a load or a write-allocate respectively: see below. Where both a load and write-allocate are waiting on the same cache-line refill, both could be set. These bits are “sticky”, remaining set until explicitly written zero.

### 6.3.6 Bus error exception

The CPU's "OCP" hardware interface rules permit a slave device attached to the system interface to signal back when something has gone wrong with a read. This should not be used to report a read parity error; if parity is checked externally, it would have to be reported through an interrupt. Typically a bus error means that some subsystem has failed to respond.

Bus errors are not signalled on an OCP write cycle, and (if they were) the 34K core ignores them.

The bus error is imprecise; that is, *EPC* does not necessarily (or even usually) point to the instruction causing the memory read (though it is precise for a bus error on an I-fetch).

Data-side bus errors are usually caused by a load: and the (non-blocking) load which caused it may have happened a long time ago.

If software knows that a particular read might encounter a bus error - typically it's some kind of probe - it should be careful to stall and wait for the load value immediately, by reading the value into a register, and make sure it can handle a bus error at that point.

On a load the hardware knows which TC or TCs were waiting for the load which went wrong, and the *TCBind[TBE]* bit will be set for each suffering TC.

There is an obscure corner case. The 34K core's D-cache is “write-allocate”: so a write which misses in the cache will trigger a read to fill the cache line ready to receive the new data.

After a bus error you can look at *ErrCtl[LBE]/ErrCtl[WABE]* to see whether the error was caused by a load or write-allocate.

## 6.4 Caches

Most of the time caches just work and are invisible to software... though your programs would go twenty times slower without them. But this section is about when caches aren't invisible any more.

Like most modern MIPS CPUs, the 34K core has separate primary I- and D-caches. They are virtually-indexed and physically-tagged, so you may need to deal with *cache aliases*, see [Section 6.4.7, "Cache aliases"](#). The design provides for 16Kbyte, 32Kbyte or 64Kbyte caches; but the largest of those are likely to come with some speed penalty. The 34K core's primary caches are 4-way set associative.

But don't hard-wire any of this information into your software. Instead, probe the *Config1* register defined by [\[MIPS32\]](#) to determine the shape and size of the cache.

### 6.4.1 Cacheability options

Any read or write made by the 34K core will be cacheable or not according to the virtual memory map. For addresses translated by the TLB the cacheability is determined by the TLB entry; the key field appears as *EntryLo[C]*. [Table 6.3](#) shows the code values used in *EntryLo[C]* - the same codes are used in the *Config* entries used to set the behavior of regions with fixed mappings (the latter are described in [Table C-4](#).)

Some of the undefined cacheability code values are reserved for use in cache-coherent systems.

**Table 6.3 Cache Code Values**

Code	Cached?	How it Writes	Notes
0	cached	write-through	An unusual choice for a high-speed CPU, probably only for debug
2	uncached		
3	cached	writeback	All normal cacheable areas
7	uncached	“Uncached Accelerated”	Unusual and interesting mode for high-bandwidth write-only hardware; see Section 6.4.2, “Uncached accelerated writes”.

### 6.4.2 Uncached accelerated writes

The 34K core permits memory regions to be marked as “uncached accelerated”. This type of region is useful to hardware which is “write only” - perhaps video frame buffers, or some other hardware stream. Sequential word stores in such regions are gathered into cache-line-sized chunks, before being written with a single burst cycle on the CPU interface.

Such regions are uncached for read, and partial-word or out-of-sequence writes have “unpredictable” effects - don’t do them. The burst write is normally performed when software writes to the last location in the memory block or does an uncached-accelerated write to some other block; but it can also be triggered by a **sync** instruction, a **prefnudge**, a matching load or any exception. If the block is not completely written by the time it’s pushed out, it will be written using a series of doubleword or smaller write cycles over the 34K core’s 64-bit memory interface.

### 6.4.3 The cache instruction and software cache management

The 34K core’s caches are not fully “coherent” and require OS intervention at times. The **cache** instruction is the building block of such OS interventions, and is required for correct handling of DMA data and for cache initialization. Historically, the **cache** instruction also had a role when writing instructions (unless the programmer takes some action, those instructions may only be in the D-cache whereas you need them to be fetched through the I-cache when the time comes). See Section 6.4.4 “Cache management when writing instructions - the “synci” instruction” below.

A cache operation instruction is written: **cache op, addr** where the **addr** is just an address format, written as for a load/store instruction. Cache operations are privileged and can only run in kernel mode (but see the note on the user-privilege **synci** instruction at the end of this section). Generally we’re not showing you instruction encodings in this book (you have software tools for that stuff) but in this case it’s probably necessary, so take a look at Figure 6-3.

**Figure 6-3 Fields in the encoding of a cache instruction**

31	26	25	21	20	18	17	16	15	0
<b>cache</b>		base		<b>op</b>		offset			
47		register		what	which				
				to do	cache				

The **op** field packs together a 2-bit field which selects which cache to work on:


- 0 L1 I-cache
- 1 L1 D-cache
- 2 reserved for L3 cache
- 3 reserved for L2 cache

## Memory map, caching, reads and writes and translation

and then adds a 3-bit field which encodes a command to be carried out on the line the instruction selects.

Before we list out the individual commands; the cache commands come in three flavors which differ in how they specify the cache entry (the “cache line”) they will work on:

- *Hit-type cache operation*: presents an address (just like a load/store), which is looked up in the cache. If this location is in the cache (it “hits”) the cache operation is carried out on the enclosing line. If this location is not in the cache, nothing happens.
- *Address-type cache operation*: presents an address of some memory data, which is processed just like a cached access - if the cache was previously invalid the data is fetched from memory.
- *Index-type cache operation*: as many low bits of the virtual address as are required are used to select the byte within the cache line, then the cache line address inside one of the four cache ways, and then the way. You have to know the size of your cache (discoverable from *Config1-2*) to know exactly where the field boundaries are, but your address is used something like this:



Don’t define your own C names for cache manipulation operation codes, at least not if you can use [\[m32c0.h\]](#)

Once you’ve picked your cache and cache line you have a choice of operations you can perform on it. In fact, the 34K core implements all the **cache** instructions defined by [\[MIPS32\]](#) (whether said there to be “required”, “recommended” or “optional”) and all of those are shown in [Table 6.4](#).

### 6.4.3.1 Read/write synchronization and the cache instruction

Before any **cache** instruction is allowed to execute, any outstanding loads and cache refills are completed, and any outstanding stores or cache line writebacks are sent to the write buffer - that is, the LDQ, FSB and WBB queues are all drained. This is somewhat like the 34K core-specific results of the **sync** instruction - but portable code should not make that assumption.

### 6.4.4 Cache management when writing instructions - the “synci” instruction

The **synci** instruction (new to the MIPS32 Release 2 update) provides a clean mechanism - available to user-level code, not just at kernel privilege level - for ensuring that instructions you’ve just written are correctly presented for execution (it combines a D-cache writeback with an I-cache invalidate). You should use it in preference to the traditional alternative of a D-cache writeback followed by an I-cache invalidate.

### 6.4.5 Cache management and multithreaded CPUs

The cache management registers are all replicated per-VPE but not per-TC, so obviously you have to avoid multiple threads on the same VPE attempting to use cache operations concurrently.

Moreover, in 34K family cores the two VPEs share the cache. In general write-back operations and the kind of invalidate which automatically writes-back a dirty line may be safely run by either VPE at any time. All other operations may cause undesirable effects unless you make sure they’re done by only one VPE at a time; and in particular, you should get the cache initialized by one VPE running alone.

Table 6.4 Operations on a cache line available with the cache instruction

Value	Command	What it does
0	Index invalidate	Sets the line to "invalid". If it's a D-cache line which is valid and "dirty" (has been written by CPU since fetched from memory), then write the contents back to memory first. This is the best and simplest way to invalidate an I-cache when initializing the CPU - though if your cache is parity-protected, you also need to fill it with good-parity data, see <b>Fill</b> below. And this is not suitable for initializing D-caches, where it might cause random write-backs: see <b>Index Store Tag</b> type below.
1	Index Load Tag	Read the cache line tag bits and addressed doubleword data into the I- or D-side cache tag registers <i>ITagLo</i> , <i>DTagLo</i> etc. For diagnostics and geeks only.
2	Index Store Tag	Set the cache tag from the <i>ITagLo</i> / or <i>DTagLo</i> registers. To initialize a D-cache from an unknown state, set the <i>DTagLo</i> / <i>DTagHi</i> registers to zero and then do this to each line.
3	Index Store Data	Write cache-line data. Not available for caches, but it is used for management of scratchpad RAM regions described in Section 6-8, "SPRAM (scratchpad RAM) configuration information in TagLo" and ITC regions described in Section 3.3.1, "Configuring ITC base address and cell repeat interval".
4	Hit invalidate	hit-type invalidate - do not writeback the data even if dirty. May cause data loss unless you know the line is not dirty.
5		<i>Sorry, different meanings for code "5" on I- and D-caches.</i>
	Writeback invalidate	<i>On a D-cache:</i> (hit-type operation) invalidate the line but only after writing it back, if dirty. This is the recommended way of invalidating a D-cache line in a running cache.
	Fill	<i>On an I-cache:</i> (address-type operation) fill a suitable cache line from the data at the supplied address - it will be selected just as if you were processing an I-cache miss at this address. Used to initialize an I-cache line's data field, which should be done when setting up the CPU when the cache is parity protected.
6	Hit writeback	If the line is dirty, write it back to memory but leave it valid in the cache. Used in a running system where you want to ensure that data is pushed into memory for access by a DMA device or other CPU.
7	Fetch and Lock	An address-type operation. Get the addressed data into the same line as would be used on a regular cached reference (if the data wasn't already cached that might involve writing back the previous occupant of the cache line). Then lock the line. Locked lines are not replaced on a cache miss. It stays locked until explicitly invalidated with a <b>cache</b> instruction.

### 6.4.6 Cache initialization and tag/data registers

The *ITagLo*, *DTagLo*, *IDataLo*, *IDataLo* and *IDataHi* registers are used for staging tag information being read from or written to the cache by some particular **cache** instructions (the 34K core has no "TagHi" registers, which are only needed for CPUs with a bigger physical address range). [MIPS32] declares that the contents of these registers is implementation dependent, so they are described here.

*ITagLo* is used for the I-cache and *DTagLo* for the D-cache. *TagLo2* is reserved for secondary cache management, and is not yet defined for the 34K family. Some other MIPS CPUs use the same staging register for the I- and D-cache, and initialization software written for such CPUs is not portable to the 34K core.

For the cache itself, you will only need to use the data and tag registers for initialization and diagnostics. But you will also need them when using cache instructions to configure and manage scratchpad memory (see the next section) and ITC locations as described in Section 3.3 "Inter-thread communication storage (ITC)".

cache line. Only diagnostic and test software will need to know details; but Figure 6-6 shows all the fields:

**Figure 6-4 Fields in the ITagLo and DTagLo Registers**

31	12 11	8 7 6 5 4	1 0
PTagLo	0	V D L	0 P

*ITagLo* and *DTagLo* can be used in a special mode; when *ErrCtl[WST]* is 1, the appropriate *TagLo* register's fields change completely, as shown in [Figure 6-7](#) below. But let's look at the standard fields first:

*PTagLo*: the cache address tag - a physical address because the 34K core's caches are physically tagged. It holds bits 31–12 of the physical address - the low 12 bits of the address are implied by the position of the data in the cache.

*V*: 1 when this cache line is valid.

*D*: 1 when this cache line is dirty (that is, it has been written by the CPU since being read from memory).

*L*: 1 when this cache line is locked, see [Section 6.4.8, "Cache locking"](#).

*P*: parity bit for tag fields other than the *TagLo[D]* bit, which is actually held separately in the "way-select" RAM. When you use the *TagLo* register to write a cache tag with **cacheIndexStoreTag** the

*TagLo[P]*: bit is generally not used - instead the hardware puts together your other fields and ensures it writes correct parity. However, it is possible to force parity to exactly this value by first setting *ErrCtl[PO]*.

**Figure 6-5 Fields in DTagLo/ITagLo when used for way-select RAM**

31	24 23	20 19	16 15	10 9	8 7	5 4	1 0
×	WSDP	WSD	WSLRU	0	×	0	×

When *ErrCtl[WST]* is set, **cacheIndexLoadTag** and **cacheIndexStoreTag** operations read/write the separate "way-select RAM" used in the 34K core's caches. Then the fields in *ITagLo* / *DTagLo* change to those shown in [Figure 6-7](#) above. These are:

*WSDP*: parity check for each of the "dirty" bits. It's like the regular tag parity bit, in that this field is not normally used when you're writing into the way-select RAM. If you want to force these values in, you need to set *ErrCtl[PO]* to 1.

*WSD*: dirty bits - found on D-side only, so not in *ITagLo*.

*WSLRU*: LRU bits.

### 6.4.7 Cache aliases

34K core has caches which are virtually indexed. Since it's quite routine to have multiple virtual mappings of the same physical data, it's possible for such a cache to end up with two copies of the same data. That's not a problem for the normal operation of caches for read-only data, but becomes troublesome:

- *When you want to write the data*: if a line is stored in two places, you'll only update one of them and some data will be lost (at least, there's a 50% chance it will be lost!) This is obviously disastrous: systems generally work hard to avoid aliases in the D-cache.
- *When you want to invalidate the line in the cache*: there's a danger you might invalidate one copy but not the other. This (more subtle) problem can affect the I-cache too.

It can be worked around. There's no problem for different virtual mappings which generate the same cache index; those lines will all compete for the 4 ways at that index, and they're distinguished through the physical tag.



The 34K CPU's smallest page size is 4Kbytes, that's  $2^{12}$  bytes. The paged memory translation means that the low 12 bits of a virtual address is always reproduced in the physical address. Since a 16Kbyte, 4-way set-associative, cache gets its index from the low 12 bits of the address, the 16Kbyte cache is alias-free. You can't get aliases if each cache "way" is no larger than the page size.

In 32Kbyte and 64Kbyte caches, one or two top bits used for the index are not necessarily the same as the corresponding bits of the physical address, and aliases are possible. The value of the one or two critical virtual address bits is sometimes called the *page color*.

It's possible for software to avoid aliases if it can ensure that where multiple virtual mappings to a physical page exist, they all have the same color. You do that by enforcing virtual-memory alignment rules (to at least a 16Kbyte boundary) for shareable regions. It turns out this is practicable over a large range of OS activities: sharing code and libraries, and deliberate interprocess shared memory. It is not so easy to do in other circumstances, particularly when pages to be mapped start their life as buffers for some disk or network operation<sup>1</sup>...

So the 34K core contains logic to make the popular 32Kbyte D-cache alias-free (effectively one index bit is from the physical address, and some ingenious tricks used to prevent that slowing the whole process excessively). The *Config7[AR]* flag should read 1 if your 32Kbyte-D-cached core was built to be alias-free.

A 32Kbyte I-cache, or any 64Kbyte I- or D-cache, are subject to aliases.

### 6.4.8 Cache locking

[MIPS32] provides for a mechanism to lock a cache line so it can't be replaced. This avoids cache misses on one particular piece of data, at the cost of reducing overall cache efficiency.

**Caution:** in complex software systems it is hard to be sure that cache locking provides any overall benefit - most often, it won't. You should probably only use locking after careful measurements have shown it to be effective for your application.

Lock a line using a **cache FetchAndLock** (it will not in fact re-fetch a line which is already in the cache).

Unlock it using any kind of relevant **cache "invalidate"** instruction<sup>2</sup> - but note that **synci** won't do the job, and should not be used on data/instruction locations which are cache-locked.

### 6.4.9 Cache control for Multithreading CPU

In normal circumstances all threads on an MT CPU simply share the cache. Because the L1 caches are 4-way set associative, they should behave quite well even though the multiple threads will generally have a larger and more complex working set of data and instructions. But if you really need to prevent threads (at least in different VPEs) from competing for the same cache resources, you can do that by setting up the *VPEOpt* register, described in Section 2.9.10, "VPEOpt register - reserve some cache "way" for use of one VPE" on page 40.

### 6.4.10 Cache management and multithreaded CPUs

The cache management registers are all replicated per-VPE but not per-TC, so obviously you have to avoid multiple threads on the same VPE attempting to use cache operations concurrently.

- 
1. There's a fair amount of rather ugly code in the MIPS Linux kernel to work around aliases.
  2. It's possible to lock and unlock lines by manipulating values in the *TagLo* register and then using a **cache Index\_Load\_Tag** instruction... but highly non-portable and likely to cause trouble. Probably for diagnostics only.

Moreover, in 34K family cores the two VPEs share the cache. In general write-back operations and the kind of invalidate which automatically writes-back a dirty line may be safely run by either VPE at any time. All other operations may cause undesirable effects unless you make sure they're done by only one VPE at a time; and in particular, you should get the cache initialized by one VPE running alone.

### 6.4.11 Cache initialization and tag/data registers

The *ITagLo*, *DTagLo*, *IDataLo*, *IDataLo* and *IDataHi* registers are used for staging tag information being read from or written to the cache (the 34K core has no “TagHi” registers, which are only needed for CPUs with a bigger physical address range). [MIPS32] declares that the contents of these registers is implementation dependent, so they need some words here.

*ITagLo* is used for the I-cache and *DTagLo* for the D-cache. *TagLo2* is reserved for secondary cache management, and is not yet defined for the 34K family. Some other MIPS CPUs use the same staging register for the I- and D-cache, and initialization software written for such CPUs is not portable to the 34K core.

Before getting into the details, note that it's a strong convention that you can write all-zeros to both *ITagLo* and *DTagLo* registers and then use **cache IndexStoreTag** to initialize a cache entry to a legitimate (but empty) state. Your cache initialization software should rely on that, not on the details of the registers.

Only diagnostic and test software will need to know details; but Figure 6-6 shows all the fields:

**Figure 6-6 Fields in the ITagLo and DTagLo Registers**

31	12 11	10 9	8 7	6 5 4	1 0
PTagLo		×	0	V D L	0 P

*ITagLo* and *DTagLo* can be used in a special mode; when *ErrCtl[WST]* is 1, the appropriate *TagLo* register's fields change completely, as shown in Figure 6-7 below. But let's look at the standard fields first:

*PTagLo*: the cache address tag - a physical address because the 34K core's caches are physically tagged. It holds bits 31–12 of the physical address - the low 12 bits of the address are implied by the position of the data in the cache.

×: a field not described for the 34K core but which might not always read zero.

V: 1 when this cache line is valid.

D: 1 when this cache line is dirty (that is, it has been written by the CPU since being read from memory).

L: 1 when this cache line is locked, see Section 6.4.8, "Cache locking".

P: parity bit for tag fields other than the *TagLo[D]* bit, which is actually held separately in the "way-select" RAM. When you use the *TagLo* register to write a cache tag with **cache IndexStoreTag** the

*TagLo[P]*: bit is generally not used - instead the hardware puts together your other fields and ensures it writes correct parity. However, it is possible to force parity to exactly this value by first setting *ErrCtl[PO]*.

**Figure 6-7 Fields in DTagLo/ITagLo when used for way-select RAM**

31	24 23	20 19	16 15	10 9	8 7	5 4	1 0
×	WSDP	WSD	WSLRU	0	×	0	×

When *ErrCtl[WST]* is set, **cache IndexLoadTag** and **cache IndexStoreTag** operations read/write the separate “way-select RAM” used in the 34K core's caches. Then the fields in *ITagLo* / *DTagLo* change to those shown in Figure 6-7 above. These are:

*WSDP*: parity check for each of the "dirty" bits. It's like the regular tag parity bit, in that this field is not normally used when you're writing into the way-select RAM. If you want to force these values in, you need to set *ErrCtl[PO]* to 1.

*WSD*: dirty bits - found on D-side only, so not in *ITagLo*.

*WSLRU*: LRU bits.

## 6.5 Scratchpad memory/SPRAM

Most of MIPS Technologies' cores can be equipped with a modestly sized high speed on-chip data memory, called "scratchpad RAM" or "SPRAM". SPRAM is connected to a cache interface, alongside the I- or D-cache, so is available separately for the I- and D-side ("ISPRAM" and "DSPRAM").

MIPS Technologies provide the interface on which users can build many types and sizes of SPRAM. We also provide a "reference design" for both ISPRAM and DSPRAM, which is what is described here. If you keep the programming interface the same as the reference design, you're more likely to be able to find software support. The reference design allows for on-chip memories of up to 1Mbytes in size.

There are two possible motives for incorporating SPRAM:

- *Dedicated high-speed memory*: small SPRAM arrays run with cache timing. Larger arrays may require one or more clocks of extra latency. Although that may still very fast compared with any memory access through the OCP interface, the SPRAM is replacing single-cycle cache: multi-cycle instruction-side SPRAM is likely to reduce performance substantially, and you should think hard before specifying it.

SPRAM can be made much larger than the maximum cache size.

Even for smaller sizes, it is possible to envisage applications where some particularly heavily-used piece of data is well-served by being permanently installed in SPRAM. Possible, but unusual. In most cases heavily-used data will be handled well by the D-cache, and until you really know otherwise it's better for the SoC designer to maximize cache (compatible with his/her frequency needs.)

But there's another more compelling use for a modest-size SPRAM:

- *"DMA" accessible to external masters on the OCP interface*: the SPRAM can be configured to be accessible from an OCP interface. OCP masters will see it just as a chunk of memory which can be read or written.

Because SPRAM stands in for the cache, data passed through the SPRAM in this way doesn't require any software cache management. This makes it spectacularly efficient as a staging area for communicating with complex I/O devices: a great way to implement "push" style I/O (that is where the device writes incoming data close to the CPU).

SPRAM must be located somewhere within the physical address map of the CPU, and is usually accessed through some "cached" region of memory (uncached region accesses work with the 34K reference design, but may not do so on other implementations - better to keep it cached). It's usually better to put it in the first 512Mbytes of physical space, because then it will be accessible through the simple kseg0 "cached, unmapped" region - with no need to set up specific TLB entries.

Because the SPRAM is close to the cache, it inherits some bits of cache housekeeping. In particular the **cache** instruction and the cache tag CPO registers are used to provide a way for software to probe for and establish the size of SPRAM<sup>1</sup>, and (in the case of ISPRAM) to load instructions into it.

**Probing for SPRAM configuration**

The presence of scratchpad RAM in your core is indicated by a “1” bit in one or both of the CPO *Config[ISP,DSP]* register flags described in Figure C-4. The MIPS Technologies reference design requires that you can query the size of and adjust the location of scratchpad RAM through “cache tags”.

To access the SPRAM “tags” (where the configuration information is to be found) first set the *ErrCtl[SPR]* bit (see Figure 6-2 above).

Now a **cache Index\_Load\_Tag\_D, 0<sup>1</sup>** instruction fetches half the configuration information into *DTagLo*, and a **cache Index\_Load\_Tag, 8** gets the other half (the “8” steps to the next feasible tag location - an artefact of the 64-bit width of the cache interface.) The corresponding operations directed at the primary I-cache read the halves of the I-side scratchpad tag, this time into *ITagLo*. The fake tags for I-side and D-side SPRAM have the same format; the information appears in *TagLo* fields as shown in Figure 6-8.

**Figure 6-8 SPRAM (scratchpad RAM) configuration information in TagLo**

	31	12 11	8	7	6	5	4	1	0
<i>TagLo</i>	physical address tag		0	valid	dirty	locked	0	parity	
addr == 0	base address[31:12]		0	En	0				
addr == 8	size of region in bytes/4KB		0						

Where:

- *base address[31:12]*: the high-order bits of the physical base address of this chunk of SPRAM;
- *En*: enable the SPRAM. From power-up this bit is zero, and so long as it stays that way the SPRAM acts as though it isn’t there;
- *size of region in bytes/4KB*: the number of page-size chunks of data mapped. If you take the whole 32 bits, it returns the size in bytes (but it will always be a multiple of 4KB).

In some MIPS cores using this sort of tag setup there could be multiple scratchpad regions indicated by two or more of these tag pairs. But the reference design provided with the 34K core can only have one I-side and one D-side region.

You can load software into the ISPRAM using cacheops. Each pair of instructions to be loaded are put in the registers *IDataHi*/*IDataLo*, and then you use a **cache Index\_Store\_Data\_I** at the appropriate index. The two data registers work together to do a 64-bit transfer (the 34K core’s instruction memory really is 64 bits wide), so for a CPU configured big-endian the first instruction in sequence is loaded into *IDataHi*, but for a CPU configured little-endian the first instruction is loaded into *IDataLo*.

Don’t forget to set *ErrCtl[SPR]* back to zero when you’re done.

1. What follows is a hardware convention which SoC designers are not compelled to follow; but MIPS Technologies recommends designers to do SPRAM this way to ease software porting.
1. The instructions are written as if using C “#define” names from [m32c0.h]


## 6.6 The TLB and translation

The TLB is the key piece of hardware which MIPS architecture CPUs have for memory management. It's a table whose input is a virtual address together with the "address space identifier" from *EntryHi[ASID]* and whose output is a physical address plus cacheability attributes. System software maintains the TLB as a cache of a much larger number of possible translations, and a special-cased handler for the exception raised when there's no suitable translation in the TLB provides adequate performance. Read on for a summary of all the fields and how it gets used; but the OS ramifications are far too extensive to cover here; for a better description in context see [SEEMIPSRUN]; and for full details of the architectural specification see [MIPS32].

### 6.6.1 The TLB array

Let's start with a sketch of a TLB entry. For the 34K core, that consists of a virtual address portion to match against and two output sections, something like Figure 6-9 - which also shows which TLB fields are carried in which CP0 registers.

Figure 6-9 Fields in a 34K™ core TLB entry



Some points to make about the TLB entry:

- The input-side virtual address fields (to the left) have the fields necessary to match an incoming address against this entry. "VPN" is (by OS tradition) a "virtual page number" - the high bits of the program (virtual) address. "VPN2" is used to remind you that this address is for a pair of pages...
- The right-hand side (physical) fields are the information used to output a translation. There are a pair of outputs for each input-match, and which of them is used is determined by the highest within-match address bit. So in standard form (when we're using 4Kbyte pages) each entry translates an 8Kbyte region of virtual address, but we can map each 4Kbyte page onto any physical address (with any permission flag bits).
- The size of the input region is configurable because the "PageMask" determines how many incoming address bits to match. The 34K core allows page sizes of 4Kbytes, 16Kbytes and going on in powers of 4 up to 256Mbytes. That's expressed by the legal values of *PageMask*, shown below.
- The "ASID" field extends the virtual address with an 8-bit, OS-assigned memory-space identifier so that translations for multiple different applications can co-exist in the TLB (in Linux, for example, each application has different code and data lying in the same virtual address region).
- The "G" (global) bit is not quite sure whether it's on the input or output side - there's only one, but it can be read and written through either of *EntryLo0-1*. When set, it causes addresses to match regardless of their ASID value, thus defining a part of the address space which will be shared by all applications. For example, Linux applications share some "kseg2" space used for kernel extensions.

## 6.6.2 The TLB and the MIPS® MT ASE

Cores in the 34K family are built with just one piece of TLB hardware. However, you can configure your CPU with the TLB either shared between two VPEs, or partitioned so that each VPE sees a standard (though smaller) TLB array.

TLB sharing will usually provide the best performance for when the VPEs are running the same kernel, are closely collaborating, or when one of them makes little or no use of translated addresses. TLB sharing is not completely software-transparent, and some OS work will be needed. See [Section 4.2.2, "Sharing and not sharing the TLB"](#) for details.

## 6.6.3 Live translation and micro-TLBs

When you're really tuning out the last cycle, you need to know that in the 34K core the translation is actually done by two little tables local to the instruction fetch unit and the load/store unit - called the ITLB and DTLB respectively (or generically, they're "micro-TLBs" or "uTLB"). There are only 4 entries in the ITLB, and 8 in the DTLB and they are functionally invisible to software: they're automatically refilled from the main TLB when required, and automatically cleared whenever the TLB is updated. It costs just three extra clocks to refill the uTLB for any access whose translation is not already in the appropriate uTLB.

## 6.6.4 Reading and writing TLB entries: Index, Random and Wired

Two CP0 registers work as simple indexes into the TLB array for programming: *Index* and *Random*. The oddly-named *Wired* controls *Random*'s behavior.


Of these: *Index* determines which TLB entry is accessed by `tlbwi`. It's also used for the result of a `tlbop` (the instruction you use to see whether a particular address would be successfully translated by the CPU). *Index* only implements enough bits to index the TLB, however big that is; but a `tlbop` which fails to find a match for the specified virtual address sets bit 31 of *Index* (it's easy to test for).

*Random* is implemented as a full CPU clock-rate downcounter. It won't decrement below the value of *Wired* (when it gets there it bounces off and starts again at the highest legal index). In practice, when used inside the TLB refill exception handler, it delivers a random index into the TLB somewhere between the value of *Wired* and the top. *Wired* can therefore be set to reserve some TLB entries from random replacement - a good place for an OS to keep translations which must never cause a TLB translation-not-present exception.

## 6.6.5 Reading and writing TLB entries - staging registers

The TLB is accessed through staging registers which between them represent all the fields in each TLB entry; they're called *EntryHi*, *PageMask* and *EntryLo0-1*. The fields from *EntryHi* and *PageMask* are shown in [Figure 6-10](#).

**Figure 6-10 Fields in the EntryHi and PageMask registers**



All these fields act as staging posts for entries being written to or read from the TLB. But some of them are more magic than that...

*EntryHi*[VPN2]: is the page-pair address to be matched by the entry this reads/writes - see above.

However, on a TLB-related exception *VPN2* is automatically set to the virtual address we were trying to translate when we got the exception. If - as is most often the case - the outcome of the exception handler is to find and install a translation to that address, *VPN2* (and generally the whole of *EntryHi*) will turn out to already have the right values in it.

*EntryHi[ASID]*: does double-duty. It is used to stage data to and from the TLB, but in normal running software it's also the source of the current "ASID" value, used to extend the virtual address to make sure you only get translations for the current process. Because of that role it is replicated per-TC in MIPS MT systems, and is also visible as *TCStatus[TASID]*.

*PageMask[Mask]*: acts as a kind of backward mask, in that a 1 bit means "don't compare this address bit when matching this address". However, only a restricted range of *PageMask* values are legal (that's with "1"s filling the *PageMask[Mask]* field from low bits upward, two at a time):

PageMask	Size of each output page	PageMask	Size of each output page
0x0000.0000	4Kbytes	0x007F.E000	4Mbytes
0x0000.6000	16Kbytes	0x01FF.E000	16Mbytes
0x0001.E000	64Kbytes	0x07FF.E000	64Mbytes
0x0007.E000	256Kbytes	0x1FFF.E000	256Mbytes
0x001F.E000	1Mbyte		

Note that the uTLBs handle only 4Kbyte and 1Mbyte page sizes; other page sizes are down-converted to 4Kbyte or 1Mbyte as they are referenced. For other page sizes this may cause an unexpectedly high rate of uTLB misses, which could be noticeable in unusual circumstances.

Then moving our attention to the output side, the two *EntryLo0-1* are identical in format as shown in [Figure 6-11](#).

**Figure 6-11 Fields in the EntryLo0-1 registers**



In *EntryLo0-1*:

*PFN*: the "physical frame number" - traditional OS name for the high-order bits of the physical address. 24 bits of *PFN* together with 12 bits of in-page address make up a 36-bit physical address; but the 34K core has a 32-bit physical address bus, and does not implement the four highest bits (which always read back as zero).

*C*: a code indicating how to cache data in this page - pages can be marked uncacheable and various flavours of cacheable. The codes here are shared with those used in CP0 registers for the cacheability of fixed address regions: see [Table 6.3](#) in [Section 6.4.1](#), "Cacheability options" on page 76.

*D*: the "dirty" flag. In hardware terms it's just a write-enable (when it's 0 you can't do a store using addresses translated here, you'll get an exception instead). However, software can use it to track pages which have been written to; when you first map a page you leave this bit clear, and then a first write causes an exception which you note somewhere in the OS' memory management tables (and of course remember to set the bit).

*V*: the "valid" flag. You'd think it doesn't make much sense - why load an entry if it's not valid? But this is very helpful so you can make just one of a pair of pages valid.

*G*: the "global" bit. This really belongs to the input side, and you don't really want two values for it. So you should always make sure this is the same in *EntryLo0* and *EntryLo1*.

## 6.6.6 TLB initialization and duplicate entries

TLB entries come up to random values on power-up, and must be initialized by hardware before use. Generally, early bootstrap software should go through setting each entry to a harmless “invalid” value.

Since the TLB is a fully-associative array and entries are written by index, it’s possible to load duplicate entries - two or more entries which match the same virtual address/ASID. In older MIPS CPUs it was essential to avoid duplicate entries - even duplicate entries where all the entries are marked “invalid”. Some designs could even suffer hardware damage from duplicates. Because of the need to avoid duplicates, even initialization code ought to use a different virtual address for each invalid entry; it’s common practice to use “kseg0” virtual addresses for the initial all-invalid entries.

Most MIPS Technologies cores protect themselves and you by taking a “machine check” exception if a TLB update would have created a duplicate entry - but in the 34K core that only happens if both entries are valid.

Earlier MIPS Technologies cores suffer a machine check even if duplicate entries are both invalid. That can happen when initializing. For example, when an OS is initializing the TLB it may well re-use the same entries as already exist - perhaps the ROM monitor already initialized the TLB, and (derived from the same source code) happened to use the same dummy addresses. If you do that, your second initialization run will cause a machine check exception. The solution is for the initializing routine to check the TLB for a matching entry (using the **tlbp** instruction) before each update.

For portability you should probably include the probe step in initialization routines: it’s not essential on the 34K core or any machine conforming to the MIPS MT ASE, where we repeat that the machine check exception doesn’t happen unless both the old and new entry are both marked as valid.



## Kernel-mode (OS) programming

[MIPS32] tells you how to write OS code which is portable across all compliant CPUs. Most OS code should not be CPU-dependent, and we won't tell you how to write it here. But release 2 of the MIPS32 Specification [MIPS32] introduced a few new optional features which are not yet well known, so are worth describing here:

- A better way of managing software-visible pipeline and hardware delays associated with CP0 programming in Section 7.1, "Hazard barrier instructions".
- New interrupt facilities described in Section 7.2, "MIPS32® Architecture Release 2 - enhanced interrupt system(s)";
- The ability to use one or more extra sets of registers ("shadow sets") to reduce context-saving overhead in interrupt handlers, in Section 7.3, "Shadow registers".
- How to get at any power-saving features, in Section 7.4, "Saving Power"

### 7.1 Hazard barrier instructions

When privileged "CP0" instructions change the machine state, you can get unexpected behavior if an effect is deferred out of its normal instruction sequence. But that can happen because the relevant control register only gets written some way down the pipeline, or because the changes it makes are sensed by other instructions early in their pipeline sequence.

Traditionally, MIPS CPUs left the kernel/low-level software engineer with the job of designing sequences which are guaranteed to run correctly, usually by padding the dangerous operation with enough **nop** or **ssnop** instructions.

From Release 2 of the MIPS32 specification this is replaced by explicit *hazard barrier* instructions. If you execute a hazard barrier between the instruction which makes the change (the "producer") and the instruction which is sensitive to it (the "consumer"), you are guaranteed that the change will be seen as complete. Hazards can appear when the producer affects even the instruction fetch of the consumer - that's an "instruction hazard" - or only affecting the operation of the consuming instruction (an "execution hazard"). Hazard barriers come in two strengths: **ehb** deals only with execution hazards, while **eret**, **jr.hb** and **jalr.hb** are barriers to both kinds of hazard.

In most implementations the strong hazard barrier instructions are quite costly, often discarding most or all of the pipeline contents: they should not be used indiscriminately. For efficiency you should use the weaker **ehb** where it is enough. Since some implementations work by holding up execution of all instructions after the barrier, it's preferable to place the barrier just before the consumer, not just after the producer

For example you might be updating a TLB entry:

```
mtc Index, t0
# other stuff, if there's stuff to do
ehb
tlbwi
jr.hb ra
```

The **ehb** makes sure that the change to *Index* has been made before you attempt to write the TLB entry, which is fine. But updating the TLB might affect how instructions are fetched in mapped space, so you should not return to code which might be running in mapped space until you've cleared the "instruction hazard". That's dealt with by the **jr.hb**.

### Hazard barriers and multi-threading

Within a thread the hazard barriers work as advertised. But because TCs share many CP0 registers and other resources, some hazards can be between different threads - or more precisely, an instruction can produce some effect on other threads which affect the behavior of subsequent instructions.

In particular, the operations which disable other threads (instructions like **dmt** or **dvpe** or direct manipulation of the associated CP0 bits *VPEct[TE]* and *MVPCt[EVP]*, or writes to *TCHalt*) may not be immediate. Instructions after the other-thread-disable instruction in the stream might - according to the MT ASE specification [MIPSMT] - see evidence of other threads continuing to run for a while. The MT ASE defines this as an instruction hazard. However, no hazard of this kind exists in 34K family CPUs, so if you're prepared to make your software CPU-dependent you may make it a bit more efficient.

### Porting software to use the new instructions

If you know your software will only ever run on a MIPS32 Release 2 or higher CPU, then that's great. But to maintain software which has to continue running on older CPUs:

- *ehb is a no-op*: on all previous CPUs. So you can substitute an **ehb** for the last no-op in your sequence of "enough no-ops", and your software is now safe on all future CPUs which are compliant with Release 2.
- *jr.hb and jalr.hb*: are decoded as plain jump-register and call-by-register instructions on earlier CPUs. Again, provided you already had enough no-ops for your worst-case older CPU, your system should now be safe on Release 2 and higher CPUs.

## 7.2 MIPS32® Architecture Release 2 - enhanced interrupt system(s)

The features for handling interrupts include:

- **Vectored Interrupt (VI) mode** offers multiple entry points (one for each of the interrupt sources), instead of the single general exception entry point.

**External Interrupt Controller (EIC) mode** goes further, and reinterprets the six core interrupt input signals as a 64-value field - potentially 63 distinguished interrupts each with their own entry point (the zero code, of course, is reserved to mean "no interrupt active").

Both these modes need to be explicitly enabled by setting bits in the *Config3* register; if you don't do that, the CPU behaves just as the original (release 1) MIPS32 specification required.

- Shadow registers - alternate sets of registers, often reserved for interrupt handlers, are described in [Section 7.3, "Shadow registers"](#). Interrupt handlers using shadow registers avoid the overhead of saving and restoring user GPR values.
- New readable *Cause[TI]* and *Cause[PCI]* bits provide a direct indication of pending interrupts from the on-core timer and performance counter subsystems (these interrupts are potentially shared with other interrupt inputs,

and it previously required system-specific programming to discover the source of the interrupt and handle it appropriately).

The new interrupt options are enabled by the *IntCtl* register, whose fields are shown in Figure 7-1.

**Figure 7-1 Fields in the IntCtl register**

31	29	29	26	25	10	9	5	4	0
IPTI	IPPCI	0			VS	0			

Notes:

*IntCtl*[IPTI,IPPCI]: are read-only fields, telling you how timer and performance counter interrupts (generated inside the core) are wired up. It's relevant in non-vectorized and simple-vectorized ("VI") interrupt modes.

The timer and performance counter interrupts are taken out to the core interface, where they are generally sent back again down one of the interrupt signals. The SoC designer who wires up the interrupts is also supposed to hardware code values which turn into the *IntCtl*[IPTI,IPPCI] fields. Each is a 3-bit binary number identifying which CPU interrupt input is shared by the internal timer interrupt (*IPTI*) or the performance counter overflow interrupt (*IPPCI*).

The interrupt is specified by giving the number of the *Cause*[IPnn] where the resulting interrupt is seen. Because *Cause*[IP0-1] are software interrupt bits, unconnected to any input, legal values for *IntCtl*[IPTI] and *IntCtl*[IPPCI] are between 2 and 7.

The timer interrupt output is per-VPE, so there are two of them from the 34K core. The *IntCtl* register is also per-VPE, reflecting the local setup. The performance counter registers are not replicated (there's just one set per CPU).

*IntCtl*[VS]: is writable to give you software control of the vector spacing; the spacing you get between consecutive entries is *IntCtl*[VS]×32 bytes. Only values of 1, 2, 4, 8 and 16 work (to give spacings of 32, 64, 128, 256, and 512 bytes respectively). A value of zero does give a zero spacing, so all interrupts arrive at the same address.

### 7.2.1 Traditional MIPS interrupt signalling and priority

Before we discuss the new features, we should remind you what was there already. On traditional MIPS systems the CPU takes an interrupt exception on any cycle where one of the eight possible interrupt sources visible in *Cause*[IP] is active, enabled by the corresponding enable bit in *Status*[IM], and not otherwise inhibited. When that happens control is passed to the general exception handler (see Table C.4 for exception entry point addresses), and is recognized by the "interrupt" value in *Cause*[ExcCode]. All interrupt are equal in the hardware, and the hardware does nothing special if two or more interrupts are active and enabled simultaneously. All priority decisions are down to the software.

Six of the interrupt sources are hardware signals brought into the CPU, while the other two are "software interrupts" taking whatever value is written to them in the *Cause* register.

The original MIPS32 specification adds an option to this. If you set the *Cause*[IV] bit, the same priority-blind interrupt handling happens but control is passed to an interrupt exception entry point which is separate from the general exception handler.

## 7.2.2 VI mode - interrupt signalling and priority

The traditional interrupt system fits with a RISC philosophy (it leaves all interrupt priority policy to software). It's also OK with complex operating systems, which commonly have a single piece of code which does the housekeeping associated with interrupts prior to calling an individual device-interrupt handler.

A single entry point doesn't fit so well with embedded systems using very low-level interrupt handlers to perform small near-the-hardware tasks. So Release 2 of the MIPS32 architecture adds "VI interrupt mode" where interrupts are despatched to one of eight possible entry points. To make this happen:

1. *Config3[VInt]* must be set, to indicate that your core has the vectored-interrupts feature - but all cores in the 34K family have it;
2. You write *Cause[IV]* = 1 to request that interrupts use the special interrupt entry point; and:
3. You set *IntCtl[VS]* non-zero, setting the spacing between successive interrupt entry points.

Then interrupt exceptions will go to one of eight distinct entry points. The bit-number in *Cause[IP]* corresponding to the highest-numbered active interrupt becomes the "vector number" in the range 0-7. The vector number is multiplied by the "spacing" implied by the OS-written field *IntCtl[VS]* (see above) to generate an offset. This offset is then added to the special interrupt entry point (already an offset of 0x200 from the value defined in *EBase*) to produce the entry point to be used.

If multiple interrupts are active and enabled, the entry point will be the one associated with the higher-numbered interrupt: in VI mode interrupts are no longer all equal, and the hardware now has some role in interrupt "priority".

## 7.2.3 External Interrupt Controller (EIC) mode

Embedded systems have lots of interrupts, typically far exceeding the six input signals traditionally available. Most systems have an external interrupt controller to allow these interrupts to be masked and selected. If your interrupt controller is "EIC compatible" and you use these features, then you get 63 distinct interrupt entry points.

To do this the same six hardware signals used in traditional and VI modes are redefined as a bus with 64 possible values<sup>1</sup>: 0 means "no interrupt" and 1-63 represent distinct interrupts. That's "EIC interrupt mode", and you're in EIC mode if you would be in VI mode (see previous section) and additionally the *Config3[VEIC]* bit is set. EIC mode is a little deceptive: the programming interface hardly seems to change, but the meaning of fields change quite a bit.

Firstly, once the interrupt bits are grouped the interrupt mask bits in *Status[IM]* can't just be bitwise enables any more. Instead this field (strictly, the 6 high order bits of this field, excluding the mask bits for the software interrupts) is recycled to become a 6-bit *Status[IPL]* ("interrupt priority level") field. Most of the time (when running application code, or even normal kernel code) *Status[IPL]* will be zero; the CPU takes an interrupt exception when the interrupt controller presents a number higher than the current value of *Status[IPL]* on its "bus" and interrupts are not otherwise inhibited.

As before, the interrupt handler will see the interrupt request number in *Cause[IP]* bits - see [Section C-2, "Fields in the Cause register"](#); the six MS of those bits are now relabelled as *Cause[RIPL]* ("requested IPL"). In EIC mode the software interrupt bits are not used in interrupt selection or prioritization: see below. But there's an important difference; *Cause[RIPL]* holds a snapshot of the value presented to the CPU when it decided to take the interrupt, whereas the old *Cause[IP]* bits simply reflected the real-time state of the input signals<sup>2</sup>.

1. The resulting system will be familiar to anyone who's used a Motorola 68000 family device (or further back, a DEC PDP/11 or any of its successors).

When an exception is triggered the new IPL - as captured in *Cause[RIPL]* - is used directly as the interrupt number; it's multiplied by the interrupt spacing implied by *IntCtl[RS]* and added to the special interrupt entry point, as described in the previous section. *Cause[RIPL]* retains its value until the CPU next takes any exception.

**Software interrupts:** the two bits in *Cause[IP1-0]* are still writable, but now become real signals which are fed out of the CPU core, and in most cases will become inputs - presumably low-priority ones - to the EIC-compliant interrupt controller.

In EIC mode the usual association of the internal timer and performance-counter overflow interrupts with individual bits of *Cause[IP]* is lost. Timer and performance counter interrupts are turned into output signals from the core, and will themselves become inputs to the interrupt controller. Ask your system integrator how they are wired.

## 7.3 Shadow registers

In hardware terms, shadow registers are deceptively simple: just add one or more extra copies of the register file. If you can automatically change register set on an exception, the exception handler will run with its own context, and without the overhead of saving and restoring the register values belonging to the interrupted program. On to the details...

MIPS shadow registers come as one or more extra complete set of 32 general purpose registers. The CPU only changes register sets on an exception or when returning from an exception with **eret**.

In the 34K core (and possibly other CPUs conforming to [MIPSMT]) there are no dedicated shadow registers, but you can configure the CPU to make the registers of one or more TCs available as shadow sets, as described in Section 7.3.1.

### Selecting shadow sets - SRSCtl

The shadow set selectors are in the *SRSCtl* register, shown in Figure 7-2.

**Figure 7-2 Fields in the SRSCtl register (shadow register set control)**

31	30	29	26	25	22	21	18	17	16	15	12	11	10	9	6	5	4	3	0
0	HSS	0	EICSS	0	ESS	0	PSS	0	CSS										

In *SRSCtl*:

*SRSCtl[HSS]*: read-only field showing the highest-numbered register set available on this VPE/CPU. The ordinary register set is set #0, so this is the number of available register sets minus one.

On single-threaded CPUs this field is fixed. However, on the 34K core this field can change when configuration software - that is, when *VPEConf0[VPC]* is set - changes the way the shadow sets are shared. See Section 7.3.1 below for how multithreading TCs can be used as shadow sets.

*SRSCtl[CSS]*: the register set currently in use. It's read-only here; set on any exception, replaced by the value in *SRSCtl[PSS]* on an **eret**.

*SRSCtl[ESS]*: this writable field is the software-selected register set to be used for "all other" exceptions; that's other than an interrupt in VI or EIC mode (both have their own special ways of selecting a register set).

- Since the incoming IPL can change at any time - depending on the priority views of the interrupt controller - this is essential if the handler is going to know which interrupt it's servicing.

## Kernel-mode (OS) programming

*SRSCtl[PSS]*: the "previous" register set, which will be used following the next **eret**.

You can get at the values of registers in this set using **rdpgpr** and **wrpgpr**.

*SRSCtl[PSS]* is writable, allowing the OS to dispatch code in a new register set; load this value and then execute an **eret**.

*SRSCtl[EICSS]*: will be explained in the [next section](#).

Just a note: *SRSCtl[PSS]* and *SRSCtl[CSS]* are not updated by *all* exceptions, but only those which write a new return address to *EPC* (or equivalently, those occasions where the exception level bit *Status[EXL]* goes from zero to one). Exceptions where *EPC* is *not* written include:

- Exceptions occurring with *Status[EXL]* already set;
- Cache parity error exceptions, where the return address is loaded into *ErrorEPC*;
- EJTAG debug exceptions, where the return address is loaded into *DEPC*.

### How new shadow sets get selected on an interrupt

In EIC mode, the external interrupt controller proposes a shadow register set number with each requested interrupt (nonzero IPL). When the CPU takes an interrupt, the externally-supplied set number determines the next set and is made visible in *SRSCtl[EICSS]* until the next interrupt.

In VI mode (no external interrupt controller) the core sees only eight possible interrupt numbers; the *SRSMap* register contains eight 4-bit fields, defining the register set to use for each of the eight interrupt levels, as shown in [Figure 7-3](#).

**Figure 7-3 Fields in the SRSMap register**

31	28	27	24	23	20	19	16	15	12	11	8	7	4	3	0
SSV7	SSV6	SSV5	SSV4	SSV3	SSV2	SSV1	SSV0								

In *SRSMap*, each of the *SSV7-0* fields has the shadow set number to be used when handling the interrupt for the corresponding *Cause[IP7-0]* bit. A zero shadow set number means not to use a shadow set.

If you are remaining with "classic" interrupt mode, it's still possible to use one shadow set for all exception handlers - including interrupt handlers - by setting *SRSCtl[ESS]* non-zero.

### Software support for shadow registers

Shadow registers work "as if by magic" for short interrupt routines which run entirely in exception mode (that is, with *Status[EXL]* set). The shadow registers are not just efficient because there's no need to save user registers; the shadow registers can also be used to hold contextual information for one or more interrupt routines which uses a particular shadow set. For more ambitious interrupt nesting schemes, software must save and stack copies of *SRSCtl[PSS]* alongside its copies of *EPC*; and it's entirely up to the software to determine when an interrupt handler can just go ahead and use a register set, and when it needs to save values on entry and restore them on exit. That's at least as difficult as it sounds: shadow sets are probably best used purely for very low-level, high-speed handlers.

### 7.3.1 Recycling multi-threading CPU's TCs as shadow sets

This recycling is controlled by some TC control bits and the *SRSCConf0-4* registers.

**Figure 7-4 Fields in the SRSCConf0 register**

31	30	29	20	19	10	9	0
M	0	SRS3	SRS2	SRS1			

In *SRSCConf0*:

*M*: is a "continuation" indication. Since there is no *SRSCConf1* in the 34K core, it will read zero.

In general there need be no more of these registers than are required to map your core's maximum complement of shadow register sets.

*SRS1-3*: are each set to the GPR set to be used for the putative shadow set number (1-3).

Shadow set 0 refers (in a MIPS MT CPU) to the register set normally associated with the current TC.

A value of all-ones in any of the (10-bit) *SRS1-3* fields (decimal 1023) indicates that this shadow set number is not usable - it won't select a set of registers.

The fact that there are no more "SRSCConf" registers means that shadow set numbers above 4 are never usable for the 34K core.

These fields may be writable (waiting to receive the number of a TC you sacrifice to provide a shadow set) or hard-wired (representing dedicated shadow register sets, whose "GPR number" will be larger than the maximum TC# of the machine.)

From reset, the writable fields take the value 1022. You just write the number of the TC you're sacrificing. Unless the donor TC is already bound to the same VPE as owns this *SRSCConf* register, nothing happens. You should also make sure the donor TC is halted, inactive and not usable by fork.

It's possible to reverse this process and seize back a TC, so long as the shadow set concerned is no longer in use.

Note that *SRSCConf0* is replicated per-VPE.

## 7.4 Saving Power

There are basically just a couple of facilities:

- The **wait** instruction: this puts the thread running to sleep. When this happens when all other threads are sleeping, halted or suspended, the core goes into a low-power mode with many clocks stopped, from which it will only emerge when it senses an interrupt. The interrupt will be delivered to any sleeping thread, but *all* sleeping threads will wake and return from their **wait**. That will usually be OK; it's normal practice to loop over **wait**.
- The *Status[RP]* bit: this doesn't do anything inside the core, but its state is made available at the core interface as *SI\_RP*. Logic outside the core is encouraged to use this to control any logic which trades off power for speed - most often, that will be slowing the master clock input to the CPU.





## 34K™ core features for debug and profiling

In this chapter you'll find:

- [Section 8.1, "EJTAG on-chip debug unit"](#)
- [Section 8.2, "PDtrace™ instruction trace facility"](#)
- [Section 8.3, "CP0 Watchpoints"](#) - monitor code and data access without using EJTAG.
- [Section 8.4, "Performance counters"](#) - gather statistics about events, useful for understanding where your program spends its time.

The description here is terse and leaves out some information about EJTAG and PDtrace facilities which is not visible to programmers. We will document it here if it's software visible, or is implementation-dependent information not found in the detailed manuals ([\[EJTAG\]](#), [\[PDTRACEUSAGE\]](#) and [\[PDTRACETCB\]](#))

### 8.1 EJTAG on-chip debug unit

This is a collection of in-CPU resources to support debug. Debug logic serves no direct purpose in the final end-user application, so it's always under threat of being omitted for cost reasons. A debug unit must have virtually no performance impact when not in use; it must use few or no dedicated package pins, and should not increase the logic gate count too much. EJTAG solves the pin issue (and gets its name) by recycling the JTAG pins already included in every SoC for chip test<sup>1</sup>.

So the debug unit requires:

- Physical communications with some kind of "probe" device (which is itself controlled by the debug host), achieved through the JTAG pins.
- The ability for a probe to "remote-control" the CPU. The basic trick is to get the CPU to execute instructions that the probe supplies. In turn that's done by directing the CPU to execute code from the magic "dmseg" region where CPU reads and writes are made down the wire to the probe. "dmseg" is itself a part of "dseg", see [Section 8.1.5, "The "dseg" memory decode region"](#).
- A distinguished debug exception. In MIPS EJTAG, this is a special "super-exception" marked by a special debug-exception-level flag, so you can use an EJTAG debugger even on regular exception handler code. See [Section 8.1.2, "Debug mode"](#) below;
- A number of "hardware breakpoints". Their numerous control registers can't be accommodated in the CP0 register set, so are memory mapped into "dseg";

---

1. It can actually be quite useful to provide EJTAG with its own pins, if your package permits.

- You can take a debug exception from a special breakpoint instruction **sddbnp**, on a match from an EJTAG hardware breakpoint, after an EJTAG single-step, when the probe writes the break bit *EJTAG\_CONTROL[EjtagBrk]*, or by asserting the external *DINT* (debug interrupt) signal.
- You can configure your hardware to take periodic snapshots of the address of the currently-executing instruction (“PC sampling”) and make those samples available to an EJTAG probe, as described in the next section.

On these foundations powerful debug facilities can be built.

The multi-vendor [EJTAG] specification has many independent options, but MIPS Technologies cores tend to have fewer options and to implement the bulk of the EJTAG specification. The 34K core can be configured by your SoC designer with either four instruction breakpoints (or none), and with two data breakpoints (or none). It is also optional whether the dedicated debug-interrupt signal *DINT* is available in your SoC.

### 8.1.1 Debug communications through JTAG

The chip’s JTAG pins give an external probe access to a special registers inside the core. The JTAG standard defines a serial protocol which lets the probe run one of a number of JTAG “instructions”, each of which typically reads/writes one of a number of registers. EJTAG’s instructions are shown in [Table 8.1](#).

**Table 8.1 JTAG instructions for the EJTAG unit**

<i>JTAG "Instruction"</i>	<i>Description</i>
IDCODE	Reads out the MIPS core and revision - not very interesting for software, not described further here.
ImpCode	Reads bit-field showing what EJTAG options are implemented - see <a href="#">Figure 8-5</a> below.
EJTAG_ADDRESS	(read/write) together, allow the probe to respond to instruction fetches and data reads/writes in the magic “dseg” region described in <a href="#">Section 8.1.5</a> , “The “dseg” memory decode region”.
EJTAG_DATA	
EJTAG_CONTROL	Package of flags and control fields for the probe to read and write; see <a href="#">Figure 8-6</a> below.
EJTAGBOOT	The “EJTAGBOOT” instruction causes the next CPU reset to lead to CPU booting from probe; see description of the <i>EJTAG_CONTROL</i> bits <i>ProbEn</i> , <i>ProbTrap</i> and <i>EjtagBrk</i> in the notes <a href="#">Figure 8-6</a> .
NORMALBOOT	The “NORMALBOOT” instruction reverts to the normal CPU bootstrap.
FASTDATA	Special access used to accelerate multi-word data transfers with probe. The probe reads/writes the 33-bit register formed of <i>EJTAG_CONTROL[PrAcc]</i> with <i>EJTAG_DATA</i> .
TCBCONTROLA	Access registers used to control “PDtrace” instruction trace output, if available. See <a href="#">Section 8.2.1</a> , “34K core-specific fields in PDtrace™ JTAG-accessible registers” - only the core-specific fields in these registers are documented here.
TCBCONTROLB	
TCBCONTROLC	
TCBADDRESS	
PCSAMPLE	Access register which holds PC sample value, see <a href="#">Section 8.1.12</a> , “PC Sampling with EJTAG”.

### 8.1.2 Debug mode

A special CPU state; the CPU goes into debug mode when it takes any debug exception - which can be caused by an **sddbnp** instruction, a hit on an EJTAG breakpoint register, from the external “debug interrupt” signal *DINT*, or single-stepping (the latter is peculiar and described briefly below). Debug mode state is visible as *Debug[DM]* (see [Figure 8-1](#) below). Debug mode (like exception mode, which is similar) disables all normal interrupts. The address map changes in debug mode to give you access to the “dseg” region, described below. Quite a lot of exceptions just won’t happen in debug mode: those which do, run peculiarly - see the relevant paragraphs in [Section 8.1.2](#), “Debug mode”.

A CPU with a suitable probe attached can be set up so the debug exception entry point is in the “dmseg” region, running instructions provided by the probe itself. With no probe attached, the debug exception entry point is in the ROM - see Section C.4, “Exception entry points”.

### 8.1.3 The debug unit and multi-threading

The software-visible resources of the EJTAG unit are replicated per VPE, and each VPE has its own distinct JTAG “tap”. Just two bits are replicated per-TC: *Debug[SSSt]* controls the single-step exception, and *Debug[OffLine]* provides a debugger with a way of controlling exactly which TCs run in between breakpoints of a debug session.

When any TC executes in debug mode, all other TCs (even in other VPEs) are suspended. There is nothing software can do to prevent a debug-mode TC from issuing instructions: it runs regardless of the state of *TCStatus[A]*, *TCHalt*, the *VPEControl[TE]* bit set by **dmr**, the *MVPControl[EVP]* bit set by **dvpe**, the *VPEConf0[VPA]* bit, or even the debugger’s own *Debug[OffLine]*. However, when you return from debug mode with a **deret** and one of these software inhibit bits is active, the TC will not execute any non-debug-mode instruction.

When you execute a debug breakpoint (**sdbbp**) instruction or hit a synchronous (address-testing only) breakpoint, the debug exception will be handled by the TC which ran the exception-causing instruction. But an asynchronous entry into debug mode caused by the assertion of *DINT* or hitting a data-testing breakpoint may use any TC affiliated with the VPE which owns the signal or set the breakpoint: and again, this TC is chosen regardless of its software-settable state, so you are guaranteed that the debug condition will be serviced.

When any TC is already executing in debug mode *DINT* (even if directed at another VPE) is ignored.

For non-debug code some MT facilities are protected by “safety catch” control bits. Debug-mode code is all-powerful, as if *VPEConf0[MVP]* was set.

#### Exceptions in debug mode

Software debuggers will probably be coded to avoid causing exceptions (testing addresses in software, for example, rather than risking address or TLB exceptions).

While executing in debug mode many conditions which would normally cause an exception are ignored: interrupts, debug exceptions (other than that caused by executing **sdbbp**), and CP0 watchpoint hits.

But other exceptions are turned into “nested debug exceptions” when the CPU is in debug mode - a facility which is probably mostly valuable to debuggers using the EJTAG probe.

On such a nested debug exception the CPU jumps to the debug exception entry point, remaining in debug mode. The *Debug[DExcCode]* field records the cause of the nested exception, and *DEPC* records the debug-mode-code restart address. This will not be survivable for the debugger unless it saved a copy of the original *DEPC* soon after entering debug mode, but it probably did that! To return from a nested debug exception like this you don’t use **deret** (which would inappropriately take you out of debug mode), you grab the address out of *DEPC* and use a jump-register.

### 8.1.4 Single-stepping

When control returns from debug mode with a **deret** and the (per-TC) single-step bit *Debug[SSSt]* is set, the instruction selected by *DEPC* will be executed in non-debug context<sup>1</sup>; then a debug exception will be taken on the thread’s very next instruction in sequence.

1. If *DEPC* points to a branch instruction, both the branch and branch-delay instruction will be executed normally.

Since at least one instruction is run in normal mode it can lead to a non-debug exception; in that case the “very next instruction in sequence” will be the first instruction of the exception handler, and you’ll get a single-step debug exception whose *DEPC* points at the exception handler.

In a multithreaded CPU any number of instructions from other threads might run before you get the single-step exception. A debugger wanting to avoid that can use the various TC’s *Debug[OffLine]* controls to inhibit TCs other than the one under debug,

### 8.1.5 The “dseg” memory decode region

EJTAG needs to use memory space both to accommodate lots of breakpoint registers (too many for CP0) and for its probe-mapped communication space. This memory space pops into existence at the top of the CPU’s virtual address map when the CPU is in debug mode, as shown in [Table 8.2](#).

Table 8.2 EJTAG debug memory region map ("dseg")

Virtual Address	Region/sub-regions	Location/register	Virtual Address	
0xE000.0000	kseg2		0xE000.0000	
0xFF1F.FFFF			0xFF1F.FFFF	
0xFF20.0000	dseg	dmseg	fastdata	0xFF20.0000
0xFF20.000F				0xFF20.000F
0xFF20.0010			0xFF20.0010	
0xFF20.0200		debug entry	0xFF20.0200	
0xFF2F.FFFF			0xFF2F.FFFF	
0xFF30.0000	drseg	DCR register	0xFF30.0000	
0xFF30.1000		IBS register	0xFF30.1000	
		<i>I-breakpoint #1 regs</i>		
0xFF30.1100		IBA1	0xFF30.1100	
0xFF30.1108		IBM1	0xFF30.1108	
0xFF30.1110		IBASID1	0xFF30.1110	
0xFF30.1118		IBC1	0xFF30.1118	
		<i>I-breakpoint #2 regs</i>		
0xFF30.1200		IBA2	0xFF30.1200	
0xFF30.1208		IBM2	0xFF30.1208	
0xFF30.1210		IBASID2	0xFF30.1210	
0xFF30.1218		IBC2	0xFF30.1218	
		<i>same for next two</i>		
		...		
0xFF30.2000		DBS register	0xFF30.2000	
		<i>D-breakpoint #1 regs</i>		
0xFF30.2100		DBA1	0xFF30.2100	
0xFF30.2108		DBM1	0xFF30.2108	
0xFF30.2110		DBASID1	0xFF30.2110	
0xFF30.2118		DBC1	0xFF30.2118	
0xFF30.2120		DBV1	0xFF30.2120	
0xFF30.2124		DBVHi1	0xFF30.2124	
		<i>D-breakpoint #2 regs</i>		
0xFF30.2200		DBA2	0xFF30.2200	
0xFF30.2208		DBM2	0xFF30.2208	
0xFF30.2210		DBASID2	0xFF30.2210	
0xFF30.2218		DBC2	0xFF30.2218	
0xFF30.2220		DBV2	0xFF30.2220	
0xFF30.2224		DBVHi2	0xFF30.2224	
0xFF30.2228			0xFF30.2228	
0xFFFF.FFFF			0xFFFF.FFFF	

Notes on Table 8.2:

- *dseg*: is the whole debug-mode-only memory area.

It's possible for debug-mode software to read the "kseg2"-mapped locations "underneath" by setting *Debug[LSNM]* (see Table 8-1 below).

## 34K™ core features for debug and profiling

- *dmseg*: is the memory region where reads and writes are implemented by the probe. But if no active probe is plugged in, or if *DCR[PE]* is clear, then accesses here cause reads and writes to be handled like regular "kseg3" accesses.
- *drseg*: is where the debug unit's main register banks are accessed. Accesses to "drseg" don't go off core. Registers in "drseg" are word-wide, and should be accessed only with 32-bit word-wide loads and stores.
- *fastdata*: is a corner of "dmseg" where probe-mapped reads and writes use a more JTAG-efficient block-mode probe protocol, reducing the amount of JTAG traffic and allowing for faster data transfer. There's no details about how it's done in this document, see [EJTAG].
- *debug entry*: is the debug exception entry point. Because it lies in "dmseg", the debug code can be implemented wholly in probe memory, allowing you to debug a system which has no physical memory reserved for debug.

### 8.1.6 EJTAG CP0 registers, particularly Debug

In normal circumstances (specifically, when not in debug mode), the only software-visible part of the debug unit is its set of three CP0 registers:

- *Debug* which has configuration and control bits, and is detailed below;
- *DEPC* keeps the restart address from the last debug exception (automatically used by the **deret** instruction);
- *DESAVE* is a CP0 register which is just 32-bits of read/write space. It's available for a debug exception handler which needs to save the value of a first general-purpose register, so that it can use that register as an address base to save all the others.

*Debug*, *DEPC* and *DESAVE* are replicated per-VPE, giving each VPE the impression of having its own EJTAG unit.

*Debug* is the most complicated and interesting. It has so many fields defined that we've taken them in three groups: debug exception cause bits in [Figure 8-2](#), information about regular exceptions which want to happen but can't because you're in debug mode in [Figure 8-3](#), and everything else. The "everything else" category includes the most important fields and comes first, in [Figure 8-1](#).

Figure 8-1 Fields in the EJTAG CP0 Debug register

31	30	29	28	27	26	25	24	21	20	19	18	17	15	14	10	9	8	7	6	5	0
DB D	D M	NoDC R	LSN M	Doz e	Hal t	CountD M	<i>pending</i> (Figure 8-3)	IEX I	<i>cause</i> (Table 8-2)	EJTAGv er	DExc- Code	NoS St	SSSt	OffLi ne	0	<i>cause</i> (Table 8-2)					

These fields are:

*DBD*: exception happened in branch delay slot. When this happens *DEPC* will point to the branch instruction, which is usually the right place to restart.

*DM*: debug mode - set on debug exception from user mode, cleared by **deret**.

Then some configuration and control bits:

*NoDCR*: read-only - 0 if there is a memory-mapped *DCR* register. MIPS Technologies cores will always have one. Any EJTAG unit implementing "dseg" at all implements *DCR*.

*LSNM*: Set this to 1 if you want debug-mode accesses to "dseg" addresses to be just sent to system memory. This makes most of the EJTAG unit's control system unavailable, so will probably only be done around a particular load/store.

*Doze*: before the debug exception, CPU was in some kind of reduced power mode.

*Halt*: before the debug exception, the CPU was stopped - probably asleep after a **wait** instruction.

*CountDM*: 1 if and only if the count register continues to run in debug mode. Writable for the 34K core, so you get to choose. On some other implementations it's read-only and just tells you what the CPU does.

*IEX*: set to 1 to defer imprecise exceptions. Set by default on entry to debug mode, cleared on exit, but writable. The deferred exception will come back when and if this bit is cleared: until then you can see that it happened by looking at the "pending" bits shown in Figure 8-3 below.

*EJTAGver*: read-only - tells you which revision of the specification this implementation conforms to. On the 34K core it reads 3 for version 3.1. The full set of legal values are:

0	Version 2.0 and earlier
1	Version 2.5
2	Version 2.6
3	Version 3.1

*DExcCode*: Cause of any non-debug exception you just handled from within debug mode - following first entry to debug mode, this field is undefined. The value will be one of those defined for *Cause[ExcCode]*, as shown in Section C.3, "Exception Code values in Cause[ExcCode]".

*NoSSSt*: read-only - reads 0 because single-step is implemented (it always is on MIPS Technologies cores).

*SSSt*: set 1 to enable single-step.

*OffLine*: prevents a TC from running any instructions (except in debug mode, but then debug mode overrides *all* software inhibitions on thread scheduling). It's there for debuggers which may need to selectively stop some threads, and should not be used by application or OS code. This bit has to be replicated per-TC.

Figure 8-2 Exception cause bits in the debug register

31	20	19	18	17	6	5	4	3	2	1	0
<i>Debug</i>		DDBSImpr	DDBLImpr	DINT		DIB	DDBS	DDBL	DBp	DSS	

*DDBSImpr*: imprecise store breakpoint - see Section 8.1.11, "Imprecise debug breaks" below. *DEPC* probably points to an instruction some time later in sequence than the store which triggered the breakpoint. The debugger or user (or both) have to cope as best they can.

*DDBLImpr*: imprecise load breakpoint. (See note on imprecise store breakpoint, above).

*DINT*: debug interrupt: either the *DINT* signal got asserted or the probe wrote *EJTAG\_CONTROL[EjtagBrk]* through the JTAG signals.

*DIB*: instruction breakpoint. If *DBp* is clear, that must have been from an *sddb*.


*DDBS*: precise store breakpoint.

*DDBL*: precise load breakpoint.

*DBp*: any sort of match with a hardware breakpoint.

*DSS*: single-step exception.

**Figure 8-3 Debug register - exception-pending flags**



These note exceptions caused by instructions run in debug mode, but which have not happened yet because they are imprecise and *Debug[|EXI]* is set. They remain set until *Debug[|EXI]* is cleared explicitly or implicitly by a **deret**, when the exception is delivered and the pending bit cleared:

*IBusEP*: bus error on instruction fetch pending. This exception is precise on the 34K core, so this can't happen and the field is always zero.

*MCheckP*: machine check pending (usually an illegal TLB update). As above, the machine check is always precise on the 34K core, so this is always zero.


*CacheEP*: cache parity error pending.

*DBusEP*: bus error on data access pending.

### 8.1.7 The DCR (debug control) memory-mapped register

This is the only memory-mapped EJTAG register apart from the hardware breakpoints (described in the next section). It's found in "drseg" at location 0xFF30.0000 as shown in Table 8.2 (but only accessible if the CPU is in debug mode). The fields are in Figure 8-4:

**Figure 8-4 Fields in the memory-mapped DCR (debug control) register**



Where:

*ENM*: (read only) reports CPU endianness (1 == big).

*DB/IB*: (read only) 1 if data/instruction hardware breakpoints are available, respectively. The 34K core has either 0 or 2 data breakpoints, and either 0 or 4 instruction breakpoints.

*PCS, PCR*: *PCS* reads 1 if the PC sampling feature is available, as it can be on the 34K core. Then *PCR* is a three-bit field defining the sampling frequency as one sample every  $2^{(5+PCR)}$  cycles. See Section 8.1.12, "PC Sampling with EJTAG" for details.



*INTE/NMIE*: set *DCR[INTE]* zero to disable interrupts in non-debug mode (it's a separate bit from the various non-debug-mode visible interrupt enables). The idea is that the debugger might want to step through kernel code or run kernel subroutines (perhaps to discover OS-related information) without losing control because interrupts start up again.

*DCR[NMIE]* masks non-maskable interrupt in non-debug mode (a nice paradox). Both bits are "1" from reset.

*NMIP*: (read-only) tells you that a non-maskable interrupt is pending, and will happen when you leave debug mode (and according to *DCR[NMIE]* as above).

*SRE*: if implemented, write zero to mask soft-reset causes. This signal has no effect inside the 34K core but is presented at the interface, where customer reset logic could be influenced by it.

*PE*: (read only) software-readable version of the probe-controlled enable bit *EJTAG\_CONTROL[ProbEn]*, which you could look at in [Figure 8-6](#).

### 8.1.8 JTAG-accessible registers

We're wandering away from what is relevant to software here: this register is available for read and write only by the probe, and is not software-accessible.

But you can't really understand the EJTAG unit without knowing what dials, knobs and switches are available to the probe, so it seems easier to give a little too much information.

First of all there are two informational fields provided to the probe, *IDCODE* (just reflects some inputs brought in to the core by the SoC team, not very interesting) and *ImpCode* ([Figure 8-5](#)); then there's the main CPU interaction control/status register *EJTAG\_CONTROL* ([Figure 8-6](#)).

**Figure 8-5 Fields in the JTAG-accessible ImpCode register**

31	29	28	25	24	23	21	20	17	16	15	14	13	1	0
EJTAGver	0	DINTsup	ASIDsize	0	MIPS16	0	NoDMA	0	MIPS32/64					
2 = 2.6		see note	see note		1		1							0

Notes on the *ImpCode* fields:

*EJTAGver*: same value (and meaning) as the *Debug[EJTAGver]* field, see the notes on [Figure 7-2](#).

*DINTsup*: whether JTAG-connected probe has a *DINT* signal to interrupt the CPU. Configured by your SoC designer (who should know) by hard-wiring the core interface signal *EJ\_DINTsup*.

The probe can always interrupt the CPU by a JTAG command using the *EJTAG\_CONTROL[EjtagBrk]*, but *DINT* is much faster, which is useful if you're cross-triggering one piece of hardware from another. However, it is fed to both VPEs at once, and it's unpredictable which of them will take the resulting debug exception (only one can).

*ASIDsize*: usually 2 (indicating the 8-bit *EntryHi[ASID]* field size required by the MIPS32 standard), but can be 0 if your core has been built with the no-TLB option (i.e. a fixed-mapping MMU).

*MIPS16*: 1 because the 34K core always supports the MIPS16 instruction set extension.

*NoDMA*: 1 - MIPS Technologies cores do not provide EJTAG "DMA" (which would allow a probe to directly read and write anything attached to the 34K core's OCP interface).

*MIPS32/64*: the zero indicates this is a 32-bit CPU.

*Rocc*: "reset occurred" - reads 1 while a reset signal is applied to the CPU - and then the 1 value persists until overwritten with a zero from the JTAG side. Until the probe reads this as zero most of the other fields are nonsense.

The EJTAG\_CONTROL register is shown in Figure 8-6:

**Figure 8-6 Fields in the JTAG-accessible EJTAG\_CONTROL register**

31	30	29	28	23	22	21	20	19	18	17	16	15	14	13	12	11	4	3	2	0
Rocc	Psz	0	Doze	Halt	PerRst	PRn W	PrAcc	0	PrRst	ProbEn	Prob- Trap	0	Ejtag- Brk	0	D M	0				

Notes on the fields:

*Psz*: (read-only) when software reads or writes "dmseg" this tells the probe whether it was a word, byte or whatever-size transfer:

<i>Byte-within-word address</i>	<i>Size code</i>	<i>Transfer Size</i>
EJTAG_ADDRESS[1-0]	EJTAG_CONTROL[Psz]	
¥	0	Byte
00	1	Halfword
10		
00	2	Word
00	3	Tri-byte (lowest address 3 bytes)
01		Tri-byte (highest address 3 bytes)

*Doze/Halt*: (read-only) indicates CPU not fully awake. *Doze* reflects any reduced-power mode, whereas *Halt* is set only if the CPU is asleep after a **wait** or similar.

*PerRst*: write to set the *EJ\_PerRst* output signal from the core, which can be used to reset non-core logic (ask your SoC designer whether it's connected to anything).

For this and all other fields which change core state, we recommend that the probe should write the field and then poll for the change to be reflected in this register, which may take a short while. In some cases the bit is just an output one, when the readback will be pointless (but harmless).

*PRnW/PrAcc*: *PrAcc* is 1 when the CPU is doing a read/write of the "dmseg" region, and the probe should service it. The "slow" read/write protocol involves the probe flipping this bit back to zero to tell the CPU the transfer is ready.

While *PrAcc* is active the read-only *PRnW* bit distinguishes writes (1) from reads (0).

*PrRst*: controls the *EJ\_PrRst* signal from the core, which may be wired back to reset the CPU and related logic. Write a 1 to reset. If it works, the probe will eventually see the bit fall back to 0 by itself, as the CPU resets. Most probes are wired up with a direct CPU reset signal, which is more reliable.

*ProbEn, ProbTrap, EjtagBrk*: *ProbEn* must be set before CPU accesses to "dmseg" will be sent to the probe. It can be written by the probe directly. *ProbTrap* relocates the debug exception entry point from 0xBFC0.0480 (when 0) to the "dmseg" location 0xFF20.0200 - required when the debug exception handler itself is supplied by the probe.

*EjtagBrk* can be written 1 to "interrupt" the CPU into debug mode.

The three come together into a trick to support systems wanting to boot from EJTAG. The value of all these three bits is preset by the "EJTAGBOOT" JTAG instruction. When the CPU resets with all of these set to 1, then the CPU will immediately enter debug mode and start reading instructions from the probe.

*DM*: (read-only) indicates the CPU is in debug mode, a probe-readable version of *Debug[DM]*.

### 8.1.9 EJTAG breakpoint registers

It's optional whether the 34K core has EJTAG breakpoint registers. But if it has instruction breakpoints, it has four of them; and if it has data breakpoints, it has two. The breakpoints:

- Work only on virtual addresses, not physical addresses. However, you can restrict the breakpoint to a single address space by specifying an "ASID" value to match. Debuggers will need the co-operation of the OS to get this right.
- Use a bit-wise address mask to permit a degree of fuzzy matching.
- On the data side, you can break only when a particular value is loaded or stored. However, such breakpoints are imprecise in a CPU like the 34K core - see [Section 8.1.11, "Imprecise debug breaks"](#) below.

There are instruction-side and data-side breakpoint status registers (they're located in "drseg", accessible only when in debug mode, and their addresses are in [Section 8.2, "EJTAG debug memory region map \("dseg"\)"](#).) They're called *IBS* and *DBS*. The latter has, in theory, two extra fields (bits 29-28) used to flag implementations which can't do a load/store break conditional on the data value. However, MIPS cores with hardware breakpoints always include the value check, so these bits read zero anyway. So the registers are as shown in [Figure 8-7](#).

**Figure 8-7 Fields in the IBS/DBS (EJTAG breakpoint status) registers**

	31	30	29	28	27	24	23	4	3	2	1	0
<i>DBS</i>	0	ASID-sup	0	BCN = 2			0				BS1-0	
<i>IBS</i>				BCN = 4			0			BSD3-0		

Where:

*ASIDsup*: is 1 if the breakpoints can use ASID matching to distinguish addresses from different address spaces; on the 34K core that's available if and only if a TLB is fitted.

*BCN*: the number of hardware breakpoints available (two data, four instructions).

*BS1-0*, *BSD3-0*: bitfields showing breakpoints which have been matched. Debug software has to clear down a bit after a breakpoint is detected.

Then each EJTAG hardware breakpoint ("n" is 0-3 to select a particular breakpoint) is set up through 4-6 separate registers:

- *IBCn*, *DBCn*: breakpoint control register shown at [Figure 7-9](#) below;
- *IBAn*, *DBAn*: breakpoint address;
- *IBAMm*, *DBAMm*: bitwise mask for breakpoint address comparison. A "1" in the mask marks an address bit which will be *excluded from* comparison, so set this zero for exact matching.

Ingeniously, *IBAMm*[0] corresponds to the slightly-bogus instruction address bit zero used to track whether the CPU is running MIPS16 instructions, and allows you to determine whether an EJTAG I-breakpoint may apply only in MIPS16 (or non-MIPS16) mode.

- *IBASIDn*, *DBASIDn* specifies an 8-bit ASID, which may be compared against the current *EntryHi*[ASID] field to filter breakpoints so that they only happen to a program in the right "address space". The ASID check can be

enabled or disabled using *IBCn[ASIDuse]* or *DBCn[ASIDuse]* respectively - see Figure 7-9 and its notes below. ID (so that the break will only affect one Linux process, for example).


The higher 24 bits of each of these registers is always zero.

- *DBVn, DBVHin* the value to be matched on load/store breakpoints. *DBCHin* defines bits 63-32 to be matched for 64-bit load/stores: the 32-bit<sup>1</sup> 34K has 64-bit load/store instructions for floating point.

Note that you can disable data matching (to get an address-only data breakpoint) by setting the value byte-lane comparison mask *DBCn[BLM]* to all 1s.

So now let’s look at the control registers in Figure 8-8.

**Figure 8-8 Fields in the hardware breakpoint control registers (IBCn, DBCn)**



The fields are:

*ASIDuse*: set 1 to compare the ASID as well as the address.

*BAI7-0*: "byte (lane) access ignore"<sup>2</sup> - which sounds mysterious. But this is really an *address* filter.

When you set a data breakpoint, you probably want to break on any access to the data of interest. You don’t usually want to make the break conditional on whether the access is done with a load byte, load word, or even load-word-left: but the obvious way of setting up the address match for a breakpoint has that effect.

To make sure you catch any access to a location, you can use the address mask to disable sub-doubleword address matching and then use *DBCn[BAI]* to mark the bytes of interest inside the doubleword: well, except that zero bits mark the bytes of interest, and 1 bits mark the bytes to ignore (hence the mnemonic).

The *DBCn[BAI]* bits are numbered by the byte-lane within the 64-bit on-chip data bus; so be careful, the relationship between the byte address of a datum and its byte lane is endianness-sensitive.

*NoSB, NoLB*: set 0 to *enable*<sup>3</sup> breakpoint on store/load respectively.

*BLM7-0*: a per-byte mask for data comparison. A zero bit means compare this byte, a 1 bit means to ignore its value. Set this field all-ones to disable the data match.

*TE*: set 1 to use as trigger for "PDtrace" instruction tracing as described in Section 8.2, "PDtrace™ instruction trace facility" below.

*BE*: set 1 to activate breakpoint. This fields resets to zero, to avoid spurious breakpoints caused by random register settings: don’t forget to set it!

1. A JTAG hardware breakpoint for a real 64-bit CPU would have 64-bit *DBVn* registers, so wouldn’t need *DBVHin*.
2. Why are there 8 bytes, when the 34K core is a 32-bit CPU with only 32-bit general purpose registers? Well, the *DBCn[BAI]* and *DBCn[BLM]* fields each have a bit for each byte-lane across the data bus, and the 34K core has a 64-bit data bus (and in fact can do 64-bit load and store operations, for example for floating point values).
3. "1-to-enable" would feel more logical. The advantage of using 0-to-enable here is that the zero value means "break on either read or write", which is a better default than "never break at all".

### 8.1.10 Understanding breakpoint conditions

There are a lot of different fields and settings which are involved in determining when a hardware breakpoint detects its condition and causes an exception.

In all cases, there will be no break if you're in debug mode already... but then for a break to happen:

- *For all breakpoints including instructions:* all the following must be true:
  1. The breakpoint control register enable bit  $IBAn[BE]/DBAn[BE]$  is set.
  2. the address generated by the program for instruction fetch, load or store matches those bits of the breakpoint's address register  $IBAn/DBAn$  for which the corresponding address-mask register bits in  $IBAn/DBAn$  are zero.
  3. either  $IBCn[ASIDuse]/DBCn[ASIDuse]$  is zero (so we don't care what address space we're matching against), OR the address-space ID of the running program, i.e.  $EntryHi[ASID]$ , is equal to the value in  $IBASIDn/DBASIDn$ .

That's all for instruction breakpoints, but for data-side breakpoints also:

- *Data compare break conditions (not value related):* both the following must be true:
  4. It's a load and  $DBCn[NoLB]$  is zero, or it's a store and  $DBCn[NoSB]$  is zero.
  5. The load or the store touches at least one byte-within-doubleword for which the corresponding  $DBCn[BAI]$  bit is zero.

If you didn't want to compare the load/store value then  $DBCn[BLM]$  will be all-ones, and you're done. But if you also want to consider the value:

- *Data value compare break conditions:*
  6. the data loaded or stored, as it would appear on the system bus, matches the 64-bit contents of  $DBVHin$  with  $DBVn$  in each of those 8-bit groups for which the corresponding bit in  $DBCn[BLM]$  is zero.

That's it.

### 8.1.11 Imprecise debug breaks

Instruction breakpoints, and data breakpoints filtering only on address conditions are *precise*; that means that:

1.  $DEPC$  will point at the fetched or load/store instruction itself (except if it's in a branch delay slot, will point at the branch instruction);
2. The instruction will not have caused any side effects; in particular, the load/store will not reach the cache or memory.

Most exceptions in MIPS architecture CPUs are precise. But because of the way the 34K core optimizes loads and stores by permitting the CPU to run on at least until it needs to use the data from a load, data breakpoints which filter on the data value are *imprecise*. The debug exception will happen to whatever instruction (typically later in the

instruction stream) is running when the hardware detects the match, and not necessarily to the same TC. The debugging software must cope.

### 8.1.12 PC Sampling with EJTAG

A valuable trick available with recent revisions of the EJTAG specification and probes, "PC sampling" provides a non-intrusive way to collect statistical information about the activity of a running system. You can tell whether PC sampling is enabled by looking at *DCR[PCS]*, as shown in Figure 7-5 above.

The hardware snapshots the "current PC" (and the TC number of that instruction) periodically, and stores that value where it can be retrieved by a debug probe. It's then up to software to construct a histogram of samples over a period of time, which (statistically) allows a programmer to see where the CPU has spent most cycles. Not only is this useful, but it's also familiar: systems have used intrusive interrupt-based PC-sampling for many years, so there are tools which can readily interpret this sort of data.

When PC sampling is configured in to your core, it runs continuously. It doesn't even stop when the CPU is hanging on a **wait** instruction (time spent waiting is still time you might want to measure). You can choose to sample as often as once per 32 cycles or as rarely as once per 4096 cycles<sup>1</sup>; at every sampling point the address of the instruction completing in that cycle (or if none completes, the address of the next instruction to complete) is deposited in a JTAG-accessible register. Sampling rate is controlled by the *DCR[PCR]* field of the debug control register shown in Figure 7-5.

The hardware stores not only 32 bits of the instruction address, but also the then-current ASID (so you can interpret the virtual PC) and an always-written-1 "new" bit which a probe can use to avoid double-counting the same sample.

## 8.2 PDtrace™ instruction trace facility

An instruction trace is a set of data generated when a program runs which allows you to recreate the sequence of instructions executed, possibly with additional information included about data values. Instruction traces rapidly become enormous, and are typically generated in some kind of abbreviated form, which may be reconstructed by software which is in possession of a copy of the binary code of your system.

34K family cores can be configured with PDtrace logic, which provides a non-intrusive way of finding out what instructions your CPU ran. If your system includes PDtrace logic, *Config3[TL]* will read 1.

With a very high-speed CPU like the 34K core this is challenging, because you need to send data so fast. The PDtrace system deals with this by:

- *Compressing the trace*: a software tool in possession of the binary of your program can predict where execution will go next, following sequential instructions and fixed branches. To trace your program it needs only to know whether conditional branches were taken, and the destination of computed branches like jump-register.
- *Switching the trace on and off*: the 34K core can be configured with up to 8 "trace triggers", allowing you to start and stop tracing based on EJTAG breakpoint matches: see [Section 8.1.9, "EJTAG breakpoint registers"](#) above and [Table 8-14](#) below.
- *High-speed connection to a debug/trace probe*: optional. But if fitted, it uses advanced signalling techniques to get trace data from the CPU core, out of dedicated package pins to a probe. Good probes have generous amounts of high-speed memory to store long traces.

---

1. Since it runs continuously, it's a good thing that from reset the sampling period defaults to its maximum.

*TraceControl2[ValidModes,TBI,TBU]* (described below at Figure 7-10 and following) tell you whether you have such a connection available on your core. You'll have to ask the hardware engineers whether they brought out the connector, of course.

- *Very high-speed on-chip trace memory:* if fitted, you may find between 256bytes and 8Mbytes of trace memory in your system (larger than a few Kbytes is unlikely). Again, see *TraceControl2[ValidModes,TBI,TBU]* to find out what facilities you have.
- *Option to slow the CPU to match the tracing speed:* when you really, really need a full trace, and are prepared to slow down your program if necessary to wait while the trace information is sent to the probe. This is controlled by *TraceControl[IO]*, see below.

In practice the PDtrace logic depends on the existence of an EJTAG unit (described in the previous section) and an enhanced EJTAG probe. To benefit from on-probe trace memory, the probe will need to attach to PDtrace-specific signals.

This manual describes only the lowest-level building blocks as visible to software. For real hardware information refer to [PDTRACETCB]; for guidance about how to use the PDtrace facilities for software development see [PDTRACEUSAGE]. To use PDtrace facilities, you'll have to read the software manuals which come with a probe.

### 8.2.1 34K core-specific fields in PDtrace™ JTAG-accessible registers

The PDtrace system is controlled by the JTAG-accessible registers TCBCONTROLA and TCBCONTROLB. They are not visible to software running on the CPU, but we'll document fields and configured values which are specific to 34K family CPUs.

**Figure 8-9 Fields in the TCBCONTROLA register**

31	26	25	24	23	22	20	19	18	17	16	15	14	13	12		5	4	3	2	1	0
0	VModes 1 0		ADW	SyP	TD	IO	D	E	S	K	U	ASID				G	TFCR	TLSM	TIM	On	

In TCBCONTROLA:

*VModes:* reads "1 0", showing that 34K family cores support all tracing modes.

*ADW:* reads "1" to indicate that we support the wide (32-bit) internal trace bus.

**Figure 8-10 Fields in the TCBCONTROLB register**

31	30	28	27	26	25		21	20	19	17	16	15	14	13	12	11	10	8	7	6		3	2	1	0
WE	0	TWSrcWidth		REG	WR	0	RM	TR	BF	TM		TLSIF	CR	Cal	TWSrcVal	CA	OfC	EN							

In TCBCONTROLB:

*TWSrcWidth:* "0 1", which indicates that a 2-bit "source" field is included in the trace word to identify the VPE running the instruction, just as a multicore system would identify the CPU.

*TWSrcVal:* becomes writable, so the probe can set this value to a distinguishable one for each VPE."

**Figure 8-11 Fields in the TCBCONTROLC register**

31	28	27		23	22	21		14	13	12		5	4		2		1		0
0	Mode		CPUvalid	CPUId		TCvalid	TCnum	TCbits		MTtraceType	MTtraceTC								

TCBCONTROLC contains new fields for multi-threading trace support, as described in [PDTRACETCB].

*MTraceTC*: can be set to 1 to include the TC ID in trace data. Powers up as zero.

**Figure 8-12 Fields in the TCBCONFIG register**

31	30	25	24	21	20	17	16	14	13	11	10	9	8	6	5	4	3	0	
CF1	0	TRIG	SZ	CRMax	CRMin	PW	PiN	OnT	OfT	REV									

In TCBCONFIG:

*CF1*: read-only, reads zero because there are no more TCB configuration registers.

*PiN*: read-only, reads zero because the 34K core is a single-issue (single pipeline) processor.

*REV*: reads 1, denoting compliance with revision 4.xx of the TCB specification.

### 8.2.2 CP0 registers for the PDtrace™ logic

There are four:

- **TraceControl** and **TraceControl2**: allow the software to take charge of what is being traced.
- **UserTraceData**: allows software to send a “user format” trace record, which can be interpreted by suitable trace analysis software to build interesting facilities.
- **TraceBPC**: controls whether and how individual EJTAG breakpoint trace triggers take effect.

**Figure 8-13 Fields in the TraceControl and TraceControl2 registers**

31	30	29	28	27	26	25	24	23	22	21	20	13	12	5	4	3	2	1	0			
<i>TraceControl</i>	TS	UT	0	TPC	TB	IO	D	E	K	S	U	ASID_M		ASID		G	TFCR	TLSM	TIM	On		
31	30	29	28								21	20	19	12	11	7	6	5	4	3	2	0
<i>TraceControl2</i>	0	CPU- ldv	CPUid							TCV	TCNum		Mode	Valid- Modes	TBI	TBU	SyP					

*TS*: set 1 to put software (manipulating this register) in control of tracing. Zero from reset.

*UT*: software can output a "user triggered record" (just write any 32-bit value to the *UserTraceData* register). There have been two types of user-triggered record, and this bit says which to output: 0 → Type 1 record, 1 → Type 2.

*TPC*: turns on PC Sampling, where the current PC value is periodically sent to the trace memory (a different feature from the EJTAG “PC Sampling” feature described above [Section 8.1.12, "PC Sampling with EJTAG"](#).)

*TB*: "trace all branch" - when 1, output all branch addresses in full. Normally, predictable branches need not be sent.

*IO*: "inhibit overflow" - slow the CPU rather than lose trace data because you can't capture it fast enough.

*D, E, K, S, U*: do trace in various CPU modes: separate bits independently filter for debug, exception, kernel, supervisor and user mode. Set 1 to trace.

*ASID\_M, ASID, G*: controls ability to trace for just one (or some) processes, recognized by their current ASID value as found in *EntryHi[ASID]*. Set the *G* ("global") to trace instructions from all and any ASIDs. Otherwise set *TraceControl[ASID]* to the value you want to trace and *ASID\_M* to all 1s (you can also use *ASID\_M* as a bit mask to select several ASID values at once).

*TFCR*: switch on to generate full PC addresses for all function call and return instructions.

*TLSM*: switch on to trace all D-cache misses (potentially including the miss address).



*TIM*: switch on to trace all I-cache misses.

*On*: master trace on/off switch - set 0 to do no tracing at all.

The read-only fields in *TraceControl2* provide information about the capabilities of your PDtrace system. That system may include a plug-in probe, and in that case the *TraceControl2[SyP]* field may read as garbage until the probe is plugged in.

The first four fields are for tracing code running on MT CPUs:

*CPUIdV*, *CPUId*: when *CPUIdV* is set, trace data will only be generated by code run by the VPE identified in *CPUId*. Ignored if *TCV* is set.

*TCV*, *TCNum*: when *TCV* is set, trace only instructions run by the TC whose number is stored in *tTCNum*.

*Mode*: whenever trace is turned on, you capture an instruction trace. *Mode* is a bit mask which determines what load/store tracing will be done<sup>1</sup>. It's coded like this:

<i>Bit No Set</i>	<i>What gets traced</i>
0	PC
1	Load addresses
2	Store addresses
3	Load data
4	Store data

However, see *TraceControl2[ValidModes]* (description below) for what your PDtrace unit is actually capable of doing. Bad things can happen if you request a trace mode which isn't available.

*TraceControl2[ValidModes]*: what is this PDtrace unit capable of tracing?

<i>ValidModes</i>	<i>What can we trace?</i>
00	PC trace only
01	Can trace load/store addresses
10	Can trace load/store addresses and data

*TraceControl2[TBI, TBU]*: best considered together, these read-only bits tell you whether there is an on-chip trace memory, on-probe trace memory, or both - and which is currently in use.

<i>TBI</i>	<i>TBU</i>	<i>On-chip or probe trace memory?</i>
0	0	only on-chip memory available
0	1	only probe memory available
1	0	Both available, currently using on-chip
1	1	Both available, currently using probe

*TraceControl2[SyP]*: read-only field which lets you know how often the trace unit sends a complete PC address for synchronization purposes, counted in CPU pipeline clock cycles. The period is  $2^{(SyP + 5)}$

### 8.2.3 JTAG triggers and local control through TraceIBPC/TraceDBPC

Recent revisions of the PDtrace specification have defined much finer controls on tracing. In particular, you can now trace only cycles matching some “breakpoint” criteria, and there is a two-stage process where cycles are traced only after an “arm” condition is detected. The new fields are shown in [Figure 8-14](#)

- 
1. Prior to v4 of the PDtrace specification, this field was in *TraceControl*, and was too small to allow all conditions to be specified independently.

**Figure 8-14 Fields in the TraceIBPC/TraceDBPC registers**

	31	30	29	28	27	26	24	23	21	20	18	17	15	14	12	11	9	8	6	5	3	2	0
<i>TraceIBPC</i>	0	0		IE	ATE														IBPC3	IBPC2	IBPC1	IBPC0	
<i>TraceDBPC</i>				DE																	DBPC1	DBPC0	

In either *TraceIBPC* or *TraceDBPC*:

*IE,DE*: master 1-to-enable bit for triggers from EJTAG instruction and data breakpoints respectively.

*ATE*: Read-only bit which lets you know whether the additional trigger controls such as ARM, DISARM, and data-qualified tracing (introduced in v4.00 of the PDtrace specification) are available - which they may be on the 34K core.

*IBPC8-0, DBPC8-0*: each three-bit field encodes tracing options independently, for up to nine EJTAG I- and D-side breakpoints (this is generous: your 34K core will typically have no more than 4 I- and 2 D-breakpoints).

Each entry can be set as follows:

<i>xBPC field</i>	<i>Description</i>
0	Stop tracing (no effect if off already).
1	Start tracing (no effect if on already).
2	Trace instructions which cause this trigger.

However, do *TraceIBPC/TraceDBPC* exist in your system? They will be there only if you have an EJTAG unit (does *Config1[EP]* read 1?), and that unit has at least one breakpoint register - check that at least one of *DCR[DB,IB]* is set (as described in).

### 8.2.4 UserTraceData reg

Write any 32-bit value you like here and the trace unit will send a "user" record (there are two "types" of user record, and which you output depends on *TraceControl[UT]*, see above). You need to send something your trace analysis system will understand, of course! Perhaps it's worth noting that this "user" is local debug software, and doesn't mean low-privilege software running in "user mode" - which of course would not be able to access this register.

### 8.2.5 Summary of when trace happens

The many different enable bits which control trace add up to (or strictly "and" up to) a whole bunch of reasons why you won't get any trace output. So it may be worth summarizing them here. So:

- If software is in charge (that is, if *TraceControl[TS]==1*) then:
  - *TraceControl[On]* must be set.
  - At least one of the CPU mode filter bits *TraceControl[D,E,S,K,U]* must be set 1 to trace instructions in debug, exception, supervisor, kernel or user-mode respectively. Mostly likely either just *TraceControl[U]* will be set (to follow just one process in a protected OS), or *TraceControl[E,S,K,U]* to follow all the software at bare-iron level (but not to trace EJTAG debug activity);
  - Either *TraceControl[G]* is set (to trace everything regardless of current ASID) or *TraceControl[ASID]* (as masked by *TraceControl[ASID\_M]*) matches the current value of the core-under-test's *EntryHi[ASID]* field.
  - The signal *PDI\_TraceOn* is asserted by the trace block. This will typically be true whenever the probe is plugged in and connected to software.
  - As above there are *D,E,S,K,U,G* and *ASID* bits (there isn't an "ASID\_M" in this case) which must be set appropriately in the JTAG-accessible *TCBCTRL0LA* register, which is not otherwise described here.

Whether JTAG or *TraceControl* is in charge, then:

- There must have been a cycle recently when there was an "on trigger", that is:
  - The CPU tripped an EJTAG breakpoint with the *IBCn[TE]/DBCn[TE]* bit set to request a trace trigger (for I-side and D-side respectively);
  - *TraceIBPC[IE]/TraceDBPC[DE]* (respectively) was set to enable triggers from EJTAG breakpoints;
  - the appropriate *TraceBPC[IBPCx]/TraceBPC[DBPCx]* field has some kind of "on" trigger - and if this trigger is conditional on "arm" there must have been an arm event since system reset or any disarm event;
- And since the on-trigger time, there must not have been a cycle which acted as an "off trigger", that is:
  - The CPU tripped an EJTAG breakpoint with the *IBCn[TE]/DBCn[TE]* bit set, and *TraceBPC[IE]/TraceBPC[DE]* (respectively) were still set;
  - where the appropriate *TraceIBPC[IBPCn]/TraceDBPC[DBPCn]* fields is set to disable triggering (subject to arming).

If there is more than one breakpoint match in the same cycle, an "on" trigger wins out over any number of "off".

## 8.3 CP0 Watchpoints

Some cores may be built with no EJTAG debug unit to save space, and some debug software may not know how to use EJTAG resources. So it may be worth configuring some of the non-EJTAG CP0 watchpoint registers. If so they're described in [Appendix C, "Watchpoint registers" on page 151](#) below.

## 8.4 Performance counters

Performance counters are provided to allow software to monitor the occurrence of events of interest within the core, and can be very useful in analyzing system performance.

34K family CPUs are fitted with four counters, each of which can be set up to count one of a large choice of different events. Unlike almost all the other CP0 registers, the performance counters are *not* replicated per-VPE: the CPU has four counters, which either VPE may use. Each 32-bit counter is accompanied by a control register whose layout is shown in [Figure 8-15](#).

34K is a multi-threading CPU, and you can optionally count only events associated with a particular thread (by its TC number), or even those events associated with threads affiliated to a particular VPE. After some thought, I haven't documented in detail when you might get a different count if you narrow to a particular VPE or TC. In most cases it's obvious whether it makes sense to count a particular event for just one TC or VPE: where it's not obvious, experiment.

**Figure 8-15 Fields in the PerfCtl register**

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

There are usually four counters, but software should check using the *PerfCtl[M]* bit (which indicates "at least one more").

Then the fields are:

*TCID*: the TC number of the thread whose events should be counted, if just-one-TC counting is enabled (i.e. *MT\_EN*==10 binary.)

*MT\_EN*: available to restrict counting to events which are attributable to a particular VPE or TC:

<b>MT_EN</b> value	<i>What gets counted?</i>
00	Events from all TCs & VPEs (i.e., don't filter)
01	Count events from all TCs affiliated to the VPE specified in the <i>VPEID</i> field. Some events can't be tied to a particular VPE - use common sense.
10	Count events only for the TC specified by the <i>TCID</i> field. Again, some events are not TC-specific.
11	Reserved

*VPEID*: defines the VPE all of whose TC's events should be counted, if just-this-VPE counting is enabled (i.e. *MT\_EN*==01 binary.)

*Event*: determines which event this counter will count; see [Table 8.3](#) below. Note that the odd-numbered and even-numbered counters mostly count different events, though some particularly important events can use any of the four counters.

*IE*: set to cause an interrupt when the counter "overflows" into its bit 31. This can either be used to implement an extended count, or (by presetting the counter appropriately) to notify software after a certain number of events have happened. The interrupt is implemented by taking a set of signals (usually *SI\_PCI* - one per VPE) out of the core, which the system integrator will have sent back in, each as one of the core's interrupt inputs. The output signal activated will depend on the VPE affiliation of the thread which last wrote to the control register, which will normally be what you want.

*U, S, K, EXL*: count events in User mode, Supervisor mode, Kernel mode and Exception mode (i.e. when *Status[EXL]* is set) respectively. Set multiple bits to count in all cases.

The events which can be counted in the 34K core are in [Table 8.3](#). Blank fields are reserved. But before you get there, take a look at the next sub-section...

### 8.4.1 Reading the event table.

There are a lot of events you can count. It's relatively cheap to wire another signal from the internals of the core into a counter. It's time consuming and expensive to formulate a signal which represents exactly what a software engineer might want to count, and even more expensive to test it. Where the definitions in [Table 8.3](#) are clear and simple, they're usually exactly right. Where they seem more obscure, tread carefully, and don't just blame the author of this manual (though sometimes it is my fault!) When you use a counter, use it first on a piece of code where you know the answer, and check you're really counting what you think you are.

When reading the table:

- *T, V, P*: in the "Type" column, mark an event which can be filtered per-TC, per-VPE or is just global (respectively). Per-TC events can be counted per-VPE, and per-VPE events can be counted globally. When you count per-TC events per-VPE or globally the counter will advance in any cycle where the event happens for any TC under consideration. Counters never advance faster than once per clock.
- *IFU*: is the "instruction fetch unit" of the CPU pipeline. We can't describe some events without referring to the inside of the CPU. You might like to look back at [Section 3.1 "The 34K™ core pipeline and multithreading"](#).
- *Replay*: when an instruction will block for a long period, sequentially-later instructions from the same TC which have got into the main pipeline must be discarded. These instructions will usually have been retained in the "skid buffer" of the IFU, so the IFU queues can be adjusted so that when the instruction unblocks, the TC can continue correctly from the following instruction. This sequence is called a "replay" and these events count the pipeline bubbles which result.
- *Refetch*: if you'd like to do a replay but the relevant instructions are not available in the skid buffer, the IFU must be instructed to discard all stored instructions for the TC and fetch them again. This event counts the number of pipeline bubbles which result.
- *Stall*: in general, "stall" counters count the cycles when the whole pipeline is blocked and no TC can make forward progress. If this type of counter is set for a particular TC, it will only count if this TC is causing the stall.

But subunits causing a stall can also signal a "long stall", and the main pipeline takes that as a cue to deschedule the blocked TC until the condition is resolved. The counters documented as "stall" or "stalled" do not count time while one TC is blocked but others continue to run.

- *Blocked cycle*: "events" like this count all and any cycles when a TC is blocked by something.
- *LDQ, FSB, WBB*: CPU queues, described in [Section 6.3.1, "Read/write ordering and queues in the 34K core"](#).
- *Instruction fetch events*: these include I-cache, ITLB and JTLB events. They are not as directly related to the instructions in your program as you might think:
  - 24K/34K CPUs have a 64-bit wide interface to the I-cache and fetch two instructions at once.
  - After a cache miss is resolved, the IFU re-fetches the missed data; the counters will count this twice.
  - The IFU always reads instructions ahead, and on a branch or exception some of the instructions fetched will never be executed. Moreover, the IFU's branch predictors sometimes cause it to fetch speculatively from a predicted branch target which turns out to be wrong; those speculative instructions will never be executed either.
  - If there's an exception-causing address error during I-fetch, it won't be counted.

- *Single-threaded mode (ST mode)*: when only one TC is eligible for scheduling, 34K enters ST mode. In ST mode some blocking events which would have been dealt with by suspending the thread and possibly replaying an instruction are handled by a whole-pipeline stall instead, which has less overhead.

Table 8.3 Performance counter events

Event No	Counter 0 and 2	Type	Counter 1 and /3	Type
0	Cycles			P
1	Instructions completed			T
2	Branch instructions completed.	T	Branch mispredictions	T
3	<b>jr</b> \$31 (return) instructions	T	<b>jr</b> \$31 predicted but guessed wrong	T
4	<b>jr</b> (not \$31) instructions	T	<b>jr</b> \$31 not predicted (the return predictor only works for one TC at a time).	T
5	ITLB accesses. There will be one for every I-fetch in a translated address region.	T	ITLB misses. Note that if two TCs cause “the same” ITLB miss in quick succession, that will only be counted once.	T
6	DTLB accesses.	T	DTLB misses	T
7	JTLB instruction accesses (same as ITLB misses).	T	JTLB I-misses: this counts TLB misses <i>and</i> TLB invalid conditions on I-fetch.	T
8	JTLB data accesses (same as DTLB misses)	T	JTLB D-miss: counts TLB misses + TLB invalid on D-access.	T
9	Instruction cache accesses. That's <i>every</i> access including replays (and as above, including instructions which are never executed). But more: for example, following a branch which is correctly predicted taken, one or more instructions on the straight-through path may be accessed.	T	Instruction cache misses. Includes misses resulting from fetch-ahead and speculation.	T
10	Data cache load/stores	T	D-cache writebacks (actually counts number of D-misses or cacheops which trigger writeback.)	T
11	Loads/stores which miss in D-cache			T
12-13	reserved			
14	Integer instructions completed	T	FPU instructions completed (not including loads and stores)	T
15	Loads completed (including FP loads)	T	Stores completed (includes FP stores)	T
16	<b>j/jal</b> instructions completed	T	MIPS16 instructions completed	T
17	no-ops completed. Early revision cores count only strict <b>nop</b> instructions, but later ones count any 3-operand instruction which discards its output by writing register \$0.	T	Integer multiply/divide unit instructions completed	T
18	Cycles where the main pipeline (RF stage) does not advance. This is either because there is no instruction scheduled, or because the ALU is backed up and can't accept an instruction	P	Replays: that is, events where IFU is made to re-issue instructions which were already scheduled once.	T
19	<b>sc</b> instructions completed	T	<b>sc</b> instructions failed	T
20	Prefetch instructions to cached addresses	T	Prefetch instructions completed with cache hit	T
21	L2 cache writebacks	P	L2 cache accesses	P

Event No	Counter 0 and 2	Type	Counter 1 and /3	Type
22	L2 cache misses			P
23	Exceptions taken	T	Cycles spent in "Single Threaded Mode".	T
24	Cycles when main pipeline is stalled while the LSU has to do a "replay". A good example of a replay is when the fill buffer gets full and needs to be emptied out to make forward progress. To empty out the buffer, the LSU has to take control of the cache which is currently being accessed by other in-flight LSU instructions. To accomplish this, the pipeline is stalled, the FSB accesses the cache to empty out its data, and then the instructions that were in flight are replayed to get their data from the cache.	T	"Refetches": Counts all replayed instructions (instructions which are send back to IFU to be refetched (and reissued)). If an instruction has been replayed multiple times, you get a count for each event.	T
25	Cycles when no instructions are available to issue for any TC	P	Cycles when main pipeline stops because an ALU operation is taking more than one clock	P
26	DSP Instructions Completed	T	ALU-DSP Saturations Done	T
27			MDU-DSP Saturations Done	T
28-31	<i>Available to count implementation-specific events signalled by wires from configurable interfaces.</i>			
28	Available for customer PM event	T	Available for customer CP2 event	T
29	Available for customer ISPRAM event	T	Available for customer DSPRAM event	
30	Available for CorExtend event	T		
31	Available for external yield manager event.	T	Custom ITC event	T
32	ITC Loads. If a TC is halted or takes an exception, a pending ITC operation will be aborted, then later retried. Each retry is counted.	T	ITC Stores issued. Invisible retries counted too, as for loads.	T
33	Uncached Loads	T	Uncached Stores	T
34	<b>fork</b> Instructions completed	T	<b>yield</b> instructions completed	T
35	CP2 register-to-register instructions completed	T	<b>mfc2/mtc2</b> instructions completed	T
36	reserved			
37-46	Count number of cycles (most often "stall cycles", i.e. time lost), not just number of events. See note on stall cycles above.			
37	I-cache miss blocked cycles - counts cycles when the TC has no instruction to issue following an I-fetch miss. This ignores the stalls due to ITLB misses as well as the 4 cycles following a redirect.	T	D-cache miss blocked cycles - counts cycles when TC is blocked when an instruction uses a register value which is subject to a load miss.	T
38	L2 I-miss stall cycles	P	L2 data miss stall cycles	P
39	D-miss cycles	P	L2 miss cycles	P
40	Uncached access block cycles	T	ITC stall cycles: when no instruction for any TC can be issued, and a TC selected for counting is waiting for an ITC operation	T
41	MDU stall cycles - note that it's possible for the MDU to indicate a "long stall" where the TC waiting for the MDU gets suspended - that wait will <i>not</i> be counted here.	T	FPU stall cycles	T
42	CP2 stall cycles	T	CorExtend stall cycles	T

34K™ core features for debug and profiling

Event No	Counter 0 and 2	Type	Counter 1 and /3	Type
43	ISPRAM stall cycles - when no instruction can be issued because the IFU has run out of instructions, after the ISPRAM sent a "not ready" indication (which requires a retry). Doesn't include a count for the 4 cycles after a redirect.	T	DSPRAM stall cycles	T
44	CACHE instruction stall cycles	P		
45	Load to Use stalls	T	Stalls when a load/store base register was computed by the preceding instruction.	T
46	Read-CP0-value interlock stalls.	T	Branch mispredict: lost cycles resulting from a mispredict (24K only).	Pi
47	Relax bubbles	V		
48	IFU FB full refetches: count up when the IFU has to refetch an address because the FB was full on a miss.	T	FB entry allocated	P
49	EJTAG Instruction triggers	T	EJTAG data triggers	T
50-55	Monitor the state of various FIFO queues relating to loads and stores, as described in <a href="#">Section 6.3.1, "Read/write ordering and queues in the 34K core"</a> .			
50	FSB < 1/4 full	P	FSB 1/4-1/2 full	P
51	FSB > 1/2 full	P	FSB full pipeline stalls	P
52	LDQ < 1/4 full	P	LDQ 1/4-1/2 full	P
53	LDQ > 1/2 full	P	LDQ full pipeline stalls	P
54	WBB < 1/4 full	P	WBB 1/4-1/2 full	P
55	WBB > 1/2 full	P	Cycles when whole CPU is stopped because an instruction needs to write data out of the core, but all write buffer entries are full.	P



## Programming the 34K™ core in user mode

Sections include:

- [Section 9.1, "The multiplier"](#): multiply, multiply/accumulate and divide timings.
- [Section 9.2, "User-mode accessible "Hardware registers""](#)
- [Section 9.3, "Prefetching data"](#): how it works.
- [Section 9.4, "Using "synci" when writing instructions"](#): writing instructions without needing to use privileged cache management instructions.
- [Section 9.5, "Tuning software for the 34K family pipeline"](#): for determined programmers, and for compiler writers. It includes information about the timing of the DSP ASE instructions.
- [Section 9.6 "Floating point instruction timing and data dependencies"](#): the floating-point unit often runs at half speed, and some of its interactions (particularly about potential exceptions) are complicated. This section offers some guidance about the timing issues you'll encounter.

### 9.1 The multiplier

As is traditional with MIPS CPUs, the integer multiplier is a semi-detached unit with its own pipeline. All MIPS32 CPUs implement:

- **mult/multu**: a 32×32 multiply of two GPRs (signed and unsigned versions) with a 64-bit result delivered in the multiply unit's pseudo-registers *hi* and *lo* (readable only using the special instructions **mfhi** and **mflo**, which are interlocked and stall until the result is available).
- **madd, maddu, msub, msubu**: multiply/accumulate instructions collecting their result in *hi/lo*.
- **mul/mulu**: simple 3-operand multiply as a single instruction.
- **div/divu**: divide - the quotient goes into *lo* and the remainder into *hi*.

Many of the most powerful instructions in the MIPS DSP ASE are variants of multiply or multiply-accumulate operations, and are described in [Section , "The MIPS32® DSP ASE"](#). The DSP ASE also provides three additional "accumulators" which behave like the *hi/lo* pair).

No multiply/divide operation ever produces an exception - even divide-by-zero is silent - so compilers typically insert explicit check code where it's required.

The 34K core multiplier is high performance and pipelined; multiply/accumulate instructions can run at a rate of 1 per clock, but a 32×32 3-operand multiply takes four clocks longer than a simple ALU operation. Divides use a bit-per-clock algorithm, which is short-cut for smaller dividends. The result is that many of these operations will not be

finished in time for the next instruction to proceed without delay. For details see [Section 9.5, "Tuning software for the 34K family pipeline"](#), and in particular [Section 9.5.4, "Data dependency delays classified"](#).

## 9.2 User-mode accessible "Hardware registers"

The 34K core complies with Revision 2 of the MIPS32 specification, which introduces *hardware registers*; CPU-dependent registers which are readable by unprivileged user space programs, usually to share information which is worth making accessible to programs without the overhead of a system call.

The hardware registers provide useful information about the hardware, even to unprivileged (user-mode) software, and are readable with the **rdhwr** instruction. [MIPS32] defines four registers so far. The OS can control access to each register individually, through a bitmask in the CP0 register *HWREna* - (set bit 0 to enable register 0 etc). *HWREna* is cleared to all-zeroes on reset, so software has to explicitly enable user access. Privileged code can access any hardware register.

The four standard registers are:

- *CPUNum* (0): Number of the CPU on which the program is currently running. This comes directly from the coprocessor 0 *EBase[CPUNum]* field.
- *SYNCL\_Step* (1): the effective size of an L1 cache line<sup>1</sup>; this is now important to user programs because they can now do things to the caches using the **synci** instruction to make instructions you've written visible for execution. Then *SYNCL\_Step* tells you the "step size" - the address increment between successive **synci**'s required to cover all the instructions in a range.

If *SYNCL\_Step* returns zero, that means that you don't need to use **synci** at all.

- *CC* (2): user-mode read-only access to the CP0 *Count* register, for high-resolution counting. Which wouldn't be much good without.
- *CCRes* (3): which tells you how fast *Count* counts. It's a divider from the pipeline clock (if you read a value of "2", then *Count* increments every 2 cycles, at half the pipeline clock rate - which is what you'll find for any 34K family core defined so far.

## 9.3 Prefetching data

MIPS32 CPUs are being increasingly used for computations where you'd once have needed a DSP. These computations often feature loops accessing large arrays, and the run-time is often dominated by cache misses.

These are excellent candidates for using the **pref** instruction, which gets data into the cache without affecting the CPU's other state. In a well-optimized loop with prefetch, data for the next iteration can be fetched into the cache in parallel with computation for the last iteration.

It's a pretty major principle that **pref** should have *no software-visible effect* other than to make things go faster.

**pref** is logically a no-op<sup>2</sup>.

The **pref** instruction comes with various possible "hints" which allow the program to express its best guess about the likely fate of the cache line. In 34K family cores the "load" and "store" variants of the hints do the same thing; but

---

1. Strictly, it's the lesser of the I-cache and D-cache line size, but it's most unusual to make them different.  
2. This isn't quite true any more; **pref** with the "PrepareForStore" hint can zero out some data which wasn't previously zero.

it makes good sense to use the hint which matches your program’s intention - you might one day port it to a CPU where it makes a difference, and it can’t do any harm.

The 34K core acts on hints as summarized in [Table 9.1](#).

## 9.4 Using “synci” when writing instructions

**Table 9.1 Hints for “pref” instructions**

<i>No</i>	<i>Hint Name</i>	<i>What happens in the 34K core</i>	<i>Why would you use it?</i>
0 1	load store	Read the cache line into the D-cache if not present.	When you expect to read the data soon. Use "store" hint if you also expect to modify it.
4 5	load_streamed store_streamed	Fetch data, but always use cache way zero - so a large sequence of "streamed" prefetches will only ever use a quarter of the cache.	For data you expect to process sequentially, and can afford to discard from the cache once processed
6 7	load_retained store_retained	Fetch data, but never use cache way zero. That means if you do a mixture of "streamed" and "retained" operations, they will not displace each other from the cache.	For data you expect to use more than once, and which may be subject to competition from "streamed" data.
25	writeback_invalidate/nudge	If the line is in the cache, invalidate it (writing it back first if it was dirty). Otherwise do nothing. However (with the 34K core only): if this line is in a region marked for “uncached accelerated write” behavior, then write-back this line.	When you know you’ve finished with the data, and want to make sure it loses in any future competition for cache resources.
30	PrepareForStore	If the line is not in the cache, create a cache line - but instead of reading it from memory, fill it with zeroes and mark it as "dirty". If the line is already in the cache do nothing - <i>this operation cannot be relied upon to zero the line.</i>	When you know you will overwrite the whole line, so reading the old data from memory is unnecessary. A recycled line is zero-filled only because its former contents could have belonged to a sensitive application - allowing them to be visible to the new owner would be a security breach.

The **synci** instruction (introduced with Revision 2 of the MIPS32 architecture specification, [\[MIPS32\]](#)) ensures that instructions written by a program (necessarily through the D-cache, if you’re running cached) get written back from the D-cache and corresponding I-cache locations invalidated, so that any future execution at the address will reliably execute the new instructions. **synci** takes an address argument, and it takes effect on a whole enclosing cache-line sized piece of memory. User-level programs can discover the cache line size because it’s available in a “hardware registers” accessed by **rdhwr**, as described in [Section 9.2](#), “User-mode accessible “Hardware registers”” above.

## 9.5 Tuning software for the 34K family pipeline

This section is addressed to low-level programmers who are tuning software by hand and to those working on efficient compilers or code translators.

Note, though, that when there is a multi-threading workload some of the following issues become less important. There’s not so much need to mitigate cache miss delays (for example) when the time when one thread is waiting will be cheerfully used by another thread which keeps running.

The 34K core is a pipelined design, and the pipeline and some of its consequences are described in [Section , "How the 34K™ core implements multi-threading"](#). That leads to a class of possible delays to do with data dependencies. For software tuning purposes it's usually enough to know the delay which results when one instruction (the "producer") generates a value in some particular register for the use of the next instruction in sequence (the "consumer"). The delay is in processor cycle time units, but it makes good sense to think of that delay as a lost opportunity to run an instruction. To tune round data dependencies, the programmer or compiler needs to re-order the instructions so that enough useful but independent instructions are placed between the producer and consumer that the consumer runs without delay.

There are times when interactions are more complicated than that. While you can pore over hardware books to try to figure out what the pipeline is doing, when it gets that difficult we advise that you should obtain a cycle-accurate simulator or other well-instrumented test environment, and try your software out.

But before getting on to data delays, we'll look at the most important causes of slow-down: cycles lost to cache misses and branches.

### 9.5.1 Cache delays and mitigating their effect

In a typical 34K CPU implementation a cache miss which has to be refilled from DRAM memory (in the very next chip on the board) will be delayed by a period of time long enough to run 50-100 instructions. A miss or uncached read (perhaps of a device register) may easily be several times slower. These really are important!

Of course, this is one of the main motivations for having a multithreading CPU: while one thread is stopped because of a cache miss, other threads can keep running, greatly improving the total throughput.

Because these delays are so large, there's not a lot you can do to help a cache-missing thread make progress. But every little helps. The 34K core has non-blocking loads, so if you can move your load instruction producer away from its consumer, you won't start paying for your memory delay until you try to run the consuming instruction.

Compilers and programmers find it difficult to move fragments of algorithm backwards like this, so the architecture also provides prefetch instructions (which fetch designated data into the D-cache, but do nothing else). Because they're free of most side-effects it's easier to issue prefetches early. Any loop which walks predictably through a large array is a candidate for prefetch.

The **pref PrepareForStore** prefetch saves a cache refill read, for cache lines which you intend to overwrite in their entirety. Read more about prefetch in [Section 9.3, "Prefetching data"](#) above.

#### Tuning data-intensive common functions

Bulk operations like `bcopy()` and `bzero()` will benefit from some CPU-specific tuning. To get excellent performance for in-cache data, it's only necessary to reorganize the software enough to cover the address-to-store and load-to-use delays. But to get the loop to achieve the best performance when cache missing, you probably want to use some prefetches.

### 9.5.2 Branch delay slot

It's a feature of the MIPS architecture that it always attempts to execute the instruction immediately following a branch. The rationale for this is that it's extremely difficult to fetch the branch target quickly enough to avoid a delay, so the extra instruction runs "for free"...

Most of the time, the compiler deals well with this single delay slot. MIPS low-level programmers find it odd at first, but you get used to it!

### 9.5.3 Branch misprediction delays

In a long-pipeline design like this, branches would be expensive if you waited until the branch was executed before fetching any more instructions. See Section 3.1, "The 34K™ core pipeline and multithreading" for what is done about this: but the upshot is that where the fetch logic can't compute the target address, or guesses wrong, that's going to cost five or more lost cycles - but most of the pain is felt by the thread which executes the branch; so long as there are other running threads the CPU can keep busy. It does depend what sort of branch: the conditional branch which closes a tight loop will almost always be predicted correctly after the first time around.

However, too many branches in too short a period of time can overwhelm the ability of the instruction fetch logic to keep ahead with its predictions. Where branchy code can be replaced by conditional moves, you'll get significant benefits.

The branch-likely<sup>1</sup> instructions (officially deprecated by the MIPS32 architecture because they may perform poorly on more sophisticated or wider-issue hardware) are predicted just like any other branch.

Although deprecated, the branch-likely instructions will probably improve the performance of loops where there is no other way of avoiding a no-op in a loop-closing branch's delay slot. If you're tempted to use this, we strongly recommend you make the code conditional on a `#define` variable tied specifically to the 34K family. If that's difficult in your environment and the code might need to be portable, it's probably better not to use this.

### 9.5.4 Data dependency delays classified

We've attempted to tabulate all possible producer/consumer delays affecting user-level code (we're not discussing CPU registers here), but excluding floating point (which is in the next section).

In fact, we won't set out the tables exactly like that. The MIPS instruction set is efficient because, most of the time, dependent instructions can be run nose-to-tail without delay. For all registers, there is a "standard" place in the pipeline where the producer should deliver its value and another place in the pipeline where the consumer picks it up<sup>2</sup>. Producer/consumer delays happen when either the producer is late delivering a result to the register (we'll abbreviate to "lazy"), or the consumer insists on obtaining its operand early (we'll abbreviate to "eager"). Of course, both may happen: in that case the delays add up.

It's important to be clear what class of registers is involved in any of these delays. For non-floating-point user-level code, there are just three classes of registers to consider:

- General purpose registers ("GPR");
- The *hi/lo* pair together with the three additional accumulators defined by the MIP DSP ASE ("ACC");
- The fields of the *DSPControl* register.

So that gives us two tables.

#### Delays caused by "eager consumers" reading values early

- 
1. The "likely" in the instruction name is historical, and pretty misleading.
  2. These are brought closer together by the magic of register file bypasses, but we don't need to get into the details here.

**Table 9.2 Register → eager consumer delays**

<i>Reg → Eager consumer</i>	<i>Del</i>	<i>Applies when...</i>
GPR → load/store	1	the GPR value is an address operand (store data is not needed early).
ACC → multiply instructions	1	the ACC value came from any <i>non-multiply</i> or <i>multiply instructions which saturate the accumulator value</i> (values generated by other multiply instructions are made available early, and thus avoid this delay).
ACC → DSP instructions which extract selected bits from an accumulator: <b>extp...</b> , <b>extr...</b> etc. DSP instructions which write a shifted value back to the accumulator: <b>mtflip</b> , <b>shilo</b> , <b>shilov</b> .	3	Always

**Delays caused by "lazy producers" delivering values late**

**Table 9.3 Lazy producer → register delays**

<i>Lazy producer → Reg</i>	<i>Del</i>	<i>Applies when...</i>
Load → GPR	1	Always (familiar as the "load delay slot").
Integer multiply unit instructions producing a GPR result. Instructions reading accumulators and writing GPR (e.g. <b>mflo</b> ).	4	Always (because the multiply unit pipeline is longer than the integer unit's).
DSP "ALU" instructions (which neither read nor write an accumulator, nor do a multiplication).	1	Always
Integer divide instruction → ACC	7 9 15 17 23 25 31 33	8-bit dividend 8-bit dividend & negative operand to <b>div</b> 16-bit dividend 16-bit dividend & negative operand to <b>div</b> 24-bit dividend 24-bit dividend & negative operand to <b>div</b> full-size dividend full-size dividend & negative operand to <b>div</b>

**How to use the tables**

Suppose we've got an instruction sequence like this one:

```

addiu    $a0, $a0, 8
lw       $t0, 0($a0) # [1]
lw       $t1, 4($a0)
addu     $t2, $t0, $t1# [2]
mul      $v0, $t2, $t3
sw       $v0, 0($a1) # [3]
    
```

Then a look at the tables should help us discover whether any instructions will be held up. Look at the dependencies where an instruction is dependent on its predecessor:

- [1] The **lw** will be held up by one clock, because its GPR address operand *\$a0* was computed by the immediately preceding instruction (see the first box of [Table 9.2](#).) The second **lw** will be OK.
- [2] The **addu** will be one clock late, because the load data from the preceding **lw** arrives late in the GPR *\$t1* (see the first box of [Table 9.3](#).)
- [3] The **sw** will be 4 clocks late starting while it waits for a result from the multiply pipe (the second box of [Table 9.3](#).)

These can be additive. In the pointer-chasing sequence:

```
lw      $t1, 0($t0)
lw      $t2, 0($t1)
```

The second load will be held up two clocks: one because of the late delivery of load data in *\$t1* (first box of [Table 9.3](#)), plus another because that data is required to form the address (first box of [Table 9.2](#).)

### Delays caused by dependencies on DSPControl fields

Some DSP ASE instructions are dependent because they produce and consume values kept in fields of the *DSPControl* register. However, the most performance-critical of these dependencies are "by-passed" to make sure no delay will occur - those are the dependencies between:

```
addsc → DSPControl[c] → addwc
cmp.x → DSPControl[ccond] → pick.x
wrdsp → DSPControl[pos,scount] → insv
```

But other dependencies passed in *DSPControl* may cause delays; in particular the *DSPControl[outlag]* bits set by various kinds of overflow are not ready for a succeeding **rddsp** instruction. The access is interlocked, and will lead to a delay of up to three clocks. We don't expect that to be a problem (but if you know different, please get in touch with MIPS Technologies).

### More complicated dependencies

There can be delays which are dependent on the dynamic allocation of resources inside the CPU. In general you can't really figure out how much these matter by doing a static code analysis, and we earnestly advise you to get some kind of high-visibility cycle-accurate simulator or trace equipment (probably based on [Section 8.2](#), "PDtrace™ instruction trace facility").

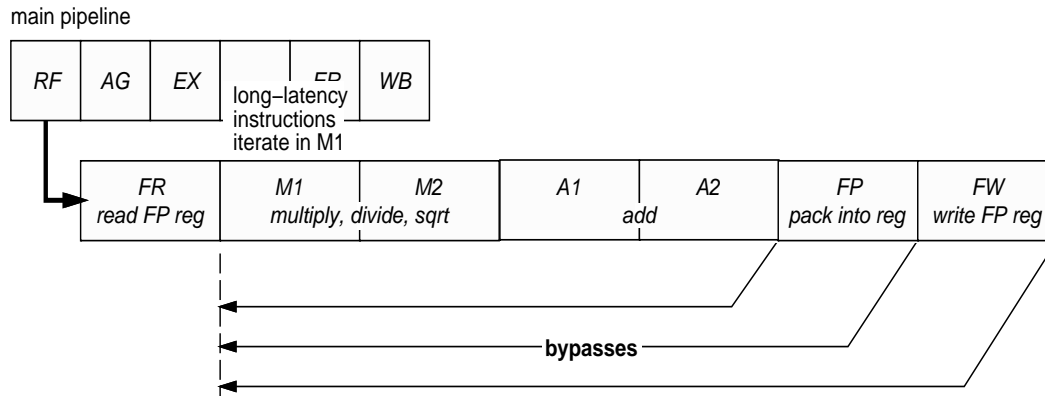
### Advice on tuning DSP ASE instruction sequences

DSP algorithm functions are often the subject of intense tuning. There is more specific and helpful advice (with examples) included in the white paper [\[DSPWP\]](#) published by MIPS Technologies.

## 9.6 Floating point instruction timing and data dependencies

This is not so simple. The floating point unit (FPU) has its own pipeline. More often than not, the FPU uses a half-rate clock compared to the integer core - a full-speed FPU is a build option, but it will then usually limit the clock rate which can be used in your design. The FPU pipeline is shown in [Figure 9-1](#).

**Figure 9-1 Overview of the FPU pipeline**



The FPU is a multiply-add pipeline, and all register-to-register instructions go through six stages:

*FR*: obtains FP register values and converts them into an expanded internal format; With a half-speed FPU, instructions issued from the integer core on an “odd” CPU cycle must wait one CPU clock time to start in the FPU.

*M1*, *M2*: multiply operation as required. Some long-latency operations “loop” in the *M1* stage until complete, holding up any subsequent FP instruction which would otherwise enter *M1*. Earlier instructions continue to run, leaving bubbles in the FP pipeline stages *M2* through *FW*.

*A1*, *A2*: add operation as required.

*FP*: convert result back to standard stored form and round.

*FW*: write back to FP register.

### 9.6.1 FPU register dependency delays

Any FPU instruction must go through pipeline stages from *M1* through *A2* before it produces a result, which can then (as shown by the “bypass” lines in the pipeline diagram) be used by a dependent instruction reaching the *M1* stage. If you want to keep the FPU pipeline full, that means there must be three non-dependent instructions between the consumer and producer of an FP value. However, other FP instruction delays can create bubbles in the FP pipeline, and then you’ll need less than three intervening instructions.

### 9.6.2 Delays caused by “long-latency” instructions looping in the *M1* stage

Instructions which take only one clock in *M1* go through the pipeline smoothly and can be completed one per FPU clock period. Instructions which take longer in *M1* always prevent the next instruction from starting in the next clock, regardless of any data dependency. Those long-latency instructions - double-precision multiplies and all division and square root operations - are listed in [Table 9.4](#). An instruction which runs for 2 cycles in *M1* holds up the FPU pipeline for one clock and so on - and of course the cycle counts are for FPU cycles.



Table 9.4 Long-latency FP instructions

Operand	Instruction type	Instructions	Cycles in M1
Double-precision (64-bit)	Any multiplication	<b>mul.d, madd.d, msub.d, nmadd.d, nmsub.d</b>	2
Single-precision (32-bit)	Reciprocal	<b>recip.s</b>	10
	divide, square-root	<b>div.s, sqrt.s</b>	14
	reciprocal square root	<b>rsqrt.s</b>	14
Double-precision (64-bit)	Reciprocal	<b>recip.d</b>	21
	divide, square-root	<b>div.d, sqrt.d</b>	29
	reciprocal square root	<b>rsqrt.d</b>	31

### 9.6.3 Delays on FPU load and store instructions

FPU store instructions stall in the main pipeline *EX* stage until the register data arrives from the FPU. Provided that the store instruction doesn't get behind slow FP instructions, FP stores run no more than one every two instructions should not produce further delays.

FPU load instructions are subject to the usual FPU timing. So long as the load hits in the cache, you should see no more than the usual FP producer-consumer delay from load to use.

### 9.6.4 Delays when main pipeline waits for FPU to decide not to take an exception

The MIPS architecture requires FP exceptions to be "precise", which (in particular) means that no instruction after the FP instruction causing the exception may do anything software-visible. That means that an FP instruction in the main pipeline may not proceed past the *ER* stage until the FPU can either report the exception, or confirm that the instruction will not cause an exception.

Floating point instructions cause exceptions not only because a user program has requested the system to trap IEEE exceptional conditions (which is unusual) but also because the hardware is not capable of generating or accepting very small ("denormalized") numbers in accordance with the IEEE standards. The latter ("unimplemented") exception is used to call up a software emulator to patch up some rare cases. But the main pipeline must be stalled until the FP hardware can rule out an exception, and that leads to a stall on every non-trivial FP operation. With a half-rate FPU, this stall will most likely be 6-7 clocks.

Software which can tolerate some deviation from IEEE precision can avoid these delays by opting to replace all denormalized inputs and results by zero - controlled by the *FCSR[FS,FO,FN]* register bits described in [Figure D.2](#) and its notes. If you have also disabled all IEEE traps, you get no possibility of FP exceptions and no extra main pipeline stalls.

### 9.6.5 Delays when main pipeline waits for FPU to accept an instruction

The half-speed FPU can never accept more than one instruction for every two main pipeline clocks. But if some of your FP instructions are the long-latency ones described above, the FP pipeline has room for just one more instruction before it backs up. Once it does back up, your whole CPU will stall until the long-latency instruction completes.

## 9.6.6 Delays on `mfc1/mtc1` instructions

Any FP instruction with GP register operands gets sent the GP values when it is launched, so `mtc1` instructions have standard FP instruction timing.

An `mfc1` instructions needs to write data into the GP register file. In general it will not complete quickly enough to use its main-pipeline *WB* slot, so the value returning to the integer unit must wait until the integer unit is not using the GP register write port. The instruction which uses the value obtained by the `mfc1` may stall until the data is available, but that usually won't be very long.

## 9.6.7 Delays caused by dependency on FPU status register fields

The conditional branch instructions `bc1f/bc1t` and the conditional moves `movf/movt` execute in the main pipeline, but test a FP condition bit generated by the various FPU compare instructions.

## 9.6.8 Slower operation in “MIPS I™” compatibility mode

Historic 32-bit MIPS CPUs had only 16 “even-numbered” floating point registers usable for arithmetic, with odd-numbered registers working together with them to let you load, store and transfer double-precision (64-bit) values. Software written for those old CPUs is incompatible with the full modern FPU, so there's a compatibility bit provided in *Status[FR]* - set zero to use MIPS I compatible code. This comes at the cost of slower repeat rates for FP instructions, because in compatibility mode not all the bypasses shown in the pipeline diagram above are active.

## References

### 34K core family manuals

- [**34KC\_DATA**]:“MIPS32 34Kc Datasheet”, MIPS Technologies document MD00418.
- [**34KF\_DATA**]:“MIPS32 34Kf Datasheet”, MIPS Technologies document MD00418: these are the basic hardware introductions, for family members without and with floating point unit.
- [**34K\_INT**]:““MIPS32® 34K™ Processor Core Family Integrator's Guide”, MIPS Technologies document MD00415, available to core licensees, describes the options available with the core.
- [**34K\_ERR**]:“MIPS32 34K Processor Core Family Errata Sheet”, MIPS Technologies document MD00417, lists any known deviations from spec. Restricted access.

### Other Manuals from MIPS Technologies

First of all, basic architectural documents:

[MIPS32]:the MIPS32 architecture definitions, in three volumes:

- [**MIPS32V1**]: “Introduction to the MIPS32 Architecture”, MIPS Technologies document MD00080.
- [**MIPS32V2**]: “The MIPS32 Instruction Set”, MIPS Technologies document MD00084.
- [**MIPS32V3**]: “The MIPS32 Privileged Resource Architecture”, MIPS Technologies document MD00088.

[MIPS64]:the MIPS64 architecture definition. Although the 34K core is a MIPS32 CPU, its floating point unit is described in these documents:

- [**MIPS64V1**]: “Introduction to the MIPS64 Architecture”, MIPS Technologies document MD00081.
- [**MIPS64V2**]: “The MIPS64 Architecture Instruction set”, MIPS Technologies document MD00085.
- [**MIPS64V3**]: “The MIPS64 Privileged Resource Architecture”, MIPS Technologies document MD00089.

Then there are some architectural extensions:

- [MIPSMT]:“The MIPS MT Application-Specific Extension to the MIPS32 Architecture”, MIPS Technologies document MD00376.
- [MIPSDSP]:“The MIPS DSP Application-Specific Extension to the MIPS32 Architecture”, MIPS Technologies document MD00372.
- [**DSPWP**]: “Effective Programming of the 24KE and 34K Cores for DSP Code”, MIPS Technologies white paper, document number MD00475.
- [MIPS16e]:“The MIPS16e™ Application-Specific Extension to the MIPS32 Architecture”, MIPS Technologies document MD00074.
- [CorExtend]:“How To Use CorExtend™ User-Defined Instructions”, MIPS Technologies document MD00333.
- [**CorExtend-24K**]:“CorExtend Instructions Integrator's Guide for MIPS32 24K, 24KE and 34K Pro Series™ cores”, MIPS Technologies document MD00348.

## References

[MDMX]: “The MDMX™ Application-Specific Extension to the MIPS64 Architecture”, MIPS Technologies document MD00095.

[EJTAG]: “EJTAG Specification”, MIPS Technologies document MD00047.

[PDTRACEIF]: “PDtrace™ Interface Specification”, MIPS Technologies document MD00136.

[EJTAGTRACE]: “EJTAG Trace Control Block Specification”, MIPS Technologies document MD00148.

[PDTRACEUSAGE]: “PDtrace™ and TCB Usage Guidelines”, MIPS Technologies document MD00365.

[PDTRACETCB]: “The PDtrace™ Interface and Trace Control Block Specification”, MIPS Technologies document MD00439. Current revision is 4.30: you need revision 4 or greater to get multithreading trace information.

Lastly, the “sister” manual to this one describes the non-MT 24K core:

[PROG24K]: “Programming the MIPS32 24K core family”, MIPS Technologies document MD00355.

### Books about programming the MIPS® architecture

[SEEMIPSRUN]: “See MIPS Run”, author Dominic Sweetman, Morgan Kaufmann ISBN 1-55860-410-3. A general and wide-ranging programmers introduction to the MIPS architecture.

[MIPSROG]: “MIPS Programmers Handbook”, Erin Farquar & Philip Bunce, Morgan Kaufmann ISBN 1-55860-297-6. Restricted to the MIPS I instruction set but with a lot of assembler examples.

### Other references

[IEEE754]: “IEEE Standard 754 for Binary Floating-Point Arithmetic”, published by the IEEE, widely available on the web. Surprisingly comprehensible.

### C language header files

These files are available as part of the free-for-download “SDE Lite” subset available from MIPS Technologies’ website. You’ll find them under.../sde/include/mips/.

[m32c0.h]: the C definitions referred to in this manual for the names and fields of standard MIPS32 CP0 registers.

[mt.h]: the C definitions for CP0 registers and other programmable resources of the MIPS MT system.

## Glossary

*ASE*: "Application-Specific Extension" to an instruction set. The acronym is used by MIPS Technologies to describe optional add-ons to the core MIPS32/MIPS64 architecture. The multi-threading package is the "MIPS MT ASE" and there's a bunch of others including the recent "DSP ASE" which adds computational instructions relevant to media-stream signal processing.

*Co-processor*: the MIPS architecture reserves some parts of the instruction set for "co-processors" - which have a few standard instructions, some instruction encoding space and standard registers. Co-processors can be standard but optional (like the floating point unit); a space for customers to build their own logic (like CP2); or, in the case of "co-processor zero", just a way to separate the encodings of critical (and certainly not optional) processor control operations and registers.

*Co-processor zero*: see CP0 below.

*CP0*: MIPS computers use a bunch of register fields for most CPU control purposes. They're accessible only in high-privilege mode, since they're part of the protection system for a protected OS. The registers and the instructions used to access them are defined using a built-in instruction set extension mechanism which conceives of four sets of instruction encodings reserved for "co-processors": the control register set, which must be present in any MIPS32 CPU, are "co-processor zero".

*CP0 hazard*: a hazard which makes some instruction sequences involving privileged operations (and particularly privileged "CP0" registers) illegal. Until quite recently OS programmers were expected to deal with CP0 hazards by inserting "enough" **nop** instructions between producers and consumers of CP0 values and state; but with Revision 2 of [MIPS32] there are better ways described in [Section 7.1, "Hazard barrier instructions"](#).

*Dispatch Scheduler*: the logical block of a MIPS MT multithreading CPU which determines which thread to favor when issuing instructions into the sequential main pipeline.

*EMT*: ("Explicit Multithreading") software which is deliberately written in terms of closely coupled (i.e. memory sharing) concurrent threads. and therefore can directly benefit from multi-threading features of the underlying CPU.

*Gating Storage*: a kind of special uncacheable memory recognized by a multithreading CPU. It's suitable for use for accessing locations where the load/store will not be completed until some event external to the thread, with no obvious maximum waiting time.

From the software's point of view, gating storage is synchronous: no instruction, side-effect or exception from the after the gating load/store is permitted unless and until the load/store completes. A load/store to gating storage may be aborted at any time before it completes, and this will be signalled as a precise exception whose return address is the load/store instruction<sup>1</sup>.

From the hardware's point of view, gating storage has a special interface to the core. The storage subsystem must signal a completed store, and the core can (at any time while waiting for a load/store to complete) ask the storage subsystem to abort the operation. An aborted operation must be "as if it never happened".

There will be some handshaking between the core and the storage subsystem to avoid a race condition between completion and abort. In some circumstances, software trying to abort a gated load/store will fail, and will be told

---

1. Or the branch instruction in whose delay slot the load/store lives - usual MIPS exception rules.

that the operation completed before it could be aborted, and will then have to cope with whatever side-effects the operation had.

*Hazard:*(or “pipeline hazard”) - an architectural requirement which requires you to avoid some instruction sequences. Historical MIPS CPUs had some interesting hazards (like the “load delay slot” and an exception corner case on multiply operations). For a long time MIPS CPUs have only had hazards on code sequences using privileged operations, see CPO hazard.

*Interrupt exempt:* in a MIPS MT CPU like the 34K core a TC may be marked as interrupt-exempt by setting `TCStatus[IXMT]`; then any interrupt presented to the VPE will never cause an exception to that TC. If all TCs belonging to a VPE are marked interrupt-exempt, that’s yet another way of disabling all interrupts.

*Inter-Thread Communication storage:*a generalized form of empty/full storage provided with the 34K core, and attaching to the gating storage (see above) interface. It’s described in [Section 3.3, "Inter-thread communication storage \(ITC\)"](#).

ITC: short for “Inter-Thread Communication storage” as above.

*ITC Cell:*one location of ITC storage. A cell stores 32 bits of data, but has multiple views at different memory locations, each of which behaves differently.

*Pipeline hazard:*see [Hazard](#) above.

*Redirect:*what happens in the pipeline when the 34K core encounters an unpredictable or wrongly-predicted branch instruction. The branch address and condition are finally available by the end of the "EX" pipeline stage (see [Section 3-1, "The 34K™ core pipeline"](#)); at this point all instructions in the pipeline or fetch unit for this thread must be discarded, and instructions fetched from the now-correct instruction instead. That’s a redirect.

*Relax:*used for the extra "bogus TC" on the 34K core which does nothing. The external thread scheduling "policy manager" (see below) has "relax" signals alongside those for real threads; when the "relax" condition has higher priority than any running threads the CPU does nothing for a cycle. This is a way of turning down the CPU (possibly saving energy) when no thread is urgent. See [Section 3.2.3, "Policy managers available for the 34K™ core family"](#).

*Shadow register set:*an extra set of general-purpose registers which can be automatically used in an interrupt handler (or other exception handler). Applications on MIPS32 architecture CPUs can use these shadow registers to reduce the overhead of interrupt handlers, both by retaining quickly-used state in the shadow registers and by avoiding the need to save and restore the state of the interrupted thread. See [Section 7.3, "Shadow registers"](#).

For software compatibility, the 34K core can recycle one or more otherwise-unused TCs’ registers as a shadow set; see [Section 4.4, "TCs recycled as Shadow registers"](#).

*Skid buffer:*in a busy multi-threading CPU threads will block very frequently. When a thread blocks there may well be later instructions from the same thread in the pipeline: you can’t stop the pipeline without holding up all the other threads, and you can’t let this thread’s later instructions complete until this thread is unblocked. So those instructions must be discarded. It would be a problem if we had a full [Redirect](#) every time a thread blocked, so the 34K core’s instruction fetch unit incorporates a "skid buffer" for each thread, which remembers the last couple of instructions issued. When a thread blocks and instructions are discarded from the main pipeline, the skid buffer can be backed up ready for the thread to be unblocked without having to fetch a whole lot more instructions.

*TC:* the logic and registers implementing a minimal thread state in the MIPS MT ASE (from "Thread Context"). A TC has at least its own PC, general-purpose registers and some other necessary bits and pieces. One or more TCs accessing the same complete set of CPO registers make up a [VPE](#).  
The "Tera" project used the word "stream" for this.

*Thread:*a computation consisting of a set of computer instructions read and activated in their programmed order.

Operating systems often use the word "thread" specifically for application-software-visible explicit threads scheduled by the operating system kernel. But any code which is entered by something other than an application-programmed branch forms a separate thread by this definition: an interrupt handler, for example.

*Thread context*: the complete state of a computation as held within the CPU. The thread state excludes (1) data stored in memory, (2) state which is inaccessible to the instruction stream (such as CP0 register contents as seen by a user task) and (3) state which is insignificant (such as cache contents, which generally make no difference to the underlying memory image).

What comprises the thread state varies according to what sort of software is running. For a Linux OS interrupt handler thread on a conventional MIPS CPU the CP0 registers are part of the thread context, but for a Linux application thread they're not. The thread state always (of course) includes the "program counter" ("PC").

*Policy Manager (PM)*: an implementation-dependent piece of logic (located outside of the MIPS core) which receives thread scheduling information from the CPU and hints from the *TCSchedule/VPESchedule* registers, and uses those and other customer-chosen inputs to propose a priority for the various TCs. The interface is designed to permit the policy manager to substantially define scheduling strategy, without the system being prone to failures caused by the inevitable delay between thread events and the PM's response to them reaching the in-core thread scheduler.

*Program Counter (PC)*: A software concept - the address of the next instruction that the thread will execute. It's realization in hardware is somewhat elusive in a pipelined CPU implementing the MIPS architecture. However, it makes a comeback as a hardware-visible thing with the MIPS MT ASE; it is well-defined in hardware for any thread loaded into a TC but which is currently stopped (that is, there are no non-speculative instructions in flight). Such threads keep their PC in the *TCRestart* register.

*Virtualizable*: a CPU feature which can be allocated from a user-privilege program and (transparently to the user program) provided by either the hardware or automatic OS assistance.

So when an OS offers "virtual memory" there's memory which is accessible by the user program - but when there isn't enough memory the user program wanders off the ready-mapped pages, generates an exception which the OS can catch and map some more memory before restarting the application (back exactly where it was when it tried to reference the memory which wasn't there).

MIPS MT resources - notably the TC which runs a concurrent thread - are defined to be virtualizable too. User programs can do their own thread creation and termination using the **fork/yield \$0** instructions, with an OS intervening when no TC is available.

*VSMP*: a system with multiple concurrent threads running in separate VPEs (see the next entry), which behaves much like a multi-CPU system sharing memory with coherent caches (a "symmetric multiprocessor" or SMP system).

*Virtual Processing Element* : see VPE, next

*VPE*: one or more *TCs* sharing a bank of CP0 registers and privileged-architecture resources make up a VPE. The "Tera" project called this a "team".

A single TC running in its own VPE - as seen by software unaware of the MIPS MT ASE - looks like an independent CPU compliant with the MIPS32/MIPS64 specifications. So you can run legacy software (including any OS) which is compatible with the MIPS architecture on a VPE even though the legacy software knows nothing about multi-threading.

*Yield Qualifier*: a signal presented to the core interface which is available for test by the **yield** instruction; see [Section 2.8.1, "Yield, Yield Qualifiers and threads waiting for hardware events"](#).





## CP0 register summary and reference

This appendix lists all the CP0 registers of the 34K core. You can find registers by name through [Table C.1](#), by number through [Table C.2](#) and there's our best shot at functional groupings below, under the heading [Section C.3 "CP0 registers by function"](#). The registers-by-number [Table C.2](#) tells you where to find a detailed description - if you're reading on-line it's a hot-link.

### C.0.0.1 Power-up state of CP0 registers

The traditions of the MIPS architecture regard it as software's job to initialize CP0 registers. As a rule, only fields where a wrong setting would prevent the CPU from booting are forced to an appropriate state by reset; other fields - including other fields in the same register - are random. This manual documents where a field has a forced-from-reset value; but your rule should be that all CP0 registers should be initialized unless you are quite sure that a random value will be harmless.

### C.0.0.2 A note on unused fields in CP0 registers

Unused fields in registers are marked either with a digit 0 or an "X". A field marked zero is guaranteed to read zero on cores in the 34K family. Unless stated otherwise, it's usually best to write it either as zero or with a value you previously read from it. A field marked "X" may return any value, and nothing you write there will have any effect.

## C.1 CP0 registers by name

**Table C.1 CP0 registers by name**

<i>Register Name</i>	<i>Number</i>	<i>Register Name</i>	<i>Number</i>	<i>Register Name</i>	<i>Number</i>	<i>Register Name</i>	<i>Number</i>
BadVAddr	8.0	EntryLo0	2.0	PerfCtl0	25.0	TraceBPC	23.4
CacheErr	27.0	EntryLo1	3.0	PerfCtl1	25.2	TraceControl	23.1
Cause	13.0	EPC	14.0	PerfCtl2	25.4	TraceControl2	23.2
Compare	11.0	ErrCtl	26.0	PerfCtl3	25.6	UserTraceData	23.3
Config	16.0	ErrorEPC	30.0	PRId	15.0	VPEConf0-1	1.2-3
Config1	16.1	HWREna	7.0	Random	1.0	VPEControl	1.1
Config2	16.2	IDataHi	29.1	SRSCnf0-4	6.1-5	VPEOpt	1.7
Config3	16.3	IDataLo	28.1	SRSCtl	12.2	VPEScheFBack	1.6
Config7	16.7	Index	0.0	SRSSMap	12.3	VPESchedule	1.5
Context	4.0	IntCtl	12.1	Status	12.0	WatchHi0	19.0
Count	9.0	ITagLo	28.0	TCBind	2.2	WatchHi1	19.1
DDataLo	28.3	LLAddr	17.0	TCContext	2.5	WatchHi2	19.2
DEPC	24.0	MVPCnf0-1	0.2-3	TCHalt	2.4	WatchHi3	19.3
DESAVE	31.0	MVPCntrol	0.1	TCRestart	2.3	WatchLo0	18.0
Debug	23.0	PageMask	5.0	TCScheFBack	2.7	WatchLo1	18.1
DTagLo	28.2	PerfCnt0	25.1	TCSchedule	2.6	WatchLo2	18.2
EBase	15.1	PerfCnt1	25.3	TCStatus	2.1	WatchLo3	18.3
EntryHi	10.0	PerfCnt2	25.5			Wired	6.0
		PerfCnt3	25.7			YQMask	1.4

## C.2 CP0 registers by number

Registers which are new with the multi-threading extension to the MIPS Architecture are marked with a dagger “†”.

**Table C.2 Cross-referenced of CP0 registers by number**

<i>Register No./Set</i>	<i>Register Name</i>	<i>Function</i>	<i>Refer to</i>
0.0	Index	Index into the TLB array	TLB indexing, p 86
0.1	MVPCntrol	CPU-wide multithreading control	Figure 2-5 , p. 38
0.2-3	MVPCnf0-1	CPU’s multithreading resources	Figure 2-4 , p. 37
1.0	Random	Randomly generated index into the TLB array	TLB indexing, p 86
1.1	VPEControl	VPE control and status	Figure 2-1 , p. 34
1.2-3	VPEConf0-1	Initializable per VPE resource lists	Figure 2-6 , p. 39
1.4	YQMask	Defines valid inputs for <b>yield</b> instruction	Yield etc, p 30
1.5	VPESchedule	Per-VPE thread policy hints	2.9.12 , p. 40
1.6	VPEScheFBack	Per-VPE information from policy manager	
1.7	VPEOpt	Per-VPE cache-way inhibition	Figure 2-7 , p. 40
2.0	EntryLo0	Output (physical) side of TLB entry for even-numbered virtual pages	Figure 6-11 , p. 87
2.1	TCStatus	Status and control for each TC	Figure 2-2 , p. 35
2.2	TCBind	VPE affiliation and own TC number of this TC	Figure 2-3 , p. 37
2.3	TCRestart	Where this TC will next fetch code from	MIPS MT CP0 etc
2.4	TCHalt	Set 1 to freeze the TC for inspection/modification	Table 2.6 , p. 37
2.5	TCContext	Read/write scratch register for OS to maintain thread ID	MIPS MT CP0 etc

Table C.2 Cross-referenced of CP0 registers by number

Register No./Set	Register Name	Function	Refer to
2.6	TCSchedule	Per-TC thread scheduling hints	2.9.12 , p. 40
2.7	TCScheFBack	Per-TC information from policy manager	
3.0	EntryLo1	Output (physical) side of TLB entry for odd-numbered virtual pages	Figure 6-11 , p. 87
4.0	Context	Mixture of pre-programmed and <i>BadVAddr</i> bits which can act as an OS page table pointer.	Figure C-9 , p. 151
5.0	PageMask	Control for variable page size in TLB entries	Figure 6-10 , p. 86
6.0	Wired	Controls the number of fixed ("wired") TLB entries	TLB indexing, p 86
6.1-5	SRSCnf0-4	Write these to use TCs as shadow registers	Figure 7-4 , p. 95
7.0	HWREna	Select which hardware registers are readable using the <b>rdhwr</b> instruction in user mode.	H/W registers, p 122
8.0	BadVAddr	Reports the address for the most recent TLB-related exception	Exception control... and TLB address registers, p 151
9.0	Count	Free-running counter at pipeline or sub-multiple speed	Count/Compare p 145
10.0	EntryHi	High-order portion of the TLB entry	Figure 6-10 , p. 86
11.0	Compare	Timer interrupt control	Count/Compare p 145
12.0	Status	Processor status and control	Figure C-1 , p. 142
12.1	IntCtl	Setup for interrupt vector and interrupt priority features.	Figure 7-1 , p. 91
12.2	SRSCtl	Shadow register set selectors	Figure 7-2 , p. 93
12.3	SRSMap	In VI (vectored interrupt) mode, determines which shadow set is used for each interrupt source.	Figure 7-3 , p. 94
13.0	Cause	Cause of last general exception	Figure C-2 , p. 143
14.0	EPC	Restart address from exception (no subfields, not described further in this manual)	[MIPS32]
15.0	PRId	Processor identification and revision	Figure C-3 , p. 146
15.1	EBase	Exception entry point base address and CPU/VPE ID	Figure C-8 , p. 149
16.0	Config	Configuration register	Figure C-4 , p. 146
16.1	Config1	Configuration for MMU, caches etc	Figure C-5 , p. 147
16.2	Config2		
16.3	Config3	Interrupt and ASE capabilities 34K family-specific configuration	Figure C-6 , p. 148 Figure C-7 , p. 149
16.7	Config7		
17.0	LLAddr	Address associated with last <b>ll</b> instruction of the "load-linked/store-conditional" instruction pair.	ll/sc, p 49
18.0	WatchLo0	I-Watchpoint address	Figure C-10 , p. 152
18.1	WatchLo1		
18.2	WatchLo2		
18.3	WatchLo3		
19.0	WatchHi0	I-Watchpoint control	Figure C-10 , p. 152
19.1	WatchHi1		
19.2	WatchHi2		
19.3	WatchHi3		
23.0	Debug	EJTAG Debug register	Figure 8-1 , p. 103
23.1	TraceControl	Control fields for the PDTrace unit.	Figure 8-13 , p. 112
23.2	TraceControl2		
23.3	UserTraceData	Software-generated PDTrace information register	UserTraceData reg, p. 114
23.4	TraceBPC	Additional controls for PDTrace start/stop	Figure 8-14 , p. 114

**Table C.2 Cross-referenced of CP0 registers by number**

<i>Register No./Set</i>	<i>Register Name</i>	<i>Function</i>	<i>Refer to</i>
24.0	DEPC	Restart address from last EJTAG debug exception	EJTAG CP0 registers p 102
25.0 25.1 25.2 25.3 25.4 25.5 25.6 25.7	PerfCtl0 PerfCnt0 PerfCtl1 PerfCnt1 PerfCtl2 PerfCnt2 PerfCtl3 PerfCnt3	Performance counter 0 control Performance counter 0 Performance counter 1 control Performance counter 1 Performance counter 2 control Performance counter 2 Performance counter 3 control Performance counter 3	Figure 8-15 , p. 116
26.0	ErrCtl	Software parity control and test modes for cache RAM arrays	Figure 6-2 , p. 75
27.0	CacheErr	Cache parity exception control and status	Figure 6-1 , p. 74
28.0 28.2 28.4	ITagLo DTagLo TagLo2	Cache tag read/write interface for I- and D-cache respectively ( <i>TagLo2</i> is reserved for L2 cache)	6.4.11 , p. 82
28.1 28.3 29.1	IDataLo DDataLo IDataHi	Low-order data read/write interface for I-, D- and L2 cache respectively...  ... and high-order data for the I-cache, which is only accessible in 64-bit units.	
30.0	ErrorEPC	Restart location from a cache parity error exception	Cache error exceptions, p 74
31.0	DESAVE	Scratch read/write register for EJTAG debug exception handler	EJTAG CP0 registers p 102

## C.3 CP0 registers by function

<i>Basic modes</i>	Status	12.0	<i>EJTAG debug</i>	Debug	23.0
<i>Exception control</i>	Cause	13.0		DEPC	24.0
	EPC	14.0		DESAVE	31.0
	BadVAddr	8.0	<i>PDTrace block</i>	TraceControl	23.1
<i>Timer</i>	Count	9.0		TraceControl2	23.2
	Compare	11.0		UserTraceData	23.3
<i>Configuration</i>	PRId	15.0		TraceBPC	23.4
	Config	16.0	<i>debug/ analysis</i>	PerfCtl0-1	25.0/2
	Config1-3	16.1-3		PerfCnt0-1	25.1/3
	Config7	16.7		WatchHi0-3	19.0-3
	EBase	15.1		WatchLo0-3	18.0-3
	IntCtl	12.1	<i>regulate user-mode access to hardware registers</i>	HWREna	7.0
	SRSCtl	12.2			
	SRSSMap	12.3	<i>Parity/ECC control</i>	CacheErr	27.0
<i>TLB maintenance (only if TLB)</i>	Context	4.0		ErrCtl	26.0
	BadVAddr	8.0		ErrorEPC	30.0
	EntryHi	10.0	<i>Multithreading (global)</i>	MVPCControl	0.1
	EntryLo0	2.0		MVPCConf0-1	0.2-3
	EntryLo1	3.0	<i>(per-VPE)</i>	VPEControl	1.1
	PageMask	5.0		VPEConf0-1	1.2-3
	Index	0.0		YQMask	1.4
	Random	1.0		VPEScheFBack	1.6
Wired	6.0		VPEOpt	1.7	
<i>Cache management</i>	ITagLo	28.0		VPESchedule	1.5
	DTagLo	28.2		SRSCConf0-4	6.1-5
	TagLo2	28.4	<i>(per-TC)</i>	TCStatus	2.1
	IDataLo	28.1		TCBind	2.2
	DDataLo	28.3		TCRestart	2.3
	DataLo2	28.5		TCHalt	2.4
	IDataHi	29.0		TCContext	2.5
				TCSchedule	2.6
			TCscheFBack	2.7	

## C.4 Miscellaneous CP0 register descriptions

Many CP0 registers in the 34K core are already described earlier in this manual, in a relevant section. But those which got missed are described below, to make sure that every CP0 register field is at least mentioned in this manual.

### C.4.1 Status register

The *Status* register is the most basic (and most diverse, for historical reasons) control register in the MIPS architecture, and its fields are squashed into [Figure C-1](#). All fields are writable unless noted otherwise.

**Figure C-1 All Status register fields**

31	28	27	26	25	24	23	22	21	20	19	18	17	16	15	10	9	8	7	6	5	4	3	2	1	0
CU3-0	RP	FR	RE	MX	PX	BEV	TS	SR	NMI	0	CEE	0	IM7-0		KX	SX	UX	KSU	ERL	EXL	IE				
<i>In EIC (external int controller) mode</i>													IPL	IM1-0											

The 34K family *Status* has no non-standard fields - they're all as defined by [\[MIPS32\]](#). Here and elsewhere these field descriptions are fairly terse, and you should read behind this if you're new to the MIPS architecture. Few of the fields in *Status* are guaranteed to be initialized by hardware on a CPU reset; bootstrap code should write a reasonable value to it early on (the same is true of many other CP0 registers, and the rule is "unless you know it's safe to leave it random, initialize it").

A few fields are somewhat core-specific, and they are described at more length.

*CU3-0*: enables for different *Co-processor* instruction sets (replicated per-TC). Writable when such a coprocessor exists. Since no 34K family CPU has a co-processor3, *Status[CU3]* is hard-wired zero.

Setting *Status[CU0]* to 1 has the peculiar effect of allowing privileged instructions to work in user mode; not something a secure OS is likely to allow often.

*RP*: Reduced power - standard field.

It's not connected inside the 34K core, but the state of the RP bit is available on the external core interface as the *SI\_RP* signal. The 34K core uses clocks generated outside the core, and this could be used in your design to slow the input clock(s).

*FR*: if there is a floating point unit, set 0 for MIPS I compatibility mode (which means you have only 16 real FP registers, with 16 odd FP register numbers reserved for access to the high bits of double-precision values).

*RE*: reverse endianness in user mode. Hard-wired to zero in the 34K core, which doesn't provide this feature.

*MX*: write 1 to enable instructions in *either* the MIPS DSP extension to the MIPS architecture, *or* the MDMX™ extension. The two may not be used together, so MDMX will never be available for the 34K core. But for maximum portability you can find out which by looking at *Config3[DSPP]* (1 if MIPS DSP is implemented) and *Config1[MD]* (1 if MIPS MDMX is implemented).

*PX*: see description of *UX* below (but always zero on the 32-bit 34K CPU).

*BEV*: "boot exception vectors" - when 1, relocates all exception vectors to near the reset-time start address. See [Section C.4, "Exception entry points"](#). This bit is automatically set when the CPU is reset.

*TS*: (read-only) records whether there has been any "machine check" exception (caused by duplicate valid TLB entries, generally a rather serious error) since the CPU was reset.

*SR*: MIPS32 architecture "soft reset" bit: the 34K core's interface only supports a full external reset, so this always reads zero.

*NMI*: (read-only) - non-maskable interrupt shares the "reset" handler code, this field reads 1 when it was a NMI event which caused it.

*CEE*: CorExtend Enable: read/write bit. Set zero to disable "CorExtend" user-defined instructions.

Not all CorExtend blocks implement this bit (those that don't are unconditionally enabled). But CorExtend blocks should use this facility if they store internal state and rely on the OS to save/restore the state associated with some particular task. In such blocks, running a CorExtend instruction with *Status[CEE]* set to zero will cause the CPU to take a "CorExtend Unusable" exception - *Cause[ExcCode]* value 17. A suitably aware kernel will catch the exception and use it to note that the task is one which uses CorExtend resources (and therefore will need CorExtend state saved and restored appropriately).

Do not attempt to set this bit if CorExtend is not present.

*IM7-0*: bitwise interrupt enable for the eight interrupt conditions also visible in *Cause[IP7-0]*; except in the "EIC" interrupt mode, see Section 7.2.3, "External Interrupt Controller (EIC) mode". In that case (as shown) the upper six bits become the "interrupt priority level" ("IPL") value in the range 0-63.

*KX, SX, UX*: the MIPS architecture's memory mapping system changes slightly to support 64-bit addressing, and these bits make that change for kernel-, supervisor- and user-privilege code respectively. But the 34K core is a 32-bit CPU, so these are always zero.

*KSU*: execution privilege level - basically user or kernel:

- 0 kernel
- 1 supervisor - not available on 34K cores
- 2 user

Now that the intermediate "supervisor" privilege level is rarely used, this field is often shown as two separate bits, with the bit 4 being called *UM* ("1 for user mode").

*ERL*: "cache parity error exception mode" - which is really a stronger version of the exception mode *Status[EXL]* bit whose description follows...

*EXL*: exception mode bit, set automatically when you first enter an exception handler or upon reset (reset is treated like an exception). MIPS hardware barely supports nested exceptions, so this disables interrupts and software should avoid causing an exception in the early part of the handler<sup>1</sup>.

*IE*: global interrupt enable, 0 to disable all interrupts.

### C.4.2 Exception control: Cause and BadVAddr register

The *BadVAddr* register is set following any address-related exception. TLB (address translation) exceptions are described below; but note that *BadVAddr* is also set by things like wrong alignment of addresses (but *not* on external difficulties like bus error).

Figure C-2 Fields in the Cause register

31	30	29	28	27	26	25	24	23	22	21	16	15	10	9	8	7	6	2	1	0
BD	TI	CE	DC	PCI	0	IV	WP	0	IP7-2			IP1-0	0	ExcCode	0					
<i>In EIC (external int controller) mode</i>											RIPL									

*Cause* tells you about the exception which just happened. Most fields are read-only:

1. There are some very special cases where nested exceptions are permitted, and the architecture specifies some rather special behaviors to support those. But they're beyond the scope of this manual; see [SEEMIPSRUN]; or the [MIPS32] bible.

## CP0 register summary and reference

*BD*: 1 if the exception happened on an instruction in a branch delay slot; in this case *EPC* is set to restart execution at the branch, which is usually the correct thing to do. You need only consult *Cause[BD]* when you need to look at the instruction which caused the exception (perhaps to emulate it).

*Ti*: last interrupt was from the on-core timer (see section below for *Count/Compare*.)

*CE*: if that was a "co-processor unusable" exception, this is the co-processor which you tried to use.

*DC*: (writable) set 1 to disable the *Count* register.

*PCI*: last interrupt was an overflow from the performance counters, see [Section 8.4, "Performance counters"](#).

*IV*: (writable) set 1 to use a special exception entry point for interrupts, see [Section C.4, "Exception entry points"](#). It's quite likely that if you're doing this, you're also using multiple entry points for different interrupt levels; see [Section 7-1, "Fields in the IntCtl register"](#).

*WP*: (writable to zero) - remembers that a watchpoint triggered when the CPU couldn't take the exception because it was already in exception mode (or error-exception mode, or debug mode). Since this bit automatically causes the exception to happen again, it must be cleared by the watchpoint exception handler.

*IP7-0, RIPL*: the current state of the interrupt request inputs. When one of them is active and enabled by the corresponding *Status[IM7-0]* bit, an interrupt may occur.

*IP1-0* are writable, and in fact always just reflect the value written here. They act as software interrupt bits.

When using "EIC" interrupt mode the interpretation of this field changes, hence the alternate name of *RIPL* ("requested interrupt priority level"). In EIC mode this represents a value between 0 and 63, and reflects the code presented on the incoming interrupt lines when the exception happened. For more information see [Section 7.2.3, "External Interrupt Controller \(EIC\) mode"](#).

*ExcCode*: what caused that last exception. Lots of values, listed in [Table C.3](#).

**Table C.3 Exception Code values in Cause[ExcCode]**

<i>Val</i>	<i>Code</i>	<i>What just happened?</i>
0	Int	Interrupt
1	Mod	Store, but page marked as read-only in the TLB
2	TLBL	Load or fetch, but page marked as invalid in the TLB
3	TLBS	Store, but page marked as invalid in the TLB
4	AdEL	Address error on load/fetch or store respectively. Address is either wrongly
5	AdES	aligned, or a privilege violation.
6	IBE	Bus error signaled on instruction fetch
7	DBE	Bus error signaled on load/store (imprecise)
8	Sys	System call, ie <b>syscall</b> instruction executed.
9	Bp	Breakpoint, ie <b>break</b> instruction executed.
10	RI	Instruction code not recognized (or not legal)
11	CpU	Co-processor instruction encoding for co-processor which is not enabled in <i>Status[CU3-0]</i> .
12	Ov	Overflow from trapping form of integer arithmetic instructions.
13	Tr	Condition met on one of the conditional trap instructions <b>teq</b> etc.
14	-	Reserved
15	FPE	Floating point unit exception - more details in <i>FCSR</i> .
16	-	Available for implementation dependent use
17	CeU	CorExtend instruction attempted when not enable by <i>Status[CEE]</i>



**Table C.3 Exception Code values in Cause[ExcCode]**

<i>Val</i>	<i>Code</i>	<i>What just happened?</i>
18	C2E	Reserved for precise Coprocessor 2 exceptions
19-21	-	Reserved
22	MDMX	Tried to run an MDMX instruction but <i>Status[MX]</i> wasn't set (most likely, the CPU doesn't do MDMX)
23	WATCH	Instruction or data reference matched a watchpoint
24	MCheck	"Machine check" - never happens in the 34K core.
25	Thread	Thread-related exception, as described in [MIPSM]; there's a sub-cause field in <i>VPEControl[EXCPT]</i> , as shown in Section 2-1, "Fields in the <i>VPEControl</i> register".
26	DSP	Tried to run an instruction from the MIPS DSP ASE, but it's not enabled (that is, <i>Status[MX]</i> is zero).
27-29	-	Reserved
30	CacheErr	Parity/ECC error somewhere in the core, on either instruction fetch, load or cache refill. In fact you never see this value in <i>Cause[ExcCode]</i> ; but some of the codes in this table including this one can be visible in the "debug mode" of the EJTAG debug unit - see Section 8.1, "EJTAG on-chip debug unit", and in particular the notes on the <i>Debug</i> register in Section 8-1, "Fields in the EJTAG CP0 Debug register".
31	-	Reserved

### C.4.3 Count and Compare

These two 32-bit registers form a useful and flexible timer. *Count* just counts. For the 34K core, that's usually at the full pipeline clock rate. But portable software can discover how fast *Count* counts by reading the "hardware register" called "CCRes", see Section 9.2, "User-mode accessible "Hardware registers"".

You can write *Count* to set a value in it, but it's generally more valuable for an OS to leave it as a free-running counter.

When the value of *Count* coincides with the value in *Compare*, an interrupt is raised. The interrupt is cleared every time *Compare* is written. This is handy:

- For a periodic interrupt, simply advance *Compare* by a fixed amount each time (and check for the possibility that *Count* has overrun it).
- To set a timer for some point in the future, just set *Compare* to an increment more than the current value of *Count*.

The timer interrupt is implemented as an output signal at the core interface; but it's conventional (well, pretty compulsory if you want OS' to work) to return it to the same VPE on an interrupt line - see notes on *IntCtl[IPTI]* below Figure 7-1. However, if you have an "EIC" interrupt controller (see Section 7.2, "MIPS32® Architecture Release 2 - enhanced interrupt system(s)") you'll need to send the timer interrupt all the way out to the interrupt controller and back.

### C.4.4 PRId, Configuration and EBase registers

The *PRId* register identifies the CPU to software. It's appropriately printed as part of the start-up display by any software telling the world about its start-up; but when portable software is configuring itself around different CPU attributes, it's always preferable to sense those attributes directly - look in other *Config* registers, or perhaps a directed software probe.

**Figure C-3 34K™ processor ID (PRId) register**

31	24 23	16 15	8 7	5 4	2 1	0
CompanyOpts 0	CompanyID 1	ProcessorID 0x95	Major	Revision Minor	Patch	

All these fields are read-only:

*CompanyOpts*: is as specified by the SoC builder who synthesizes the core. It should be a number between 0 and 127 - higher values are reserved by MIPS Technologies.

*CompanyID*: is 1 for MIPS Technologies

*ProcessorID*: is 0x95 for the 34K core.

*Revision*: shows the version of the core. This number (divided into Major, Minor and Patch level as shown) is referenced in the Errata Sheet provided to customers from time to time by MIPS Technologies. The following revisions have been shipped to date:

Release Identifier	PRId[Revision] Maj.min.patch/hex	Description	Date
1_0_*	0.1.1 / 0x5	Early Access (EA) release	June 13, 2005
2_0_*	1.0.0 / 0x20	General Availability (GA) release	September 30, 2005
2_1_*	2.1.0 / 0x44	MR1 release. Bug fixes, 8KB cache support, OCP Resync Support.	March 10, 2006

**Figure C-4 Fields in the Config register**

31	30	28	27	25	24	23	22	21	20	19	18	17	16	15	14	13	12	10	9	7	6	4	3	2	0
Config	M	K23	KU	ISP	DSP	UDI	SB	0	0	MM	0	BM	BE	AT	AR	MT	0	VI	K0						

Figure C-4 shows the fields of the *Config* register, which mixes read-only and writable fields:

*M*: reads 1 if *Config1* is available.

*K23*, *KU*: (writable) if your 34K core-based system uses fixed mapping instead of having a TLB, you set the cacheability attributes of chunks of the memory map by writing these fields. If you have a TLB, these fields are unused (write only zeroes to them).

*Config[K23]* is for program addresses 0xC000.0000-0xFFFF.FFFF (the "kseg2" and "kseg3" areas), while *Config[KU]* is for program addresses 0x0000.0000-0x7FFF.FFFF (the "kuseg" area)

Down at the bottom of the register *Config[K0]* sets the cacheability of kseg0, but it would be very unusual to make that anything other than cacheable.

*ISP/DSP*: (read-only) reads 1 if I-side/D-side scratchpad (SPRAM) is fitted, see Section 6.5, "Scratchpad memory/SPRAM". (Don't confuse this with the MIPS "DSP" ASE, whose presence is indicated by *Config3[DDSP]*.)

*UDI*: (read-only) reads 1 if your core implements user-defined "CorExtend" instructions, and (if this is a MT CPU) if the CorExtend unit is made available to this VPE by the setting of the *VPEConf0* register.

*SB*: read-only "SimpleBE" bus mode indicator.

If set, means that this core will only do simple partial-word transfers on its OCP interface; that is, the only partial-word transfers will be byte, aligned half-word and aligned word.

If zero, it may generate partial-word transfers with an arbitrary set of bytes enabled (which some memory controllers may not like).

*MM*: writable: set 1 if you want writes resulting from separate store instructions in write-through mode merged into a single (possibly burst) transaction at the interface. Note that the *Config[MM]* bit is not replicated per-VPE (like most CP0 fields): there's only one per CPU and anything written by one VPE affects the other one.

This doesn't affect cache writebacks (which are always whole blocks together) or uncached writes (which are never merged).

*BM*: read-only - tells you whether your bus uses sequential or sub-block burst order; set by hardware to match your system controller.

*BE*: (read-only) 1 for big-endian, 0 for little-endian.

*AT*: (read-only) MIPS32 or MIPS64 compliance On 34K family cores it will read "0", but the possible values are:

- 0 MIPS32
- 1 MIPS64 instruction set but MIPS32 address map
- 2 MIPS64 instruction set with full address map

*AR*: Architecture revision level. On 34K family cores it will read "1", denoting release 2 of the MIPS32 specification.

*MT*: (read-only) MMU type (all MIPS Technologies cores may be configured as type 1 or 3):

- 0 None
- 1 MIPS32/64 compliant TLB
- 2 "BAT" type
- 3 MIPS-standard fixed mapping

*Vt*: (read-only) 0 because no 34K family core has a virtually indexed and virtually tagged I-cache

*K0*: (writeable) is the fixed kseg0 region cached or uncached? And if cached, how exactly does it behave - this field is encoded just like the "cache coherency attribute" field of a TLB entry, as it shows up in the *EntryLo0-1* register.

**Figure C-5 Fields in the Config1-2 registers**

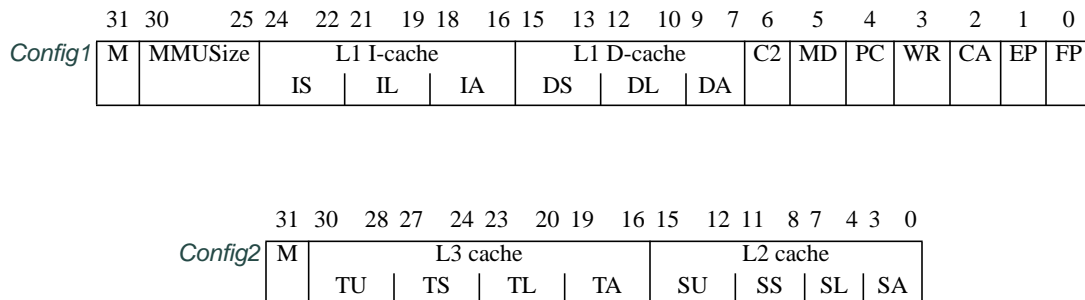


Figure C-5 shows the Config1-2 registers, both of which are wholly read-only:

*Config1[M]*: continuation bit, 1 if *Config2* is implemented.

*Config1[MMUSize]*: the size of the TLB array (the array has MMUSize+1 entries). On the 34K core this is a read-only field which automatically returns the number of entries available to your VPE - unless the TLB is shared, in which case it returns the size of the whole array.

## CP0 register summary and reference

*L1 I-cache, L1 D-cache*: for each cache this reports

- S Number of sets per way. Calculate as:  $64 \times 2^S$
- L Line size. Zero means no cache at all, otherwise calculate as:  $2 \times 2^L$
- A Associativity/number of ways - calculate as  $A + 1$

So if (IS, IL, IA) is (2,4,3) you have 256 sets/way, 32 bytes per line and 4-way set associative: that's a 32Kbyte cache.

*Config1[C2,FP]*: 1 if CP2 or CP1 is available respectively. A coprocessor 2 would be a customer-designed coprocessor, but *FP* selects CP1, the floating point unit. In an MT system these bits reflect whether the units are really available to this VPE, which depends on the setting of *VPEConf0[NCP2,NCP1]*.

*Config1[MD]*: 1 if MDMX ASE is implemented in the floating point unit (very unlikely for the 34K core).

*Config1[PC]*: there is at least one performance counter implemented, see [Section 8.4, "Performance counters"](#).

*Config1[WR]*: reads 1 because the 34K core always has watchpoint registers, see [Figure C-10](#).

*Config1[CA]*: reads 1 because the MIPS16e compressed-code instruction set is available (as it generally is on MIPS Technologies cores).

*Config1[EP]*: reads 1 because an EJTAG debug unit is always provided, see [Section 8.1, "EJTAG on-chip debug unit"](#).

*Config1[FP]*: see entry shared with *Config1[C2]* above.

*Config2[M]*: continuation bit, 1 if *Config3* is implemented.

*Config2[TU]*: implementation-specific bits related to tertiary cache, if fitted. Can be writable.

*Config2[TS,TL,TA]*: tertiary cache size and shape - encoded just like *Config1[IS,IL,IA]* which see above.

*Config2[SU]*: implementation-specific bits for secondary cache, if fitted. Can be writable.

*Config2[SS,SL,SA]*: secondary cache size and shape, encoded like *Config1[IS,IL,IA]* above.

**Figure C-6 Fields in the Config3 register**

31	30	11	10	9	7	6	5	4	3	2	1	0
M	0	DSPP	0	VEIC	VInt	SP	0	MT	SM	TL		

The *Config3* register is wholly read-only. It's fields shown in [Table C-6](#) are:

*Config3[M]*: continuation bit, zero because there is no *Config4*.

*DSPP*: reads 1 if the MIPS DSP extension is implemented, as described in [Chapter 5, "The MIPS32® DSP ASE"](#) on [page 57](#).

*VEIC*: read-only bit from the core input signal *SI\_EICPresent* which should be set in the SoC to alert software to the availability of an EIC-compatible interrupt controller, see [Section 7.2, "MIPS32® Architecture Release 2 - enhanced interrupt system\(s\)"](#). The core interface signal is replicated per-VPE: it is possible (if peculiar) to have only one VPE provided with an EIC-compatible interrupt controller.

*VInt*: reads 1 to tell you that the 34K core can handle vectored interrupts.

*SP*: reads 0 to tell you the 34K core does not support sub-4Kbyte page sizes.

*MT*: reads 1 to tell you the 34K core implements the MIPS MT (multithreading) extension.

*SM*: reads 0, the 34K core does not handle instructions from the "SmartMIPS" ASE.

TL: reads 1 if your core will do instruction trace.

**Figure C-7 Fields in the Config7 Register**

31	19	18	17	16	15	9	8	7	6	5	4	3	2	1	0
0	HCI			FPR	AR	0	ES	0	NBLSU	ULB	BP	RPS	BHT	SL	

*Config7* is a mix of read-only fields (kept in the top 16 bits) and writable fields controlling implementation-dependent options (in the low 16 bits). They are:

*Config7*[HCI]: read-only field which is always zero on 34K family cores. It reads 1 for some software-simulated CPUs, to indicate that the software-modelled cache does not require initialization. Most software should ignore this bit.

*Config7*[FPR]: read-only field. Reads 1 if an FPU is fitted but (as is common) it runs at half the main core clock rate.

*Config7*[AR]: read-only field, indicating that the D-cache is configured to avoid cache aliases (see [Section 6.4.7, "Cache aliases"](#)).

All the remaining fields are read/write, and control various functions. Only one of them is likely to find real system use:

*Config7*[ES]: defaults to zero. If set, the **sync** instruction will be signalled on the core's OCP interface as an "ordering barrier" transaction. The transaction is an extension to the OCP standards, and system controllers which don't support it will typically under-decode it as a read from the boot ROM area. But that's going to be quite slow: so set this bit only if your system understands the synchronizing transaction. This option may be set only for the whole CPU: setting it for one VPE sets it for the other.

The remaining writable fields default to zero and are uncommonly set. It is therefore always safe *not* to write *Config7*. Some of these bits are for diagnostics and experimentation only:

*Config7*[NBLSU]: set 1 to arrange that load/store pipeline stalls will stop the main pipeline too, keeping them synchronized. For debug and investigation only.

*Config7*[ULB]: set 1 to make all uncached loads blocking (a program usually only blocks when it uses the data which is loaded). You want to do this only when nothing else will work...

*Config7*[BP]: when set, no branch prediction is done, and all branches and jump stall as above.

*Config7*[RPS]: when set, the return address branch predictor is disabled, so **jr \$31** is treated just like any other jump register. Instruction fetch stalls after the branch delay slot, until the jump instruction reaches the "EX" stage in the pipeline and can provide the right address (typically adds 5 clocks compared to a successfully predicted return address).

*Config7*[BHT]: when set, the branch history table is disabled and all branches are predicted taken. This bit is don't care if *Config7*[BP] is set.

*Config7*[SL]: when set, disables non-blocking loads. Normally the 34K core will keep running after a load instruction even if it misses in the D-cache, until the data is used. With this disable bit set, the CPU will stall on any load D-cache miss.

### The EBase register

**Figure C-8 Fields in the EBase register**

31	30	29	12	11	10	9	0
1	0	ExceptionBase				0	CPUNum

*EBase* is primarily supplied for multi-CPU systems (or with a MIPS MT CPU, for systems using multiple VPEs). It does two vital jobs: one is to allow software to know which CPU it's running on and the other is to relocate the exception entry points. The latter is necessary because CPUs sharing a memory map (as SMP CPUs often do and the VPEs inside a MIPS MT CPU are obliged to do) have their exception entry points in kseg0, so they will map to the same physical location, and the CPUs would end up jumping to the same exception handlers.

So in [Figure C-8](#):

1 0: is prefixed to the base address bits to make sure the exception vector ends up in the "kseg0" region, conventionally used for OS code.

*ExceptionBase*: is the base address for the exception vectors, adjustable to a resolution of 4Kbytes. See [Table C.4](#) below for where that leaves all the exception entry points.

That means any or all of your CPUs and/or VPEs can have their own unique exception handlers.

*CPUNum*: on single-threaded CPUs this is just a single "CPU number" field (set by the core interface bus *SI\_CPUNum*, which the SoC designer will tie to some suitable value).

But on MIPS MT CPUs that is augmented by some per-VPE value - typically in the least significant bits - so that each VPE gets a distinct value returned for *EBase[CPUNum]*.

### Exception entry points

The incremental growth of exception entry points has left no one place where all the entry points are summarized; so here's [Table C.4](#). You need to accept that *BASE* is 0x8000.0000 for CPUs without an *EBase* register (or where the software, ignoring the *EBase* register, leaves it at its power-on value); and that otherwise *BASE* is the 4Kbyte-aligned address found in *EBase[ExceptionBase]*.

**Table C.4 Exception entry points**

<i>Memory region</i>	<i>Entry point</i>	<i>Exceptions handled here</i>
EJTAG probe-mapped	0xFF20.0200	EJTAG debug, when mapped to "probe" memory.
ROM-only entry points	0xBFC0.0480	EJTAG debug, when using normal ROM memory.
	0xBFC0.0000	Post-reset and NMI entry point.
ROM entry points (when <i>Status[BEV]=1</i> )	0xBFC0.0200	Simple TLB Refill ( <i>Status[EXL]=0</i> ).
	0xBFC0.0300	Cache Parity Error
	0xBFC0.0400	Interrupt special ( <i>Cause[IV]=1</i> ).
	0xBFC0.0380	All others
"RAM" entry points ( <i>Status[BEV]=0</i> )	BASE+0x100	Cache parity error - in RAM. but always through uncached kseg1 window.
	BASE+0x000	Simple TLB Refill ( <i>Status[EXL]=0</i> ).
	BASE+0x200	Interrupt special ( <i>Cause[IV]=1</i> ).
	BASE+0x200+ . . .	multiple interrupt entry points - seven more in "VI" mode, 63 in "EIC" mode.
	BASE+0x180	All others

### C.4.5 Configuring interrupts - The IntCtl and SRSCtl registers

The *IntCtl* register is defined in [Section 7.2, "MIPS32® Architecture Release 2 - enhanced interrupt system\(s\)"](#).

The registers used for shadow register control and setup (*SRSCtl* and *SRSMap*) are described in [Section , "Selecting shadow sets - SRSCtl"](#).

### C.4.6 TLB registers

Four CP0 registers represent the contents of a TLB entry: *EntryHi*, *EntryLo0-1* and *PageMask*. They're quite intricately tied in to the TLB's operation and are all shown in [Section 6.6, "The TLB and translation"](#) above.


There are also the address-exception registers *BadVAddr* and *Context*. We'll deal with them here.

#### TLB address registers: *BadVAddr* and *Context*

On any address-related exception (including all TLB-related exceptions) *BadVAddr* tells you the virtual address whose translation when wrong (some of the same bits showed up as the value loaded into *EntryHi[VPN2]* and discussed above).

*Context* just contains a mix of pre-programmed and borrowed-from-*BadVAddr* bits, as shown in [Table C-9](#).

**Figure C-9 Fields in the Context register**



The idea is that *Context[PTEBase]* can be set to the base address of a (suitably aligned) page table in memory; then the VPN number is shifted such that each ascending 8Kbyte translation unit generates another step through a page table (assuming that each entry is 2×32-bit words in size - reasonable since you need to store at least the two candidate *EntryLo0-1* values.)

An OS which can accept a page table in this format can contrive that in the time-critical simple TLB refill exception, *Context* automatically points to the right page table entry for the new translation.

This is a great idea, but modern OS' tend not to use it - the demands of portability mean it's too much of a stretch to bend the page table information to fit this model.

### C.4.7 Cache registers

Implementation-dependent, see [Section 6.4.11, "Cache initialization and tag/data registers"](#).

### C.4.8 EJTAG unit registers

see [Section 8.1, "EJTAG on-chip debug unit"](#).


### C.4.9 Watchpoint registers

Watchpoint registers are a debugging aid, allowing you to cause an exception when instructions are fetched (or the CPU loads/stores) from particular virtual addresses.

In many cases it's better to use the more extensive breakpoint/watchpoint facilities provided in the EJTAG debug unit, see [Section 8.1, "EJTAG on-chip debug unit"](#).

Each watchpoint is controlled by a pair of CP0 registers. The 34K core has two instruction watchpoints *WatchLo0/WatchHi0* and *WatchLo1/WatchHi1*; and two data watchpoints *WatchLo2/WatchHi2* and *WatchLo3/WatchHi3*.

**Figure C-10 Fields in the WatchLo/WatchHi registers**



Where:

*VAddr*: the address to match on, with a resolution of a doubleword.

*WatchLo[I]*, *WatchLo[R]*, *WatchLo[W]*: accesses to match: I-fetches, Reads (loads), Writes (stores). The 34K core uses separate I- and D-side watchpoints. So in the I-side watchpoints you'll find that *WatchLo0-1[R]* and *WatchLo0-1[W]* is fixed to zero, while on the D-side *WatchLo2-3[I]* will be zero.

*M*: the *WatchHi[M]* bit is set whenever there is one more watchpoint register pair to find; your software should use it to figure out how many watchpoints there are (and should not rely on this manual for this purpose).

*G*, *ASID*: *WatchHi[ASID]* matches addresses from a particular address space (the "ASID" is like that in TLB entries) - except that you can set *WatchHi[G]* ("global") to match the address in any address space.

*Mask*: implements address ranges. Set bits in *WatchHi[Mask]* to mark corresponding *VAddr* address bits to be ignored when deciding whether this is a match.

*WatchHi[I]*, *WatchHi[R]*, *WatchHi[W]*: read *WatchHi* after a watch exception, and these fields tell you what type of access (if anything) matched.

### C.4.10 Performance counter registers

Performance counters are provided to allow software to monitor the occurrence of events of interest within the core, and can be very useful in analyzing system performance. They're described in [Section 8.4, "Performance counters"](#) above.

### C.4.11 Parity/ECC control

The *ErrCtl* register controls parity protection of the L1 caches (if it was configured in your core in the first place) and is described in [Section 6.3.5, "ErrCtl register"](#) above (and in particular [Figure 6-2](#)).

### C.4.12 Registers added for Multithreading

See [Section 2.9, "Multithreading ASE - CP0 \(privileged\) registers"](#).



# MIPS® Architecture quick-reference sheet(s)

## D.1 General purpose register numbers and names

By ancient convention the general-purpose registers in the MIPS architecture have conventional names which remind you of their standard usage in popular MIPS ABIs. [Table D.1](#) shows those names related to both the "o32" ABI (almost universally used for 32-bit MIPS applications), but also the minor variations in the "n32" and "n64" ABIs defined by Silicon Graphics.

If you're not sure what an ABI is, just read the "o32" column!

**Table D.1 Conventional names of registers with usage mnemonics**


Register Nos	name		use	
\$0	zero	always zero		
\$1	AT	assembler temporary		
\$2-\$3	v0-v1	return value from function		
\$4-\$7	a0-a3	arguments		
	<i>o32</i>		<i>n32/n64</i>	
	<i>name</i>	<i>use</i>	<i>name</i>	<i>use</i>
\$8-\$11	t0-t3	temporaries	a4-a7	more arguments
\$12-\$15	t4-t7		t0-t3	temporaries
\$24-\$25	t8-t9		t8-t9	
\$16-\$23	s0-s7	saved registers		
\$26-\$27	k0-k1	reserved for interrupt/trap handler		
\$28	gp	global pointer		
\$29	sp	stack pointer		
\$30	s8/ fp	frame pointer if needed (additional saved register if not)		
\$31	ra	Return address for subroutine		

## D.2 Floating point information

You should read a book on floating point (the MIPS architecture is highly compliant with the IEEE754 standard), or read [\[SEEMIPSRUN\]](#): to understand MIPS floating point well. This section is just bare reference material.

### D.2.1 Data representation

[Figure D-1](#) shows how data is stored in MIPS registers. These are recommended interpretations of the 32-bit and 64-bit formats mentioned in the IEEE standards.

**Figure D-1 How floating point numbers are stored in a register**

Where:

- *sign*: FP numbers are positive numbers with a separate sign bit; "1" denotes a negative number.
- *mantissa*: represents a binary number. But this is a floating point number, so the units depend on:
- *exp*: the exponent.

When 32-bit data is held in a 64-bit register, the high 32 bits are don't care.

Floating point data in memory is endianness-dependent, in just the same way as integer data is; the higher bit-numbered bytes shown in [Section D-1, "How floating point numbers are stored in a register"](#) will be at the lowest memory location when the core is configured big-endian, and the highest memory location when the core is little-endian.

## D.2.2 Setting up the FPU and the FPU control registers

There's a fair amount of state which you set up to change the way the FPU works; this is controlled by fields in the FPU control registers, described here.

[\[IEEE754\]](#) defines five classes of exceptional result. For each class the programmer can select whether to get an IEEE-defined "exceptional result" or to be interrupted. Exceptional results are sometimes just normal numbers but where precision has been lost, but also can be an *infinity* or *NaN* ("not-a-number") value.

Control over the interrupt-or-not options is done through the *FCSR[Enable]* field (or more cleanly through *FENR*, the same control bits more conveniently presented); see [Table D.2](#) below.

It's overwhelmingly popular to keep *FCSR[Enable]* zero and thus never generate an IEEE exception.

There are five FP control registers:

Table D.2 FPU (co-processor 1) control registers

<i>Conventional Name</i>	<i>CP1 ctrl reg no.</i>	<i>Description</i>
FCSR	31	Extensive control register - the only FPU control register on historical MIPS CPUs. Contains <i>all</i> the control bits. But in practice some of them are more conveniently accessed through <i>FCCR</i> , <i>FEXR</i> and <i>FENR</i> below.
FIR	0	FP implementation register: read-only information about the capability of this FPU.
FCCR	25	Convenient partial views of <i>FCSR</i> are better structured, and allow you to update fields without interfering with the operation of independent bits. <i>FCCR</i> has FP condition codes, <i>FEXR</i> contains IEEE exceptional-condition information (cause and flag bits) you read, and <i>FENR</i> is IEEE exceptional-condition enables you write.
FEXR	26	
FENR	28	

### The FP implementation (FIR) register

Figure D-2 shows the fields in *FIR* and the read-only value they always have for 34K family FPUs:

Figure D-2 Fields in the FIR register

	31	25	24	23	22	21	20	19	18	17	16	15	8	7	0
<i>FIR</i>	0	FC	0	F64	L	W	3D	PS	D	S	Processor ID		Revisio.sp.25		
<i>34K core</i>		1		1	1	1	0	0	1	1	0xXX		<i>whatever</i>		

The fields have the following meanings:

*FC*: "full convert range": the hardware will complete *any* conversion operation without running out of bits and causing an "unimplemented" exception.

*F64/L/W/D/S*: this is a 64-bit floating point unit and implements 64-bit integer ("L"), 32-bit integer ("W"), 64-bit FP double ("D") and 32-bit FP single ("S") operations.

*3D*: does not implement the MIPS-3D ASE.

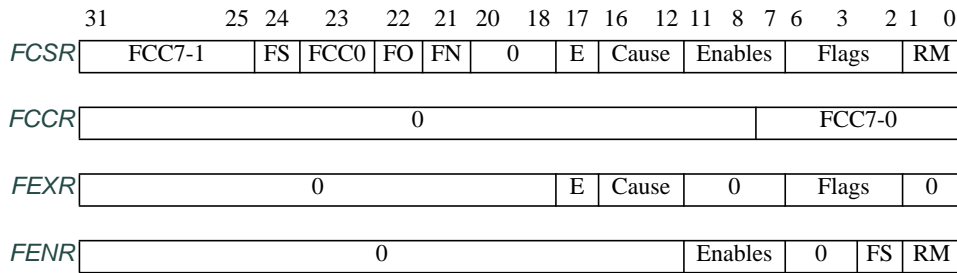
*PS*: does not implement the paired-single instructions described in [MIPS64]

*Processor ID/Revision*: major and minor revisions of the FPU - as is usual with revisions it's very useful to print these out from a verbose sign-on message, and rarely a good idea to have software behave differently according to the values.

### The FP control/status registers (FCSR, FCCR, FEXR, FENR)

Figure D-3 shows all these registers and their bits

**Figure D-3 Floating point control/status register and alternate views**



Where:

*FCC7-0*: the floating point condition codes: set by compare instructions, tested by appropriate branch and conditional move instructions.

*FS/FO/FN*: options to avoid "unimplemented" exceptions when handling tiny ("denormalized") numbers<sup>1</sup>. They do so at the cost of IEEE compatibility, by replacing the very small number with either zero or with the nearest nonzero quantity with a normalized representation.

The *FO* ("flush override") bit causes all tiny operand and result values to be replaced.

The *FS* ("flush to zero") bit causes all tiny operand and result values to be replaced, but additionally does the same substitution for any tiny intermediate value in a multiply-add instruction. This is provided both for legacy reasons, and in case you don't like the idea that the result of a multiply/add can change according to whether you use the fused instruction or a separate multiply and add.

The *FN* bit ("flush to nearest") bit causes all result values to be replaced with somewhat better accuracy than you usually get with *FS*: the result is either zero or a smallest-normalized-number, whichever is closer. Without *FN* set you can only replace your tiny number with a nonzero result if the "RP" or "RM" rounding modes (round towards more positive, round towards more negative) are in effect.


For full IEEE-compatibility you must set  $FCSR[FS,FO,FN] == [0, 0, 0]$ .

To get the best performance compatible with a guarantee of no "unimplemented" exceptions, set  $FCSR[FS,FO,FN] == [1, 1, 1]$ .

Just occasionally for legacy applications developed with older MIPS CPUs which did not have the *FO* and *FN* options, you might set  $FCSR[FS,FO,FN] == [1, 0, 0]$ .

*E*: (often shown in documents as part of the *Cause* array) is a status bit indicating that the last FP instruction caused an "unimplemented" exception.

*Cause/Enables/Flags*: each of these fields is broken up into five bits, each representing an IEEE-recognized class of exceptional results<sup>2</sup> which can be individually treated either by interrupting the computation, or substituting an IEEE-defined exceptional value. So each field contains:



1. See [SEEMIPSRUN]: for an explanation of "normalized" and "denormalized".
2. Sorry about the ugly wording. The IEEE standard talks of "exceptions" which makes more sense but gets mixed up with MIPS "exceptions", and they're not the same thing.

The bits are *V* for invalid operation (e.g. square root of -1), *Z* for divide-by-zero, *O* for overflow (a number too large to represent), *U* for underflow (a number too small to represent without loss of precision) and *I* for inexact - even  $1/3$  is inexact in binary.

Then the:

*Enables*: field is "write 1 to take a MIPS exception if this condition occurs" - rarely done. With the IEEE exception-catcher disabled, the hardware/emulator together will provide a suitable exceptional result.

*Cause*: field records what if any conditions occurred in the last-executed FP instruction. Because that's often too transient, the

*Flags*: field remembers all and any conditions which happened since it was last written to zero by software.

*RM*: is the rounding mode, as required by IEEE:

<i>RM</i>	<i>Meaning</i>
0	Round to nearest - <i>RN</i> If the result is exactly half-way between the nearest values, pick the one whose mantissa bit0 is zero.
1	Round toward zero - <i>RZ</i>
2	Round towards plus infinity - <i>RP</i> "Round up" (but unambiguous about what you do about negative numbers).
3	Round towards minus infinity - <i>RM</i>



## CP0 Registers of the 34K Core

The System Control Coprocessor (CP0) provides the register interface to the 34K processor core and supports memory management, address translation, exception handling, and other privileged operations. Each CP0 register has a unique number that identifies it; this number is referred to as the *register number*. For instance, the *PageMask* register is register number 5. A register may also have a select After updating a CP0 register there is a hazard period of zero or more instructions from the update instruction (MTC0) and until the effect of the update has taken place in the core.

This chapter contains the following sections:

- E.1 “CP0 Register Summary” below
- E.2 “CP0 Register Descriptions” on page 161

### E.1 CP0 Register Summary

Table E.1 lists the CP0 registers in numerical order. The individual registers are described throughout this chapter.

**Table E.1 CP0 Registers**

Register			Function	Per		
Number	Select	Name		VPE	TC	Proc
0	0	Index <sup>1</sup>	Index into the TLB array. This register is reserved if the TLB is not implemented.	X		
1	0	Random <sup>1</sup>	Randomly generated index into the TLB array. This register is reserved if the TLB is not implemented.	X		
2	0	EntryLo0 <sup>1</sup>	Low-order portion of the TLB entry for even-numbered virtual pages. This register is reserved if the TLB is not implemented.			
3	0	EntryLo1 <sup>1</sup>	Low-order portion of the TLB entry for odd-numbered virtual pages. This register is reserved if the TLB is not implemented.	X		
4	0	Context <sup>2</sup>	Pointer to page table entry in memory. This register is reserved if the TLB is not implemented.	X		
5	0	PageMask	PageMask controls the variable page sizes in TLB entries. This register is reserved if the TLB is not implemented.	X		
6	0	Wired <sup>1</sup>	Controls the number of fixed (“wired”) TLB entries. This register is reserved if the TLB is not implemented.	X		
7	0	HWREna	Enables access via the RDHWR instruction to selected hardware registers in non-privileged mode.	X		
8	0	BadVAddr <sup>2</sup>	Reports the address for the most recent address-related exception.	X		
9	0	Count <sup>2</sup>	Processor cycle count.	X		

Table E.1 CP0 Registers (Continued)

Register			Function	Per		
Number	Select	Name		VPE	TC	Proc
10	0	EntryHi <sup>1</sup>	High-order portion of the TLB entry. This register is reserved if the TLB is not implemented.	X	X <sup>3</sup>	
11	0	Compare <sup>2</sup>	Timer interrupt control.	X		
12	0	Status <sup>2</sup>	Processor status and control.	X	X <sup>4</sup>	
12	1	IntCtl <sup>2</sup>	Set-up for interrupt vector and interrupt priority features.	X		
12	2	SRSCtl <sup>2</sup>	Shadow register set selectors	X		
12	3	SRSMap <sup>2</sup>	In vectored interrupt mode, determines which shadow set is used for each interrupt source.	X		
13	0	Cause <sup>2</sup>	Cause of last exception.	X		
14	0	EPC <sup>2</sup>	Program counter at last exception.	X		
15	0	PRId	Processor identification and revision.	X		
15	1	EBase	Exception base address.	X		
16	0	Config	Configuration register.	X		
16	1-2	Config1-2	Configuration for MMU, caches, etc.	X		
16	3	Config3	Interrupt and ASE capabilities	X		
16	7	Config7	34K family-specific configuration register.	X		
17	0	LLAddr	Address associated with last LL instruction of a “load-linked/store-conditional” instruction pair.	X		
18	0-1	WatchLo0-1 <sup>2</sup>	Low-order watchpoint address associated with instruction watchpoints.	X		
18	2-3	WatchLo2-3 <sup>2</sup>	Low-order watchpoint address associated with data watchpoints.	X		
19	0-1	WatchHi0-1 <sup>2</sup>	High-order watchpoint address used for instruction watchpoints.	X		
19	2-3	WatchHi2-3 <sup>2</sup>	High-order watchpoint address used for data watchpoints.	X		
23	0	Debug <sup>5</sup>	EJTAG Debug register.	X		
24	0	DEPC <sup>5</sup>	Restart address from last EJTAG debug exception.	X		
25	0	PerfCtl0	Performance counter 0 control.			X
25	1	PerfCnt0	Performance counter 0.			X
25	2	PerfCtl1	Performance counter 1 control.			X
25	3	PerfCnt1	Performance counter 1.			X
25	4	PerfCtl2	Performance counter 2 control.			X
25	5	PerfCnt2	Performance counter 2.			X
25	6	PerfCtl3	Performance counter 3 control.			X
25	7	PerfCnt3	Performance counter 3.			X
26	0	ErrCtl	Software test enable of way-select and Data RAM arrays for I-Cache and D-Cache.	X		
27	0	CacheErr	Records information about cache parity errors	X		



Table E.1 CP0 Registers (Continued)

Register			Function	Per		
Number	Select	Name		VPE	TC	Proc
28	0	TagLo0	Cache tag read/write interface for I-cache.	X		
28	1	DataLo0	Low-order data read/write interface for I-cache.	X		
28	2	TagLo1	Cache tag read/write interface for D-cache.	X		
28	3	DataLo1	Low-order data read/write interface for D-cache.	X		
28	4	TagLo2	Cache tag read/write interface for L2-cache.	X		
28	5	DataLo2	Low-order data read/write interface for L2-cache.	X		
29	0	DataHi0	Upper bits for I-cache interface. This is only accessible in 64-bit units.	X		
30	3	ErrorEPC <sup>2</sup>	Program counter at last error.	X		
31	0	DeSAVE <sup>5</sup>	Debug handler scratchpad register.	X		

1. Registers used in memory management.
2. Registers used in exception processing.
3. *ASID* per-TC. See “EntryHi Register (CP0 Register 10, Select 0)” on page 168.
4. *KSU* and *CU0-3* per-TC. See “Status Register (CP0 Register 12, Select 0)” on page 170.
5. Registers used in debug.

## E.2 CP0 Register Descriptions

The CP0 registers provide the interface between the ISA and the architecture. Each register is discussed below, with the registers presented in numerical order, first by register number, then by select field number.

For the read/write properties of the field, the following notation is used:

Table E.2 CP0 Register Field Types

Notation	Hardware Interpretation	Software Interpretation
R/W	A field in which all bits are readable and writable by software and, potentially, by hardware. Hardware updates of this field are visible by software reads. Software updates of this field are visible by hardware reads. If the reset state of this field is “Undefined,” either software or hardware must initialize the value before the first read will return a predictable value. This should not be confused with the formal definition of <b>UNDEFINED</b> behavior.	
R	A field that is either static or is updated only by hardware. If the Reset State of this field is either “0” or “Preset”, hardware initializes this field to zero or to the appropriate state, respectively, on power-up. If the Reset State of this field is “Undefined”, hardware updates this field only under those conditions specified in the description of the field.	A field to which the value written by software is ignored by hardware. Software may write any value to this field without affecting hardware behavior. Software reads of this field return the last value updated by hardware. If the Reset State of this field is “Undefined,” software reads of this field result in an <b>UNPREDICTABLE</b> value except after a hardware update done under the conditions specified in the description of the field.
W	A field that can be written by software but which can not be read by software. Software reads of this field will return an <b>UNDEFINED</b> value.	

**Table E.2 CP0 Register Field Types (Continued)**

Notation	Hardware Interpretation	Software Interpretation
0	A field that hardware does not update, and for which hardware can assume a zero value.	A field to which the value written by software must be zero. Software writes of non-zero values to this field may result in <b>UNDEFINED</b> behavior of the hardware. Software reads of this field return zero as long as all previous software writes are zero. If the Reset State of this field is “Undefined,” software must write this field with zero before it is guaranteed to read as zero.

**E.2.1 Index Register (CP0 Register 0, Select 0)**

The *Index* register is a 32-bit read/write register that contains the index used to access the TLB for TLBP, TLBR, and TLBWI instructions. The width of the index field is implementation-dependent as a function of the number of TLB entries that are implemented. The minimum value for TLB-based MMUs is  $Ceiling(Log_2(TLBEentries))$ .

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

This register is only valid with the TLB. It is reserved if the FM is implemented.

**Figure E-1 Index Register Format**



**Table E.3 Index Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
P	31	Probe Failure. Set to 1 when the previous TLBProbe (TLBP) instruction failed to find a match in the TLB.	R/W	Undefined
0	30:6	Must be written as zeros; returns zeros on reads.	0	0
Index	5:0	Index to the TLB entry affected by the TLBRead and TLBWrite instructions. For 16 or 32 entry TLBs, behavior is undefined if index points to a non-existent entry.	R/W	Undefined

**E.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 which 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.

This register is only valid with the TLB. It is reserved if the FM is implemented.

**Figure E-2 Random Register Format**




**Table E.4 Random Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
0	31:6	Must be written as zero; returns zero on reads.	0	0
Random	5:0	TLB Random Index	R	TLB Entries - 1

### E.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. These registers are only valid when the TLB-based memory management unit is present. They are reserved if the FM-style MMU is present.

**Figure E-3 EntryLo0, EntryLo1 Register Format**



**Table E.5 EntryLo0, EntryLo1 Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
R	31:30	Reserved. Should be ignored on writes; returns zero on reads.	R	0
0	29:26	These 4 bits are normally part of the PFN, however, since 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	0
PFN	25:6	Page Frame Number: Contributes to the definition of the high-order bits of the physical address. The PFN field corresponds to bits 31..12 of the physical address.	R/W	Undefined
C	5:3	Coherency attribute of the page. See <a href="#">Table E.6</a> .	R/W	Undefined

Table E.5 EntryLo0, EntryLo1 Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
D	2	“Dirty” or write-enable bit: Indicates 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: Indicates 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 <i>TLB Invalid</i> 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> registers become the G bit in the TLB entry. If the TLB entry G bit is a one, then the ASID comparisons are ignored during TLB matches. On a read from a TLB entry, the G bits of both <i>EntryLo0</i> and <i>EntryLo1</i> reflect the state of the TLB G bit.	R/W	Undefined

Table E.6 lists the encoding of the *C* field of the *EntryLo0* and *EntryLo1* registers and the *K0* field of the *Config* register.

Table E.6 Cache Coherency Attributes

C[5:3] Value	Cache Coherency Attribute
0	Cacheable, non-coherent, write-through, no write allocate
1	Reserved
2	Uncached
3	Cacheable, non-coherent, write-back, write allocate
6	Reserved
7	Uncached Accelerated

## E.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 E-4 Context Register Format

31	23 22	4 3	0
PTEBase	BadVPN2	0	

Table E.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 Context 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 <sub>31:13</sub> of the virtual address that missed.	R	Undefined
0	3:0	Must be written as zero; returns zero on reads.	0	0

### E.2.5 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 E.9.

This register is only valid with the TLB. It is reserved if the FM is implemented.

Figure E-5 PageMask Register Format



Table E.8 PageMask Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:29, 12:0	Ignored on write; returns zero on read.	R	0
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 E.9 Values for the Mask Field of the PageMask Register

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

**Table E.9 Values for the Mask Field of the PageMask Register (Continued)**

Page Size	Bit															
	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13
256 MByte	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Software may determine which page sizes are supported by writing all ones to the *PageMask* register, then reading the value back. If a pair of bits reads back as ones, the processor implements that page size. The operation of the processor is **UNDEFINED** if software loads the *Mask* field with a value other than one of those listed in Table E.9, even if the hardware returns a different value on read. Hardware may depend on this requirement in implementing hardware structures.

### E.2.6 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 E-6. 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. It is reserved if the FM is implemented.

**Figure E-6 Wired and Random Entries in the TLB**



**Figure E-7 Wired Register Format**



Table E.10 Wired Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
0	31:6	Must be written as zero; returns zero on reads.	0	0
Wired	5:0	TLB wired boundary. For 16 and 32 entry TLBs, behavior is undefined if value is set to a value larger than last TLB entry.	R/W	0

### E.2.7 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 E-8 HWREna Register Format



Table E.11 HWREna Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:4	Must be written with zero; returns zero on read	0	0
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.


### E.2.8 BadVAddr Register (CP0 Register 8, Select 0)

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

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

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

**Figure E-9 BadVAddr Register Format**



**Table E.12 BadVAddr Register Field Description**

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


### E.2.9 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. If enabled, the counter increments every other clock. Setting the DC bit in the *Cause* register to 0 enables counting.

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 E-10 Count Register Format**



**Table E.13 Count Register Field Description**

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

### E.2.10 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. 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 *VPN2* field of the *EntryHi* register is not defined after an address error exception and this field may be modified by hardware during the address error exception sequence. Software writes of the *EntryHi* register (via MTC0) do not cause the implicit write of address-related fields in the *BadVAddr*, *Context* registers.

This register is only valid with the TLB. It is reserved if the FM is implemented.

**Figure E-11 EntryHi Register Format**



**Table E.14 EntryHi Register Field Descriptions**

Fields		Description	Read / Write	Reset State
Name	Bits			
VPN2	31:13	$VA_{31..13}$ 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
0	12:8	Must be written as zero; returns zero on read.	0	0
ASID	7:0	Address space identifier. This field is written by hardware on a TLB read and by software to establish the current ASID value for TLB write and against which TLB references match each entry's TLB ASID field.	R/W	Undefined


### E.2.11 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. Traditionally, this has been done by multiplexing it with hardware interrupt 5 to set interrupt bit IP(7) in the *Cause* register.

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

**Figure E-12 Compare Register Format**



**Table E.15 Compare Register Field Description**

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

## E.2.12 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.

**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.

### E.2.12.1 Operating Modes

#### **Debug Mode**

The processor is operating in Debug Mode if the *DM* bit in the CP0 *Debug* register is a one. If the processor is running in Debug Mode, it has full access to all resources that are available to Kernel Mode operation.

#### **Kernel Mode**

The processor is operating in Kernel Mode when the *DM* bit in the *Debug* register is a zero and any of the following three conditions is true:

- The *KSU* field in the CP0 *Status* register contains 2#00
- The *EXL* bit in the *Status* register is one
- The *ERL* bit in the *Status* register is one

The processor enters Kernel Mode at power-up, or as the result of an interrupt, exception, or error. The processor leaves Kernel Mode and enters User Mode or Supervisor Mode when all of the previous three conditions are false, usually as the result of an ERET instruction.

#### **Supervisor Mode**

The processor is operating in Supervisor Mode when all of the following conditions are true:

- The *DM* bit in the *Debug* register is a zero
- The *KSU* field in the *Status* register contains 2#01
- The *EXL* and *ERL* bits in the *Status* register are both zero

Supervisor mode is not supported with the Fixed Mapping MMU.

### User Mode


The processor is operating in User Mode when all of the following conditions are true:

- The *DM* bit in the *Debug* register is a zero
- The *KSU* field in the *Status* register contains 2#10
- The *EXL* and *ERL* bits in the *Status* register are both zero

#### E.2.12.2 Coprocessor Accessibility

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

**Figure E-13 Status Register Format**



**Table E.16 Status Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
CU3	31	Reserved.	R	0						
CU2	30	Controls access to Coprocessor 2 <table border="1" style="width: 100%; margin-top: 5px;"> <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>This bit can only be written when a coprocessor 2 unit is present. This bit cannot be written and will read as 0 if coprocessor 2 unit is not present.</p>	Encoding	Meaning	0	Access not allowed	1	Access allowed	R/W	Undefined
Encoding	Meaning									
0	Access not allowed									
1	Access allowed									
CU1	29	Controls access to Coprocessor 1 <table border="1" style="width: 100%; margin-top: 5px;"> <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>This bit can only be written when the Floating Point Unit is present (34Kf core); in the 34Kc core, this bit cannot be written and will read as 0.</p>	Encoding	Meaning	0	Access not allowed	1	Access allowed	R/W	Undefined
Encoding	Meaning									
0	Access not allowed									
1	Access allowed									
CU0	28	Controls access to coprocessor 0 <table border="1" style="width: 100%; margin-top: 5px;"> <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 CU0 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	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.	R/W	0						

**Table E.16 Status Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
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 datatype. 64-bit datatypes are stored in even-odd pairs of registers</td> </tr> <tr> <td>1</td> <td>Floating point registers can contain any datatype</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"> <li>• No floating point unit is implemented</li> <li>• 64-bit floating point unit is not implemented</li> </ul>	Encoding	Meaning	0	Floating point registers can contain any 32-bit datatype. 64-bit datatypes are stored in even-odd pairs of registers	1	Floating point registers can contain any datatype	R/W	0
Encoding	Meaning									
0	Floating point registers can contain any 32-bit datatype. 64-bit datatypes are stored in even-odd pairs of registers									
1	Floating point registers can contain any datatype									
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>Enables access to DSP ASE resources. An attempt to execute any DSP ASE instruction before this bit has been set to 1 will cause a DSP State Disabled exception.</p>	R							
R	23	<p>Reserved. This field is ignored on write and read as 0.</p>	R	0						
BEV	22	<p>Controls the location of exception vectors:</p> <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	<p>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 reserved if the TLB is not implemented. Software can only write a 0 to this bit to clear it and cannot force a 0-1 transition</p>	R/W0	0						
SR	20	<p>Indicates that the entry through the reset exception vector was due to a Soft Reset. Soft Reset is not supported on this processor and this bit is not writable and will always read as 0</p>	R	0						

Table E.16 Status Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
NMI	19	<p>Indicates that the entry through the reset exception vector was due to an NMI:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Not NMI (Reset)</td> </tr> <tr> <td>1</td> <td>NMI</td> </tr> </tbody> </table> <p>Software can only write a 0 to this bit to clear it and cannot force a 0-1 transition.</p>	Encoding	Meaning	0	Not NMI (Reset)	1	NMI	R/W0	1 for NMI; 0 otherwise
Encoding	Meaning									
0	Not NMI (Reset)									
1	NMI									
0	18	Must be written as zero; returns zero on read.	0	0						
CEE	17	<p>CorExtend Enable: This bit is sent to the CorExtend block to be used to enable the CorExtend block. The usage of this signal by a CorExtend block is implementation dependent. This bit is reserved if CorExtend is not present.</p>	R/W	Undefined						
R	16	Reserved. Ignored on write and read as zero.	R	0						
IM7..IM2	15:10	<p>Interrupt Mask: Controls the enabling of each of the hardware interrupts. An interrupt is taken if interrupts are enabled and the corresponding bits are set in both the Interrupt Mask field of the <i>Status</i> register and the Interrupt Pending field of the <i>Cause</i> register and the <i>IE</i> bit is set in the <i>Status</i> register.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt request disabled</td> </tr> <tr> <td>1</td> <td>Interrupt request enabled</td> </tr> </tbody> </table> <p>In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled (<math>Config3[VEIC] = 1</math>), these bits take on a different meaning and are interpreted as the <i>IPL</i> field, described below.</p>	Encoding	Meaning	0	Interrupt request disabled	1	Interrupt request enabled	R/W	Undefined
Encoding	Meaning									
0	Interrupt request disabled									
1	Interrupt request enabled									
IPL	15:10	<p>Interrupt Priority Level: In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled (<math>Config3[VEIC] = 1</math>), this field is the encoded (0..63) value of the current IPL. An interrupt will be signalled only if the requested IPL is higher than this value. If EIC interrupt mode is not enabled (<math>Config3[VEIC] = 0</math>), these bits take on a different meaning and are interpreted as the <i>IM7..IM2</i> bits, described above.</p>	R/W	Undefined						
IM1..IM0	9:8	<p>Interrupt Mask: Controls the enabling of each of the software interrupts.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt request disabled</td> </tr> <tr> <td>1</td> <td>Interrupt request enabled</td> </tr> </tbody> </table> <p>In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled (<math>Config3[VEIC] = 1</math>), these bits are writable, but have no effect on the interrupt system.</p>	Encoding	Meaning	0	Interrupt request disabled	1	Interrupt request enabled	R/W	Undefined
Encoding	Meaning									
0	Interrupt request disabled									
1	Interrupt request enabled									
R	7:5	Reserved. This field is ignored on write and read as 0.	R	0						

**Table E.16 Status Register Field Descriptions**

Fields		Description	Read / Write	Reset State										
Name	Bits													
KSU	4:3	<p>This field denotes the base operating mode of the processor. The encoding of this field is:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Base mode is Kernel Mode</td> </tr> <tr> <td>01</td> <td>Base mode is Supervisor Mode</td> </tr> <tr> <td>10</td> <td>Base mode is User Mode</td> </tr> <tr> <td>11</td> <td>Reserved</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>KSU</i> field.</p>	Encoding	Meaning	00	Base mode is Kernel Mode	01	Base mode is Supervisor Mode	10	Base mode is User Mode	11	Reserved	R/W	Undefined
Encoding	Meaning													
00	Base mode is Kernel Mode													
01	Base mode is Supervisor Mode													
10	Base mode is User Mode													
11	Reserved													
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"> <li>• The processor is running in kernel mode</li> <li>• Interrupts are disabled</li> <li>• The ERET instruction will use the return address held in <i>ErrorEPC</i> instead of <i>EPC</i></li> <li>• The lower <math>2^{29}</math> bytes of kuseg are treated as an unmapped and uncached region. This allows main memory to be accessed in the presence of cache errors. The operation of the processor is <b>UNDEFINED</b> if the ERL bit is set while the processor is executing instructions from kuseg.</li> </ul>	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"> <li>• The processor is running in Kernel Mode</li> <li>• Interrupts are disabled.</li> <li>• TLB Refill exceptions use the general exception vector instead of the TLB Refill vector.</li> <li>• <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</li> </ul>	Encoding	Meaning	0	Normal level	1	Exception level	R/W	Undefined				
Encoding	Meaning													
0	Normal level													
1	Exception level													

Table E.16 Status Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
IE	0	Interrupt Enable: Acts as the master enable for software and hardware interrupts: <table border="1" data-bbox="488 405 1060 520"> <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> In Release 2 of the Architecture, this bit may be modified separately via the DI and EI instructions.	Encoding	Meaning	0	Interrupts are disabled	1	Interrupts are enabled	R/W	Undefined
Encoding	Meaning									
0	Interrupts are disabled									
1	Interrupts are enabled									

### E.2.13 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 E-14 IntCtl Register Format




Table E.17 IntCtl Register Field Descriptions

Fields		Description	Read / Write	Reset State																					
Name	Bits																								
IPTI	31:29	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. <table border="1" data-bbox="500 1255 1049 1520"> <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> 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 <i>SI_TimerInt</i> signal is attached. 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.	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 E.17 IntCtl Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State																					
Name	Bits																								
IPPCI	28:26	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. 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.	R	Externally Set																					
VS	9:5	<p>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.</p> <table border="1" style="margin-left: auto; margin-right: 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> <p>All other values are reserved. The operation of the processor is <b>UNDEFINED</b> if a reserved value is written to this field.</p>	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	25:10, 4:0	Must be written as zero; returns zero on read.	0	0																					

**E.2.14 SRSCtl Register (CP0 Register 12, Select 2)**

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

**Figure E-15 SRSCtl Register Format**



**Table E.18 SRSCtl Register Field Descriptions**

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 34K processor are: 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 <b>UNDEFINED</b> if a value larger than the one in this field is written to any of these other fields.</p>	R	Preset



Table E.18 SRSCtl Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
EICSS	21:18	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. If <i>Config3[VEIC]</i> is 0, this field returns zero on read.	R	Undefined
ESS	15:12	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. The operation of the processor is <b>UNDEFINED</b> if software writes a value into this field that is greater than the value in the HSS field.	R/W	0
PSS	9:6	Previous Shadow Set. If GPR shadow registers are implemented, and with the exclusions noted in the next paragraph, this field is copied from the CSS field when an exception or interrupt occurs. An ERET instruction copies this value back into the CSS field if <i>Status[BEV]</i> = 0. 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. The operation of the processor is <b>UNDEFINED</b> if software writes a value into this field that is greater than the value in the HSS field.	R/W	0
CSS	3:0	Current Shadow Set. If GPR shadow registers are implemented, this field is the number of the current GPR set. With the exclusions noted in the next paragraph, this field is updated with a new value on any interrupt or exception, and restored from the PSS field on an ERET. Table E.19 describes the various sources from which the CSS field is updated on an exception or interrupt. 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. The value of CSS can be changed directly by software only by writing the PSS field and executing an ERET instruction.	R	0
0	31:30, 25:22, 17:16, 11:10, 5:4	Must be written as zeros; returns zero on read.	0	0

**Table E.19 Sources for new SRSCtl[CSS] on an Exception or Interrupt**

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

### E.2.15 SRSMap Register (CP0 Register 12, Select 3)


The *SRSMap* register contains 8 4-bit fields that provide the mapping from an vector number to the shadow set number to use when servicing such an interrupt. The values from this register are not used for a non-interrupt exception, or a non-vectored interrupt (*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 *SRSMap* 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 E-16 SRSMap Register Format**



**Table E.20 SRSMap 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

### E.2.16 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  $IP_{1..0}$ , DC, IV, and WP fields, all fields in the *Cause* register are read-only. Release 2 of the Architecture added optional support for an External Interrupt Controller (EIC) interrupt mode, in which  $IP_{7..2}$  are interpreted as the Requested Interrupt Priority Level (RIPL).

Figure E-17 Cause Register Format




Table E.21 Cause Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
BD	31	Indicates whether the last exception taken occurred in a branch delay slot: <table border="1" style="margin: 5px 0; width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 15%;">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> The processor updates <i>BD</i> only if <i>Status[EXL]</i> was zero when the exception occurred.	Encoding	Meaning	0	Not in delay slot	1	In delay slot	R	Undefined
Encoding	Meaning									
0	Not in delay slot									
1	In delay slot									
TI	30	Timer Interrupt. This bit denotes whether a timer interrupt is pending (analogous to the <i>IP</i> bits for other interrupt types): <table border="1" style="margin: 5px 0; width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 15%;">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> The state of the <i>TI</i> bit is available on the external core interface as the <i>SL_TimerInt</i> signal.	Encoding	Meaning	0	No timer interrupt is pending	1	Timer interrupt is pending	R	Undefined
Encoding	Meaning									
0	No timer interrupt is pending									
1	Timer interrupt is pending									
CE	29:28	Coprocessor unit number referenced when a Coprocessor Unusable exception is taken. This field is loaded by hardware on every exception, but is <b>UNPREDICTABLE</b> for all exceptions except for Coprocessor Unusable.	R	Undefined						
DC	27	Disable <i>Count</i> register. In some power-sensitive applications, the <i>Count</i> register is not used and is the source of meaningful power dissipation. This bit allows the <i>Count</i> register to be stopped in such situations. <table border="1" style="margin: 5px 0; width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 15%;">Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Enable counting of <i>Count</i> register</td> </tr> <tr> <td>1</td> <td>Disable counting of <i>Count</i> register</td> </tr> </tbody> </table>	Encoding	Meaning	0	Enable counting of <i>Count</i> register	1	Disable counting of <i>Count</i> register	R/W	0
Encoding	Meaning									
0	Enable counting of <i>Count</i> register									
1	Disable counting of <i>Count</i> register									

**Table E.21 Cause Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State																					
Name	Bits																								
PCI	26	<p>Performance Counter Interrupt: This bit denotes whether a performance counter interrupt is pending (analogous to the <i>IP</i> 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 performance counter interrupt is pending</td> </tr> <tr> <td>1</td> <td>Performance counter interrupt is pending</td> </tr> </tbody> </table> <p>The state of the <i>PCI</i> bit is available on the external core interface as the <i>SI_PCInt</i> signal.</p>	Encoding	Meaning	0	No performance counter interrupt is pending	1	Performance counter interrupt is pending	R	Undefined															
Encoding	Meaning																								
0	No performance counter interrupt is pending																								
1	Performance counter interrupt is pending																								
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>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)																								
WP	22	<p>Indicates that a watch exception was deferred because <i>Status[EXL]</i> or <i>Status[ERL]</i> were a one at the time the watch exception was detected. This bit both indicates that the watch exception was deferred, and causes the exception to be initiated once <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 a 0, thereby causing a 0-to-1 transition. If such a transition is caused by software, it is <b>UNPREDICTABLE</b> whether hardware ignores the write, accepts the write with no side effects, or accepts the write and initiates a watch exception once <i>Status[EXL]</i> and <i>Status[ERL]</i> are both zero.</p>	R/W	Undefined																					
IP7..IP2	15:10	<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>15</td> <td>IP7</td> <td>Hardware interrupt 5</td> </tr> <tr> <td>14</td> <td>IP6</td> <td>Hardware interrupt 4</td> </tr> <tr> <td>13</td> <td>IP5</td> <td>Hardware interrupt 3</td> </tr> <tr> <td>12</td> <td>IP4</td> <td>Hardware interrupt 2</td> </tr> <tr> <td>11</td> <td>IP3</td> <td>Hardware interrupt 1</td> </tr> <tr> <td>10</td> <td>IP2</td> <td>Hardware interrupt 0</td> </tr> </tbody> </table> <p>If EIC interrupt mode is not enabled (<i>Config3[VEIC]</i> = 0), timer interrupts are combined in a system-dependent way with any hardware interrupt. If EIC interrupt mode is enabled (<i>Config3[VEIC]</i> = 1), these bits take on a different meaning and are interpreted as the <i>RIPL</i> field, described below.</p>	Bit	Name	Meaning	15	IP7	Hardware interrupt 5	14	IP6	Hardware interrupt 4	13	IP5	Hardware interrupt 3	12	IP4	Hardware interrupt 2	11	IP3	Hardware interrupt 1	10	IP2	Hardware interrupt 0	R	Undefined
Bit	Name	Meaning																							
15	IP7	Hardware interrupt 5																							
14	IP6	Hardware interrupt 4																							
13	IP5	Hardware interrupt 3																							
12	IP4	Hardware interrupt 2																							
11	IP3	Hardware interrupt 1																							
10	IP2	Hardware interrupt 0																							

Table E.21 Cause Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State									
Name	Bits												
R IPL	15:10	Requested Interrupt Priority Level: If EIC interrupt mode is enabled ( <i>Config3[VEIC]</i> = 1), this field is the encoded (0..63) value of the requested interrupt. A value of zero indicates that no interrupt is requested. If EIC interrupt mode is not enabled ( <i>Config3[VEIC]</i> = 0), these bits take on a different meaning and are interpreted as the <i>IP7..IP2</i> bits, described above.	R	Undefined									
IP1..IP0	9:8	Controls the request for software interrupts: <table border="1" data-bbox="488 594 1062 709"> <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> 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.	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 E.22	R	Undefined									
0	25:24, 21:16, 7, 1:0	Must be written as zero; returns zero on read.	0	0									

Table E.22 Cause Register ExcCode Field

Exception Code Value		Mnemonic	Description
Decimal	Hexadecimal		
0	16#00	Int	Interrupt
1	16#01	Mod	TLB modification exception
2	16#02	TLBL	TLB exception (load or instruction fetch)
3	16#03	TLBS	TLB exception (store)
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. If an SDBBP instruction is executed while the processor is running in EJTAG Debug Mode, this value is written to the <i>Debug[DExcCode]</i> field to denote an SDBBP in Debug Mode.
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

Table E.22 Cause Register ExcCode Field (Continued)

Exception Code Value		Mnemonic	Description
Decimal	Hexadecimal		
14	16#0e	-	Reserved
15	16#0f	FPE	Floating point exception
16	16#10	IS1	Coprocessor 2 implementation specific exception
17	16#11	CEU	CorExtend Unusable
18	16#12	C2E	Precise Coprocessor 2 exception
19-22	16#13-16#16	-	Reserved
23	16#17	WATCH	Reference to <i>WatchHi/WatchLo</i> address
24	16#18	MCheck	Machine check
30	16#1e	CacheErr	Cache error. In normal mode, a cache error exception has a dedicated vector and the Cause register is not updated. If a cache error occurs while in Debug Mode, this code is written to the <i>Debug[DExcCode]</i> field to indicate that re-entry to Debug Mode was caused by a cache error.
31	16#1f	-	Reserved

### E.2.17 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 the MIPS16 ASE, a read of the *EPC* register (via *MFC0*) returns the following value in the destination GPR:

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

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

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

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the 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 E-18 EPC Register Format




Table E.23 EPC Register Field Description

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

### E.2.18 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 E-19 PRId Register Format




Table E.24 PRId Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Company Option	31:24	Implementation specific values	R	Preset
Company ID	23:16	Identifies the company that designed or manufactured the processor. In the 34K this field contains a value of 1 to indicate MIPS Technologies, Inc.	R	1
Processor ID	15:8	Identifies the type of processor. This field allows software to distinguish between the various types of MIPS Technologies processors.	R	0x93

**Table E.24 PRId Register Field Descriptions (Continued)**

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

**E.2.19 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 multi-processor system.

The *EBase* register provides the ability for software to identify the specific processor within a multi-processor system, and allows the exception vectors for each processor to be different, 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 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. Bit 29 of exception base address will be forced to 1 on Cache Error exceptions so the exception handler will be executed from the uncached kseg1 segment.

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 4KBbyte page boundary.

**Figure E-20 EBase Register Format**






Table E.25 EBase Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
1	31	This bit is ignored on write and returns one on read.	R	1
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
CPUNum	9:0	This field specifies the number of the CPU in a multi-processor 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
0	30, 11:10	Must be written as zero; returns zero on read.	0	0

### E.2.20 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. The *K0*, *KU*, and *K23* fields must be initialized by software in the Reset exception handler, if the reset value is not desired.

Figure E-21 Config Register Format — Select 0

31	30	28	27	25	24	23	22	21	20	19	18	17	16	15	14	13	12	10	9	7	6	3	2	0
M	K23	KU	ISP	DSP	UDI	SB	0	MM	0	BM	BE	AT	AR	MT	0	K0								

Table E.26 Config Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
M	31	This bit is hard-wired 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 <a href="#">Table E.27</a> for the field encoding.	FM: R/W TLB: R	FM: 010 TLB: 000						
KU	27:25	This field controls the cacheability of the kuseg and useg address segments in FM implementations. Refer to <a href="#">Table E.27</a> for the field encoding.	FM: R/W TLB: R	FM: 010 TLB: 000						
ISP	24	I-side ScratchPad RAM present	R	Preset						
DSP	23	D-side ScratchPad RAM present	R	Preset						
UDI	22	This bit indicates that CorExtend User Defined Instructions have been implemented.	R	Preset						
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No User Defined Instructions are implemented</td> </tr> <tr> <td>1</td> <td>User Defined Instructions are implemented</td> </tr> </tbody> </table>	Encoding	Description	0	No User Defined Instructions are implemented	1	User Defined Instructions are implemented		
Encoding	Description									
0	No User Defined Instructions are implemented									
1	User Defined Instructions are implemented									

**Table E.26 Config Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State								
Name	Bit(s)											
SB	21	Indicates whether SimpleBE bus mode is enabled. Set via <i>SI_SimpleBE</i> input pin. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No reserved byte enables on OCP interface</td> </tr> <tr> <td>1</td> <td>Only simple byte enables allowed on OCP interface</td> </tr> </tbody> </table>	Encoding	Description	0	No reserved byte enables on OCP interface	1	Only simple byte enables allowed on OCP interface	R	Externally Set		
Encoding	Description											
0	No reserved byte enables on OCP interface											
1	Only simple byte enables allowed on OCP interface											
MM	18	This bit indicates whether write-through merging is enabled in the 32 byte collapsing write buffer. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No Merging</td> </tr> <tr> <td>1</td> <td>Merging allowed</td> </tr> </tbody> </table>	Encoding	Description	0	No Merging	1	Merging allowed	R/W	1		
Encoding	Description											
0	No Merging											
1	Merging allowed											
BM	16	Burst order. Set via <i>SI_SBlock</i> input pin. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Sequential</td> </tr> <tr> <td>1</td> <td>SubBlock</td> </tr> </tbody> </table>	Encoding	Description	0	Sequential	1	SubBlock	R	Externally Set		
Encoding	Description											
0	Sequential											
1	SubBlock											
BE	15	Indicates the endian mode in which the processor is running. Set via <i>SI_Endian</i> input pin. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	Little endian	1	Big endian	R	Externally Set		
Encoding	Description											
0	Little endian											
1	Big endian											
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. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Release 1</td> </tr> <tr> <td>1</td> <td>Release 2</td> </tr> <tr> <td>2:7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	0	Release 1	1	Release 2	2:7	Reserved	R	001
Encoding	Description											
0	Release 1											
1	Release 2											
2:7	Reserved											
MT	9:7	MMU Type: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Standard TLB</td> </tr> <tr> <td>3</td> <td>Fixed Mapping</td> </tr> <tr> <td>0, 2, 4:7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	1	Standard TLB	3	Fixed Mapping	0, 2, 4:7	Reserved	R	Preset
Encoding	Description											
1	Standard TLB											
3	Fixed Mapping											
0, 2, 4:7	Reserved											
K0	2:0	Kseg0 coherency algorithm. Refer to <a href="#">Table E.27</a> for the field encoding.	R/W	010								
0	20:19, 17, 6:3	Must be written as zeros; returns zeros on reads.	0	0								

Table E.27 Cache Coherency Attributes

K0(2:0) Value	Cache Coherency Attribute
0	Cacheable, non-coherent, write-through, no write allocate
1	Reserved
2	Uncached
3	Cacheable, non-coherent, write-back, write allocate
6	Reserved
7	Uncached Accelerated

### E.2.21 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.

The instruction and data cache configuration parameters include encodings for the number of sets per way, the line size, and the associativity. The total cache size for a cache is therefore:

$$\text{Associativity} * \text{Line Size} * \text{Sets Per Way}$$

If the line size is zero, no cache is implemented.

Figure E-22 Config1 Register Format

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

Table E.28 Config1 Register Field Descriptions

Fields		Description	Read / Write	Reset State		
Name	Bit(s)					
M	31	This bit is hard-wired to '1' to indicate the presence of the Config2 register.	R	1		
MMU Size	30:25	This field contains the number of entries in the TLB minus one. The field is read as 0 decimal if the TLB is not implemented	R	Preset		
IS	24:22	This field contains the number of instruction cache sets per way. The corresponding total instruction cache size is shown in parentheses	R	Preset		
					Encoding	Description
					0x0	64 (8KB)
					0x1	128 (16KB)
					0x2	256 (32KB)
					0x3	512 (64KB)
0x4:0x7	Reserved					

**Table E.28 Config1 Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State												
Name	Bit(s)															
IL	21:19	<p>This field contains the instruction cache line size. The cache line size is fixed at 32 bytes when the ICache is present. A value of 0 indicates no ICache.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>No ICache present</td> </tr> <tr> <td>0x1:0x3</td> <td>Reserved</td> </tr> <tr> <td>0x4</td> <td>32 bytes</td> </tr> <tr> <td>0x5:0x7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	0x0	No ICache present	0x1:0x3	Reserved	0x4	32 bytes	0x5:0x7	Reserved	R	Preset		
Encoding	Description															
0x0	No ICache present															
0x1:0x3	Reserved															
0x4	32 bytes															
0x5:0x7	Reserved															
IA	18:16	<p>This field contains the level of instruction cache associativity. This field is fixed at 4-way set associative.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0x0:0x2</td> <td>Reserved</td> </tr> <tr> <td>0x3</td> <td>4-way</td> </tr> <tr> <td>0x4:0x7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	0x0:0x2	Reserved	0x3	4-way	0x4:0x7	Reserved	R	0x3				
Encoding	Description															
0x0:0x2	Reserved															
0x3	4-way															
0x4:0x7	Reserved															
DS	15:13	<p>This field contains the number of data cache sets per way. The corresponding total data cache size is shown in parentheses.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>64 (8KB)</td> </tr> <tr> <td>0x1</td> <td>128 (16KB)</td> </tr> <tr> <td>0x2</td> <td>256 (32KB)</td> </tr> <tr> <td>0x3</td> <td>512 (64KB)</td> </tr> <tr> <td>0x4:0x7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	0x0	64 (8KB)	0x1	128 (16KB)	0x2	256 (32KB)	0x3	512 (64KB)	0x4:0x7	Reserved	R	Preset
Encoding	Description															
0x0	64 (8KB)															
0x1	128 (16KB)															
0x2	256 (32KB)															
0x3	512 (64KB)															
0x4:0x7	Reserved															
DL	12:10	<p>This field contains the data cache line size. The cache line size is fixed at 32 bytes when a Dcache is present. This field reads 0 when a Dcache is not present.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0x0</td> <td>No DCache present</td> </tr> <tr> <td>0x1:0x3</td> <td>Reserved</td> </tr> <tr> <td>0x4</td> <td>32 bytes</td> </tr> <tr> <td>0x5:0x7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	0x0	No DCache present	0x1:0x3	Reserved	0x4	32 bytes	0x5:0x7	Reserved	R	Preset		
Encoding	Description															
0x0	No DCache present															
0x1:0x3	Reserved															
0x4	32 bytes															
0x5:0x7	Reserved															
DA	9:7	<p>This field contains the type of set associativity for the data cache. The associativity is fixed at 4-way.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0x0:0x2</td> <td>Reserved</td> </tr> <tr> <td>0x3</td> <td>4-way</td> </tr> <tr> <td>0x4:0x7</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Description	0x0:0x2	Reserved	0x3	4-way	0x4:0x7	Reserved	R	0x3				
Encoding	Description															
0x0:0x2	Reserved															
0x3	4-way															
0x4:0x7	Reserved															

Table E.28 Config1 Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State	
Name	Bit(s)				
C2	6	Coprocessor 2 present.	R	Preset	
		<b>Encoding</b>			<b>Description</b>
		0			Coprocessor2 not present
		1	Coprocessor2 present		
MD	5	MDMX implemented.	R	0	
PC	4	Performance Counter registers implemented.	R	1	
WR	3	Watch registers implemented.	R	1	
		<b>Encoding</b>			<b>Description</b>
		0			No Watch registers are present
		1	One or more Watch registers are present		
CA	2	Code compression (MIPS16) implemented.	R	1	
		<b>Encoding</b>			<b>Description</b>
		0			No MIPS16 present
		1	MIPS16 is implemented		
EP	1	EJTAG present: This bit is always set to indicate that the core implements EJTAG.	R	1	
FP	0	FPU implemented.	R	Preset	

### E.2.22 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. L2 values reflect the configuration information input from the L2 module. L3 fields are reset to 0 because L3 caches are not supported by the 34K core. All fields in the *Config2* register are read-only.

Figure E-23 Config2 Register Format

31	30	28	27	24	23	20	19	16	15	13	12	8	7	4	3	0
M	TU	TS	TL	TA	SU	SS	SL	SA								

Table E.29 Config2 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
M	31	This bit is hard-wired to '1' to indicate the presence of the Config3 register.	R	1
TU	30:28	Implementation specific tertiary cache control. Tertiary cache not supported	R	0
TS	27:24	Tertiary cache sets per way. Tertiary cache not supported	R	0
TL	23:20	Tertiary cache line size. Tertiary cache not supported	R	0

**Table E.29 Config2 Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State	
Name	Bit(s)				
TA	19:16	Tertiary cache associativity. Tertiary cache not supported	R	0	
SU	15:13	Reserved	R	0	
SS	12:8	Secondary cache sets per way	R	Preset	
		<b>Encoding</b>			<b>Sets Per Way</b>
		0			64
		1			128
		2			256
		3			512
		4			1024
		5			2048
		6			4096
		7			8192
8-15	Reserved				
SL	7:4	Secondary cache line size	R	Preset	
		<b>Encoding</b>			<b>Sets Per Way</b>
		0			No cache present
		1			4
		2			8
		3			16
		4			32
		5			64
		6			128
		7			256
8-15	Reserved				
SA	3:0	Secondary cache associativity	R	Preset	
		<b>Encoding</b>			<b>Sets Per Way</b>
		0			Direct mapped
		1			2
		2			3
		3			4
		4			5
		5			6
		6			7
		7			8
8-15	Reserved				

**E.2.23 Config3 Register (CP0 Register 16, Select 3)**

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




**Table E.30 Config3 Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
TL	0	Trace Logic implemented. This bit indicates whether MIPS trace support is implemented.	R	Preset						
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	Trace logic is not implemented	1	Trace logic is implemented		
Encoding	Description									
0	Trace logic is not implemented									
1	Trace logic is implemented									
0	30:11, 9:7, 3	Must be written as zeros; returns zeros on read	0	0						

**E.2.24 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 core.

**Figure E-25 Config7 Register Format**



**Table E.31 Config7 Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
0	31:1, 15:9, 7	These bits are unused and should be written as 0.	R	0						
FPR	17	Floating Point Ratio: Indicates clock ratio between integer core and floating point unit on 34Kf cores. Reads as 0 on 34Kc cores.	R	Based on HW present						
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>FP clock frequency is the same as the integer clock</td> </tr> <tr> <td>1</td> <td>FP clock frequency is one-half the integer clock</td> </tr> </tbody> </table>	Encoding	Description	0	FP clock frequency is the same as the integer clock	1	FP clock frequency is one-half the integer clock		
Encoding	Description									
0	FP clock frequency is the same as the integer clock									
1	FP clock frequency is one-half the integer clock									
AR	16	Alias removed: This bit indicates that the data cache is organized to avoid virtual aliasing problems. This bit is only set if the data cache config and MMU type would normally cause aliasing - i.e., only for the 32KB data cache and TLB-based MMU.	R	Based on HW present						
ES	8	Externalize Sync: If this bit is set, the SYNC instruction will cause a SYNC specific transactions to go out on the external bus. If this bit is cleared, no transaction will go out, but all SYNC handling internal to the core will still be performed. Refer to SYNC instruction description for more information.	R/W	0						



Table E.31 Config7 Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
NBLSU	5	Non-Blocking LSU: Writing 1 to this field will lock the LSU and ALU pipelines together. This forces LSU pipeline stalls to also stall the ALU pipeline.	R/W	0
ULB	4	Uncached Loads Blocking: Writing 1 to this field will make all uncached loads blocking.	R/W	0
BP	3	Branch Prediction: Writing 1 to this field will disable all speculative branch prediction. The fetch unit will wait for a branch to be resolved before fetching the target or fall-through path.	R/W	0
RPS	2	Return Prediction Stack: Writing 1 to this field will disable the use of the Return Prediction Stack. Returns (JR ra) will stall instruction fetch until the destination is calculated.	R/W	0
BHT	1	Branch History Table: Writing 1 to this field will disable the dynamic branch prediction. Branches will be statically predicted taken.	R/W	0
SL	0	Scheduled Loads: Writing 1 to this field will make load misses blocking.	R/W	0

### E.2.25 LLAddr Register (CP0 Register 17, Select 0)

Figure E-26

### E.2.26 WatchLo Register (CP0 Register 18, Select 0-3)

The *WatchLo* and *WatchHi* registers together provide the interface to a watchpoint debug facility that initiates a watch exception if an instruction or data access matches the address specified in the registers. As such, they duplicate some functions of the EJTAG debug solution. Watch exceptions are taken only if the *EXL* and *ERL* bits are both zero in the *Status* register. If either bit is a one, the *WP* bit is set in the *Cause* register, and the watch exception is deferred until both the *EXL* and *ERL* bits are zero.

There are 4 sets of Watch register pairs (*WatchLo*, *WatchHi*). Two of them (select 0, 1) are associated with instruction addresses only. Thus, only the I bit is writable, the R and W bits are tied to 0. The other two (select 2, 3) are associated with data addresses and can only be used for R or W watchpoints.

The *WatchLo* register specifies the base virtual address and the type of reference (instruction fetch, load, store) to match.

Figure E-27 WatchLo Register Format




Table E.32 WatchLo Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
VAddr	31:3	This field specifies the virtual address to match. Note that this is a doubleword address, since bits [2:0] are used to control the type of match.	R/W	Undefined
I	2	If this bit is set, watch exceptions are enabled for instruction fetches that match the address.	R/W	0
R	1	If this bit is set, watch exceptions are enabled for loads that match the address.	R/W	0
W	0	If this bit is set, watch exceptions are enabled for stores that match the address.	R/W	0

### E.2.27 WatchHi Register (CP0 Register 19, Select 0-3)

The *WatchLo* and *WatchHi* registers together provide the interface to a watchpoint debug facility that initiates a watch exception if an instruction or data access matches the address specified in the registers. As such, they duplicate some functions of the EJTAG debug solution. Watch exceptions are taken only if the *EXL* and *ERL* bits are zero in the *Status* register. If either bit is a one, then the *WP* bit is set in the *Cause* register, and the watch exception is deferred until both the *EXL* and *ERL* bits are zero.

The *WatchHi* register contains information that qualifies the virtual address specified in the *WatchLo* register: an *ASID*, a Global (*G*) bit, and an optional address mask. If the *G* bit is 1, then any virtual address reference that matches the specified address will cause a watch exception. If the *G* bit is a 0, only those virtual address references for which the *ASID* value in the *WatchHi* register matches the *ASID* value in the *EntryHi* register cause a watch exception. The optional mask field provides address masking to qualify the address specified in *WatchLo*.

There are 4 sets of Watch register pairs (*WatchLo*, *WatchHi*). Two of them (select 0, 1) are associated with instruction addresses only. Thus, only the *I* bit is meaningful, the *R* and *W* bits are tied to 0. The other two (select 2, 3) are associated with data addresses and can only be used for *R* or *W* watchpoints.

Figure E-28 WatchHi Register Format

31	30	29	24	23	16	15	12	11	3	2	0	
M	G	0	ASID			0	Mask			I	R	W

Table E.33 WatchHi Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
M	31	Indicates the presence of additional Watch registers.	R	Preset
G	30	If this bit is one, any address that matches that specified in the <i>WatchLo</i> register causes a watch exception. If this bit is zero, the <i>ASID</i> field of the <i>WatchHi</i> register must match the <i>ASID</i> field of the <i>EntryHi</i> register to cause a watch exception.	R/W	Undefined
ASID	23:16	<i>ASID</i> value which is required to match that in the <i>EntryHi</i> register if the <i>G</i> bit is zero in the <i>WatchHi</i> register.	R/W	Undefined

Table E.33 WatchHi Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
Mask	11:3	Bit mask that qualifies the address in the <i>WatchLo</i> register. Any bit in this field that is a set inhibits the corresponding address bit from participating in the address match.	R/W	Undefined
I	2	This bit is set by hardware when an instruction fetch condition matches the values in this watch register pair. When set, the bit remains set until cleared by software, which is accomplished by writing a 1 to the bit.	W1C	Undefined
R	1	This bit is set by hardware when a load condition matches the values in this watch register pair. When set, the bit remains set until cleared by software, which is accomplished by writing a 1 to the bit.	W1C	Undefined
W	0	This bit is set by hardware when a store condition matches the values in this watch register pair. When set, the bit remains set until cleared by software, which is accomplished by writing a 1 to the bit.	W1C	Undefined
0	29:24, 15:12	Must be written as zero; returns zero on read.	0	0

### E.2.28 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 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 from 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* 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 E-29 Debug Register Format

31	30	29	28	27	26	25	24	23	22	21	20	19
DBD	DM	NoDCR	LSNM	Doze	Halt	CountDM	IBusEP	MCheckP	CacheEP	DBusEP	IEXI	DDBSImpr

18	17	15	14	10	9	8	7	6	5	4	3	2	1	0
DDBLImpr	EJTAGver	DExcCode		NoSSt	SSt	R	Offline	DINT	DIB	DDBS	DDBL	DBp	DSS	

**Table E.34 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>Description</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	Description	0	Not in delay slot	1	In delay slot	R	Undefined
Encoding	Description									
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>Description</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	Description	0	Processor is operating in non-debug mode	1	Processor is operating in debug mode	R	0
Encoding	Description									
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: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	dseg is present	1	No dseg present	R	0
Encoding	Description									
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>Description</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	Description	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	Description									
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>Description</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	Description	0	Processor not in low power mode when debug exception occurred	1	Processor in low power mode when debug exception occurred	R	Undefined
Encoding	Description									
0	Processor not in low power mode when debug exception occurred									
1	Processor in low power mode when debug exception occurred									
Halt	26	Indicates that the internal system bus clock was stopped when the debug exception occurred: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	Internal system bus clock stopped	1	Internal system bus clock running	R	Undefined
Encoding	Description									
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"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	Count register stopped in debug mode	1	Count register is running in debug mode	R/W	1
Encoding	Description									
0	Count register stopped in debug mode									
1	Count register is running in debug mode									

Table E.34 Debug Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
IBusEP	24	Imprecise instruction fetch Bus Error exception Pending: All instruction bus errors are precise on the 34K core so this bit will always read as 0. 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 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 instruction fetch is taken by the processor, and <i>IBusEP</i> is cleared.	R	0						
MCheckP	23	Indicates that an imprecise Machine Check exception is pending. Set when a Machine Check exception occurs or if a 1 is written to the bit by software. Cleared when a machine check exception is taken by the processor, and by reset. If <i>MCheckP</i> is set when <i>IEXI</i> is cleared, a Machine Check exception is taken by the processor, and <i>MCheckP</i> is cleared.	R	0						
CacheEP	22	Indicates that an imprecise Cache Error is pending.	R/W1	0						
DBusEP	21	Data access Bus Error exception Pending: Set when an data bus error event occurs or if a 1 is written to the bit by software. Cleared when a Data Bus Error exception is taken by the processor, and by reset. If <i>DBusEP</i> is set when <i>IEXI</i> is cleared, a Data Bus Error exception is taken by the processor, and <i>DBusEP</i> is cleared.	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.	R	0						
DDBLImpr	18	Indicates that an imprecise Debug Data Break Load exception was taken.	R	0						
EJTAGver	17:15	EJTAG version. <table border="1" data-bbox="511 1465 1084 1549"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>3</td> <td>Version 3.x</td> </tr> </tbody> </table>	Encoding	Description	3	Version 3.x	R	011		
Encoding	Description									
3	Version 3.x									
DExcCode	14:10	Indicates the cause of the latest exception in debug mode. See <a href="#">Table E.22</a> for a list of values. Value is undefined after a debug exception.	R	Undefined						
NoSST	9	Indicates whether the single-step feature controllable by the <i>SSr</i> bit is available in this implementation: <table border="1" data-bbox="511 1738 1084 1852"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	Single-step feature available	1	No single-step feature available	R	0
Encoding	Description									
0	Single-step feature available									
1	No single-step feature available									

Table E.34 Debug Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
SSt	8	Controls if debug single step exception is enabled: <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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> <p>This is implemented per TC. Global single-step operation of a VPE can be achieved by setting SSt for all TCs.</p>	Encoding	Description	0	No debug single-step exception enabled	1	Debug single step exception enabled	R/W	0
Encoding	Description									
0	No debug single-step exception enabled									
1	Debug single step exception enabled									
Offline	7	Implemented per-TC. When this bit is 1, TC is allowed to execute only in Debug mode.	R/W	0						
R	6	Reserved. Must be written as zeros; returns zeros on reads.	R	0						
DINT	5	Indicates that a debug interrupt exception occurred. Cleared on exception in debug mode. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	No debug interrupt exception	1	Debug interrupt exception	R	Undefined
Encoding	Description									
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"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	No debug interrupt exception	1	Debug interrupt exception	R	Undefined
Encoding	Description									
0	No debug interrupt exception									
1	Debug interrupt exception									
DDBS	3	Indicates that a debug data break exception occurred on a store. Cleared on exception in debug mode. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</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	Description	0	No debug data exception on a store	1	Debug instruction exception on a store	R	Undefined
Encoding	Description									
0	No debug data exception on a store									
1	Debug instruction exception on a store									
DDBL	2	Indicates that a debug data break exception occurred on a load. Cleared on exception in debug mode. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug data exception on a load</td> </tr> <tr> <td>1</td> <td>Debug instruction exception on a load</td> </tr> </tbody> </table>	Encoding	Description	0	No debug data exception on a load	1	Debug instruction exception on a load	R	Undefined
Encoding	Description									
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. <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug software breakpoint exception</td> </tr> <tr> <td>1</td> <td>Debug software breakpoint exception</td> </tr> </tbody> </table>	Encoding	Description	0	No debug software breakpoint exception	1	Debug software breakpoint exception	R	Undefined
Encoding	Description									
0	No debug software breakpoint exception									
1	Debug software breakpoint exception									

Table E.34 Debug Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State						
Name	Bit(s)									
DSS	0	Indicates that a debug single-step exception occurred. Cleared on exception in debug mode.	R	Undefined						
		<table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No debug single-step exception</td> </tr> <tr> <td>1</td> <td>Debug single-step exception</td> </tr> </tbody> </table>	Encoding	Description	0	No debug single-step exception	1	Debug single-step exception		
Encoding	Description									
0	No debug single-step exception									
1	Debug single-step exception									

## E.2.29 Trace Control Register (CP0 Register 23, Select 1)

The *TraceControl* register configuration is shown below.

Figure E-30 TraceControl Register Format




Table E.35 TraceControl Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
TS	31	The trace select bit is used to select between the hardware and the software trace control bits. A value of zero selects the external hardware trace block signals, and a value of one selects the trace control bits in the <i>TraceControl</i> register.	R/W	0
UT	30	This bit is used to indicate the type of user-triggered trace record. A value of zero implies a user type 1 and a value of one implies a user type 2. The actual triggering of a user trace record happens on a write to the <i>UserTraceData</i> register. This is a 32-bit register for 32-bit processors and a 64-bit register for 64-bit processors.	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 1, 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. Hence, the core tracing logic must not allow a FIFO overflow and discard trace data. This is achieved by stalling the pipeline when the FIFO is nearly full, so that no trace records are ever lost.	R/W	Undefined
D	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, and either the G bit must be one, or the current process ASID must match the ASID field in this register. When set to zero, trace is disabled in Debug Mode, irrespective of other bits.	R/W	Undefined

Table E.35 TraceControl Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
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, and either the G bit must be one, or the current process ASID must match the ASID field in this register. 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, and either the G bit must be one, or the current process ASID must match the ASID field in this register. When set to zero, trace is disabled in Kernel Mode, irrespective of other bits.	R/W	Undefined
S	22	When set to one, this enables tracing in Supervisor Mode. For trace to be enabled in Supervisor mode, the On bit must be one, and either the G bit must be one, or the current process ASID must match the ASID field in this register. When set to zero, trace is disabled in Supervisor Mode, irrespective of other bits. If the processor does not implement Supervisor Mode, this bit is ignored on write and returns zero on read.	R/W	Undefined
U	21	When set to one, this enables tracing in User Mode. For trace to be enabled in User mode, the On bit must be one, and either the G bit must be one, or the current process ASID must match the ASID field in this register. When set to zero, trace is disabled in User Mode, irrespective of other bits.	R/W	Undefined
ASID_M	20:13	This is a mask value applied to the ASID comparison (done when the G bit is zero). A "1" in any bit in this field inhibits the corresponding ASID bit from participating in the match. As such, a value of zero in this field compares all bits of ASID. Note that the ability to mask the ASID value is not available in the hardware signal bit; it is only available via the software control register. If the processor does not implement the standard TLB-based MMU, this field is ignored on write and returns zero on read.	R/W	Undefined
ASID	12:5	The ASID field to match when the G bit is zero. When the G bit is one, this field is ignored. If the processor does not implement the standard TLB-based MMU, this field is ignored on write and returns zero on read.	R/W	Undefined
G	4	When set, this implies that tracing is to be enabled for all processes, provided that other enabling functions (like U, S, etc..) are also true. If the processor does not implement the standard TLB-based MMU, this field is ignored on write and returns 1 on read. This causes all match equations to work correctly in the absence of an ASID.	R/W	Undefined
TFCR	3	When asserted, used to trace function call and return instructions with full PC values.	R/W	Undefined



Table E.35 TraceControl Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bits			
TLSM	2	When asserted, used to trace data cache load and store misses with full PC values, and potentially the data address and value as well.	R/W	Undefined
TIM	1	When asserted, used to trace instruction miss with full PC values.	R/W	Undefined
On	0	This is the master trace enable switch in software control. When zero, tracing is always disabled. When set to one, tracing is enabled whenever the other enabling functions are also true.	R/W	0

### E.2.30 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). As such, these fields in the *TraceControl2* register will not have valid values until the TCB asserts these values.

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

Figure E-31 TraceControl2 Register Format




Table E.36 TraceControl2 Register Field Descriptions

Fields		Description	Read / Write	Reset State
Name	Bits			
0	31:30	Reserved for future use; Must be written as zero; returns zero on read.	0	0
CPUIdV	29	When set, this bit specifies the VPE defined in <i>CPUId</i> must be traced. Otherwise, instructions from all VPEs are traced when other conditions for tracing are valid. This bit is ignored if <i>TCV</i> is asserted.	R/W	
CPUId	28:21	This field specifies the number of the VPE to trace when <i>CPUIdV</i> is set.	R/W	
TCV	20	When set, the <i>TCNum</i> field specifies the number of the TC that must be traced. Otherwise, instructions from all TCs are traced when other conditions for tracing are valid.	R/W	
TCNum	19:12	Specifies the TC to trace when <i>TCV</i> is set. The right-most bits only are used.	R/W	

**Table E.36 TraceControl2 Register Field Descriptions**

Fields		Description	Read / Write	Reset State												
Name	Bits															
Mode	11:7	<p>These 5 bits provide the same trace mode functions as the <i>PDI_TraceMode[4:0]</i> signal, and is described here again. When tracing is turned on, this signal specifies what information is to be traced by the core. It uses 5 bits, where each bit turns on a tracing of a specific tracing when that bit value is a 1. If the corresponding bit is 0, then the Trace Value shown in column two is not traced by the processor. On the 34K core PC tracing is always enabled, regardless of the value on bit 23.ode. The table shows what trace value is turned on:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Trace the Following</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>PC</td> </tr> <tr> <td>1</td> <td>Load address</td> </tr> <tr> <td>2</td> <td>Store address</td> </tr> <tr> <td>3</td> <td>Load data</td> </tr> <tr> <td>4</td> <td>Store data</td> </tr> </tbody> </table>	Bit	Trace the Following	0	PC	1	Load address	2	Store address	3	Load data	4	Store data	R/W	Undefined
Bit	Trace the Following															
0	PC															
1	Load address															
2	Store address															
3	Load data															
4	Store data															
Valid-Modes	6:5	<p>This field specifies the subset 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	Preset		
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, as follows:</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 TBU 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 TBU bit of this register indicates to which trace buffer the traces is currently written.</td> </tr> </tbody> </table> <p>This bit is loaded from the <i>PDI_TBImpI</i> signal when the <i>PDI_SyncOffEn</i> signal is asserted.</p>	Encoding	Meaning	0	Only one trace buffer is implemented, and the TBU bit of this register indicates which trace buffer is implemented	1	Both on-chip and off-chip trace buffers are implemented by the TCB and the TBU bit of this register indicates to which trace buffer the traces is currently written.	R	Undefined						
Encoding	Meaning															
0	Only one trace buffer is implemented, and the TBU bit of this register indicates which trace buffer is implemented															
1	Both on-chip and off-chip trace buffers are implemented by the TCB and the TBU bit of this register indicates to which trace buffer the traces is currently written.															

Table E.36 TraceControl2 Register Field Descriptions


Fields		Description	Read / Write	Reset State																		
Name	Bits																					
TBU	3	<p>This bit denotes to which trace buffer the trace is currently being written and is used to select the appropriate interpretation of the <i>TraceControl2[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> <p>This bit is loaded from the <i>PDI_OffChipTB</i> signal when the <i>PDI_SyncOffEn</i> signal is asserted.</p>	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																					
SyP	2:0	<p>The period (in cycles) to which the internal synchronization counter is reset when tracing is started, or when the synchronization counter has overflowed.</p> <table border="1"> <thead> <tr> <th>SyP</th> <th>Sync Period</th> </tr> </thead> <tbody> <tr> <td>000</td> <td><math>2^5</math></td> </tr> <tr> <td>001</td> <td><math>2^6</math></td> </tr> <tr> <td>010</td> <td><math>2^7</math></td> </tr> <tr> <td>011</td> <td><math>2^8</math></td> </tr> <tr> <td>100</td> <td><math>2^9</math></td> </tr> <tr> <td>101</td> <td><math>2^{10}</math></td> </tr> <tr> <td>110</td> <td><math>2^{11}</math></td> </tr> <tr> <td>111</td> <td><math>2^{12}</math></td> </tr> </tbody> </table> <p>This field is loaded from the <i>PDI_SyncPeriod</i> signal when the <i>PDI_SyncOffEn</i> signal is asserted.</p>	SyP	Sync Period	000	$2^5$	001	$2^6$	010	$2^7$	011	$2^8$	100	$2^9$	101	$2^{10}$	110	$2^{11}$	111	$2^{12}$	R	Undefined
SyP	Sync Period																					
000	$2^5$																					
001	$2^6$																					
010	$2^7$																					
011	$2^8$																					
100	$2^9$																					
101	$2^{10}$																					
110	$2^{11}$																					
111	$2^{12}$																					

### E.2.31 User Trace Data Register (CP0 Register 23, Select 3)

A software write to any bits in the *UserTraceData* register will trigger a trace record to be written indicating a type 1 or type 2 user format. The type is based on the *UT* bit in the *TraceControl* register. This register cannot be written in consecutive cycles. The trace output data is **UNPREDICTABLE** if this register is written in consecutive cycles.

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

Figure E-32 User Trace Data Register Format



**Table E.37 UserTraceData 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

**E.2.32 TracelBPC Register (CP0 Register 23, Select 4)**

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

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

**Figure E-33 TracelBPC Register Format**



**Table E.38 TracelBPC Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
0	31:29, 27:12	Reserved for future implementation	R	0/1						
IE	28	Used to specify whether the trigger signal from EJTAG instruction breakpoint should trigger tracing functions or not: <table border="1" style="margin: 5px auto; width: 80%;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Disables trigger signals from instruction breakpoints</td> </tr> <tr> <td>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									
IBPC <sub>n</sub>	3n-1:3n-3	The three bits are decoded to enable different tracing modes. <a href="#">Table E.40</a> shows the possible interpretations. Each set of 3 bits represents the encoding for the instruction breakpoint n in the EJTAG implementation, if it exists. If the breakpoint does not exist then the bits are reserved, read as zero and writes are ignored. If bit 27 is zero, bits 3n-1:3n-2 are ignored, and only the bottom bit 3n-3 is used to start and stop tracing as specified in versions less than 4.00 of this specification.	R/W	0						

**E.2.33 TraceDBPC Register (CP0 Register 23, Select 5)**

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

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

Figure E-34 TraceDBPC Register Format

31	29	28	27	6	5	3	2	0
0	DE	0				DBPC <sub>1</sub>	DBPC <sub>0</sub>	

Table E.39 TraceDBPC Register Field Descriptions

Fields		Description	Read / Write	Reset State						
Name	Bits									
0	31:29, 27:6	Reserved for future implementation	R	0/1						
DE	28	Used to specify whether the trigger signal from EJTAG data breakpoint should trigger tracing functions or not: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>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									
DBPC <sub>n</sub>	3n-1:3n-3	The three bits are decoded to enable different tracing modes. <a href="#">Table E.40</a> shows the possible interpretations. Each set of 3 bits represents the encoding for the data breakpoint n in the EJTAG implementation, if it exists. If the breakpoint does not exist then the bits are reserved, read as zero and writes are ignored. If ATE is zero, bits 3n-1:3n-2 are ignored, and only the bottom bit 3n-3 is used to start and stop tracing as specified in versions less than 4.00 of this specification.	R/W	0						

Table E.40 BreakPoint Control Modes: IBPC and DBP

Value	Trigger Action	Description
000	Unconditional Trace Stop	Unconditionally stop tracing if tracing was turned on. If tracing is already off, then there is no effect.
001	Unconditional Trace Start	Unconditionally start tracing if tracing was turned off. If tracing is already turned off then there is no effect.
010 to 111	Not used	Reserved for future implementation

### E.2.34 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), the *DEPC* contains the virtual address of the instruction where execution should resume after the debug handler code is executed.

## CP0 Registers of the 34K Core

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

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

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

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

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the 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 E-35 DEPC Register Format**



**Table E.41 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

### E.2.35 Performance Counter Register (CP0 Register 25, select 0-3)

The 34K processor defines four performance counters and four 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 E.42](#).

**Table E.42 Performance Counter Register Selects**

Select[2:0]	Register
0	Register 0 Control
1	Register 0 Count
2	Register 1 Control
3	Register 1 Count
4	Register 2 Control
5	Register 2 Count
6	Register 3 Control

Table E.42 Performance Counter Register Selects

Select[2:0]	Register
7	Register 3 Count

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

Bit 31 of each of the counters are AND'ed with an interrupt enable bit, *IE*, of their respective control register, and then OR'ed together to create the *SI\_PCI* output. This signal is combined with one of the *SI\_Int* pins to signal an interrupt to the core. 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.

Figure E-36 Performance Counter Control Register

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

Table E.43 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 a MTC0 or MFC0 select field value of 'n+2' and 'n+3'.	R	1 for counter 0 0 for counter 1										
TCID	29:22	Specifies which TC events should be counted for if per-TC counting is enabled.	R/W	Undefined										
MT_EN	21:20	Specifies which events should be counted: <table border="1" data-bbox="490 1157 1062 1377"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Count events from all TCs &amp; VPEs</td> </tr> <tr> <td>01</td> <td>Count events from all TCs of the VPE specified in VPEID</td> </tr> <tr> <td>10</td> <td>Count events from the TC specified in TCID</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Meaning	00	Count events from all TCs & VPEs	01	Count events from all TCs of the VPE specified in VPEID	10	Count events from the TC specified in TCID	11	Reserved	R/W	Undefined
Encoding	Meaning													
00	Count events from all TCs & VPEs													
01	Count events from all TCs of the VPE specified in VPEID													
10	Count events from the TC specified in TCID													
11	Reserved													
VPEID	19:16	Specifies which VPE events should be counter for if per-VPE counting is enabled.	R/W	Undefined										
Event	11:5	Counter event enabled for this counter. Possible events are listed in <a href="#">Table 8.3 on page 118</a> .	R/W	Undefined										
IE	4	Counter Interrupt Enable. This bit masks bit 31 of the associated count register from the interrupt exception request output.	R/W	0										
U	3	Count in User Mode. When this bit is set, the specified event is counted in User Mode.	R/W	Undefined										
S	2	Count in Supervisor Mode. When this bit is set, the specified event is counted in Supervisor Mode.	R/W	Undefined										
K	1	Count in Kernel Mode. When this bit is set, count the event in Kernel Mode when <i>EXL</i> and <i>ERL</i> both are 0.	R/W	Undefined										

**Table E.43 Performance Counter Control Register Field Descriptions (Continued)**

Fields		Description	Read/ Write	Reset State
Name	Bits			
EXL	0	Count when <i>EXL</i> . When this bit is set, count the event when <i>EXL</i> = 1 and <i>ERL</i> = 0.	R/W	Undefined
0	30, 15:12	Must be written as zeroes; returns zeroes when read.	0	0

The different events which are countable are listed in [Section 8.4, "Performance counters" on page 115](#).

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

**Figure E-37 Performance Counter Count Register****Table E.44 Performance Counter Count Register Field Descriptions**

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

### E.2.36 ErrCtl Register (CP0 Register 26, Select 0)

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

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

A CACHE Index Load Tag operation to the instruction cache will update the *PCI* field with the instruction precode bits from the data array and the *PI* field with the parity bits from the data array if parity is supported. A CACHE Index Load Tag operation to the data cache will cause the *PD* bits to be updated with the byte parity for the selected word of the data array if parity is implemented. If parity is disabled or not implemented, the contents of the *PI* and *PD* fields after a CACHE Index Load Tag operation will be 0.

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

The way- selection RAM test mode is enabled by setting the *WST* bit. This mode is intended for software testing of the way-selection RAM and data RAM. It modifies the functionality of the CACHE Index Load Tag and Index Store Tag operations so that they modify the way-selection RAM instead of the TAG RAMs. In addition, when the *WST* bit is set, the CACHE Index Store Data can be used for testing the data RAM.



Setting the *SPR* bit enables scratchpad test mode. This mode allows reading and writing of the scratchpad pseudo-tags as well the scratchpad data array.

Setting the *ITC* bit enables access to the ITC pseudo-tags that control the addressing information

At most one of the *WST*, *SPR*, and *ITC* bits should be set. Setting multiple bits will lead to unpredictable behavior.

**Figure E-38 ErrCtl Register**

31	30	29	28	27	26	25	24	23			19	18					13	12									4	3				0
PE	PO	WST	SPR	PCO	ITC	LBE	WABE	0				PCI				PI				PD												

**Table E.45 ErrCtl Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
PE	31	Parity Enable. This bit enables or disables the cache parity protection for both the instruction cache and the data cache. <table border="1" data-bbox="485 779 1060 898"> <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> This field is only write-able if the cache parity option was implemented when the core was built. If cache parity is not supported, this field is always read as 0. Software can test for cache parity support by attempting to write a 1 to this field, then read back the value.	Encoding	Meaning	0	Parity disabled	1	Parity enabled	R or R/W	0
Encoding	Meaning									
0	Parity disabled									
1	Parity enabled									
PO	30	Parity Overwrite. If set, the <i>PI/PD</i> fields of this register overwrites calculated parity for the data array. In addition, the <i>P</i> field of the TagLo register overwrites calculated parity for the tag array. This bit only has significance during CACHE Index Store Tag and CACHE Index Store Data operations. <table border="1" data-bbox="485 1220 1060 1339"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Use calculated parity</td> </tr> <tr> <td>1</td> <td>Override calculated parity</td> </tr> </tbody> </table>	Encoding	Meaning	0	Use calculated parity	1	Override calculated parity	R/W	0
Encoding	Meaning									
0	Use calculated parity									
1	Override calculated parity									
WST	29	Way Selection Test. If set, way-selection RAM test mode is enabled. This affects only the CACHE instruction operation. <table border="1" data-bbox="485 1423 1060 1543"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Test mode disabled</td> </tr> <tr> <td>1</td> <td>Test mode enabled</td> </tr> </tbody> </table>	Encoding	Meaning	0	Test mode disabled	1	Test mode enabled	R/W	0
Encoding	Meaning									
0	Test mode disabled									
1	Test mode enabled									
SPR	28	ScratchPadRAM test. If set, indexed CACHE instructions operate on the ScratchPad RAM. Undefined behavior if ScratchPad RAM is not present	R/W	0						

**Table E.45 ErrCtl Register Field Descriptions**

Fields		Description	Read / Write	Reset State						
Name	Bits									
PCO	27	Precode override. If set, the contents of the <i>PCI</i> field overwrite the calculated precode bits when data is written to the instruction cache for indexed CACHE instruction operations. <table border="1" style="margin-top: 10px;"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Use calculated precode</td> </tr> <tr> <td>1</td> <td>Override calculated precode</td> </tr> </tbody> </table>	Encoding	Meaning	0	Use calculated precode	1	Override calculated precode	R/W	0
Encoding	Meaning									
0	Use calculated precode									
1	Override calculated precode									
ITC	26	InterThread Communication. If set, Index Load Tag and Index Store Tag CACHE instructions operate on the ITC tag.	R/W	0						
LBE	25	Bit indicating that the most recent Data Bus Error was involved a load instruction. A Per-TC BE bit will indicate which TCs were impacted.	R/W	Undefined						
WABE	24	Bit indicating that the most recent Data Bus Error was due to a write allocate and that store data was lost. There is no indication of which TC(s) the store request came from. It is possible for both <i>LBE</i> and <i>WABE</i> to be set if the bus error was on a line being used for both loads and stores.	R/W	Undefined						
0	23:19	Must be written as zeroes; returns zeroes when read.	0	0						
PCI	18:13	Instruction precode bits read from or written to the instruction cache data RAM.	R/W	Undefined						
PI	12:4	Parity bit read from or written to instruction cache data RAM. <table border="1" style="margin-top: 10px;"> <thead> <tr> <th>Bits</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>12</td> <td>Even parity bit for the pre-code bits</td> </tr> <tr> <td>11:4</td> <td>Per-byte even parity bits for the 64b of data</td> </tr> </tbody> </table>	Bits	Meaning	12	Even parity bit for the pre-code bits	11:4	Per-byte even parity bits for the 64b of data	R/W	Undefined
Bits	Meaning									
12	Even parity bit for the pre-code bits									
11:4	Per-byte even parity bits for the 64b of data									
PD	3:0	Parity bits read from or written to data cache data RAM. <i>PD[0]</i> is even parity for the least-significant byte of the requested data.	R/W	Undefined						

**E.2.37 CacheErr Register (CP0 Register 27, Select 0)**

The *CacheErr* register provides an interface with the cache error-detection logic. When a Cache Error exception is signalled, the fields of this register are set accordingly.

**Figure E-39 CacheErr Register**




Table E.46 CacheErr Register Field Descriptions

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									
EC	30	<p>Indicates the cache level at which the error was detected:</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Primary</td> </tr> <tr> <td>1</td> <td>Non-primary</td> </tr> </tbody> </table>	Encoding	Meaning	0	Primary	1	Non-primary	R	Undefined
Encoding	Meaning									
0	Primary									
1	Non-primary									
ED	29	<p>Error Data. Indicates a data RAM error.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No data RAM error detected</td> </tr> <tr> <td>1</td> <td>Data RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No data RAM error detected	1	Data RAM error detected	R	Undefined
Encoding	Meaning									
0	No data RAM error detected									
1	Data RAM error detected									
ET	28	<p>Error Tag. Indicates a tag RAM error.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No tag RAM error detected</td> </tr> <tr> <td>1</td> <td>Tag RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No tag RAM error detected	1	Tag RAM error detected	R	Undefined
Encoding	Meaning									
0	No tag RAM error detected									
1	Tag RAM error detected									
ES	27	<p>Error source. Indicates whether error was caused by internal processor or external snoop request.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Error on internal request</td> </tr> <tr> <td>1</td> <td>Error on external request</td> </tr> </tbody> </table>	Encoding	Meaning	0	Error on internal request	1	Error on external request	R	Undefined
Encoding	Meaning									
0	Error on internal request									
1	Error on external request									
EE	26	<p>Error external: Indicates whether a bus parity error was detected. Not supported</p>	R	0						
EB	25	<p>Error Both. Indicates that a data cache error occurred in addition to an instruction cache error.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No additional data cache error</td> </tr> <tr> <td>1</td> <td>Additional data cache error</td> </tr> </tbody> </table> <p>In the case of an additional data cache error, the remainder of the bits in this register are set according to the instruction cache error.</p>	Encoding	Meaning	0	No additional data cache error	1	Additional data cache error	R	Undefined
Encoding	Meaning									
0	No additional data cache error									
1	Additional data cache error									

**Table E.46 CacheErr Register Field Descriptions (Continued)**

Fields		Description	Read / Write	Reset State						
Name	Bits									
EF	24	<p>Error Fatal. Indicates that a fatal cache error has occurred. There are a few situations where software will not be able to get all information about a cache error from the <i>CacheErr</i> register. These situations are fatal because software cannot determine which memory locations have been affected by the error. To enable software to detect these cases, the <i>EF</i> bit (bit 24) has been added to the <i>CacheErr</i> register.</p> <p>The following 6 cases are indicated as fatal cache errors by the <i>EF</i> bit:</p> <p>E.46.1 Dirty parity error in dirty victim (dirty bit cleared in tag)</p> <p>E.46.2 Tag parity error in dirty victim</p> <p>E.46.3 Data parity error in dirty victim</p> <p>E.46.4 WB store miss and EW error at the requested index</p> <p>E.46.5 Dual/Triple errors from different transactions, e.g. scheduled and non-scheduled load.</p> <p>E.46.6 Multiple data cache errors detected before the first instruction of the cache error handler is issued.</p> <p>In addition to the above, simultaneous instruction and data cache errors as indicated by <i>CacheErr[EB]</i> will cause information about the data cache error to be unavailable. However, that situation is not indicated by <i>CacheErr[EF]</i>.</p>	R	Undefined						
SP	23	<p>Scratchpad. Indicates Scratchpad RAM parity error.</p> <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>No Scratchpad RAM error detected</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Scratchpad RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No Scratchpad RAM error detected	1	Scratchpad RAM error detected	R	0
Encoding	Meaning									
0	No Scratchpad RAM error detected									
1	Scratchpad RAM error detected									
EW	22	<p>Error Way. Indicates a way selection RAM error.</p> <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>No way selection RAM error detected</td> </tr> <tr> <td style="text-align: center;">1</td> <td>Way selection RAM error detected</td> </tr> </tbody> </table>	Encoding	Meaning	0	No way selection RAM error detected	1	Way selection RAM error detected	R	Undefined
Encoding	Meaning									
0	No way selection RAM error detected									
1	Way selection RAM error detected									
Way	21:20	<p>Way. Specifies the cache way in which the error was detected. It is not valid if a Tag RAM error is detected (ET=1) or Scratchpad RAM error is detected (SP=1).</p>	R	Undefined						
Index	19:0	<p>Index. Specifies the cache or Scratchpad RAM index of the double word in which the error was detected. The way of the faulty cache is written by hardware in the <i>Way</i> field. Software must combine the <i>Way</i> and <i>Index</i> read in this register with cache configuration information in the <i>Config1</i> register in order to obtain an index which can be used in an indexed CACHE instruction to access the faulty cache data or tag. Note that <i>Index</i> is aligned as a byte index, so it does not need to be shifted by software before it is used in an indexed CACHE instruction. <i>Index</i> bits [4:3] are undefined upon tag RAM errors and <i>Index</i> bits above the MSB actually used for cache indexing will also be undefined.</p> <p>Bits [19:16] are only used for errors in the Scratchpad RAM.</p>	R	Undefined						



Table E.47 TagLo Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State
Name	Bit(s)			
P	0	Parity. Specifies the parity bit for the cache tag. This bit is updated with tag array parity on CACHE Index Load Tag operations and used as tag array parity on Index Store Tag operations when the <i>PO</i> bit of the <i>ErrCtl</i> register is set. NOTE: For the Data cache, this parity does not cover the dirty bit; the dirty bit has a separate parity bit placed in the way selection RAM.	R/W	Undefined
WSDP	23:20	Dirty Parity (Optional, D-side only). This field contains the value read from the WS array during a CACHE Index Load WS operation. If the <i>PO</i> field of the <i>ErrCtl</i> register is asserted, then this field is used to store the dirty parity bits during a CACHE Index Store WS operation.	R/W	Undefined
WSD	19:16	Dirty bits (D-side only). This field contains the value read from the WS array after a CACHE Index Load WS operation. It is used to store into the WS array during CACHE Index Store WS operations.	R/W	Undefined
WSLRU	15:10	LRU bits. This field contains the value read from the WS array after a CACHE Index Load WS operation. It is used to store into the WS array during CACHE Index Store WS operations.	R/W	Undefined
BasePA	31:12	When reading pseudo-tag 0 of a scratchpad RAM, this field will contain bits [31:12] of the base address of the scratchpad region	R/W	Undefined
E	7	When reading pseudo-tag 0 of a scratchpad RAM, this bit will indicate whether the scratchpad is enabled	R/W	Undefined
Size	19:12	When reading pseudo-tag 1 of a scratchpad RAM, this field indicates the size of the scratchpad array. This field is the number of 4KB sections it contains. (Combined with the 0's in 11:0, the register will contain the number of bytes in the scratchpad region)	R/W	Undefined

In addition to the three uses of the *TagLo* register specified above, there is a fourth application where *TagLo* is used to access the pseudo-tags (control registers) of the ITC block. This is done by executing the Index Store Tag or Index Load Tag operation of the CACHE instruction with the *ErrCtl[ITC]* set to 1 (and *ErrCtl[SPR]/ErrCtl[WST]* set to 0).

### E.2.39 DataLo Register (CP0 Register 28, Select 1,3)

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

Note that there are separate *DataLo* registers for each of the primary caches (L1 I-cache: select 1, L1 D-cache: select 3). This register does not exist for the L2 cache.

Figure E-43 DataLo Register Format



Table E.48 DataLo Register Field Description

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

### E.2.40 DataHi Register (CP0 Register 29, Select 1)

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

The *DataHi* register only exists for the Instruction Cache. The interface to the I-cache only operates on pairs of instructions - the high instruction will be written into the *DataHi* register.

Note that *DataHi* and *DataLo* reflect the memory ordering of the instructions. Depending on the endianness of the system, Instruction0 belongs in either *DataHi* (BigEndian) or *DataLo* (LittleEndian) and vice versa for Instruction1.

Figure E-44 DataHi Register Format



Table E.49 DataHi Register Field Description

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

### E.2.41 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 non-maskable 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

## CP0 Registers of the 34K Core

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

In processors that implement the MIPS16 ASE, 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

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

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 E-45 ErrorEPC Register Format**



**Table E.50 ErrorEPC Register Field Description**

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

### E.2.42 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 where the existence of a valid stack for context saving cannot be assumed.

**Figure E-46 DeSave Register Format**



**Table E.51 DeSave Register Field Description**

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



## Revision History

Revision	Date	Description
0.50	21st October 2004	First release for the 34K core “pre-release” package. Status is preliminary; in particular note that the MIPS MT ASE description is missing some details which changed with v0.97 of [MIPSMT].
0.81	24th May 2005	For first customer access (“EA”) release of the 34K core. More information about the DSP ASE. Brought up to date with v0.98 of [MIPSMT] and [MIPSDSP].
0.99	3rd August 2005	Preview of text of v1.00 leading up to GA release of the 34K core.
1.00	9th August 2005	For GA release of the 34K core.
1.05	28th September 2005	For GA release of the 34K core. Better description of policy managers and performance counters. Compatible with v1.00 of MT ASE and DSP ASE
1.20	1st March 2006	Incremental improvements with feedback. Change bars are against 1.05.
1.30	26th May 2006	Changes to help customers recycling the manual for reference: <ul style="list-style-type: none"> <li>• Added CP0 reference-format appendix.</li> <li>• Complete review of performance counter event description.</li> <li>• Many small changes in response to feedback.</li> <li>• Converted to revised document templates.</li> </ul> Change bars vs. 1.20