

The MIPS32™ 4KEc™ core from MIPS® Technologies is a member of the MIPS32 4KE™ processor core family. It is a high-performance, low-power, 32-bit MIPS RISC core designed for custom system-on-silicon applications. The core is designed for semiconductor manufacturing companies, ASIC developers, and system OEMs who want to rapidly integrate their own custom logic and peripherals with a high-performance RISC processor. It is highly portable across processes, and can be easily integrated into full system-on-silicon designs, allowing developers to focus their attention on end-user products. The 4KEc core is ideally positioned to support new products for emerging segments of the digital consumer, network, systems, and information management markets, enabling new tailored solutions for embedded applications.

The 4KEc core implements the MIPS32 Release 2 Architecture with the MIPS16e™ ASE, and the 32-bit privileged resource architecture. The Memory Management Unit (MMU) contains 4-entry instruction and data Translation Lookaside Buffers (ITLB/DTLB) and a 16 dual-entry joint TLB (JTTLB) with variable page sizes. The synthesizable 4KEc core includes a Multiply/Divide Unit (MDU) that implements single cycle MAC instructions, which enable DSP algorithms to be performed efficiently. It allows 32-bit x 16-bit MAC instructions to be issued every cycle, while a 32-bit x 32-bit MAC instruction can be issued every 2 cycles.

Instruction and data caches are fully configurable from 0 - 64 Kbytes in size. In addition, each cache can be organized as direct-mapped or 2-way, 3-way, or 4-way set associative. Load and fetch cache misses only block until the critical word becomes available. The pipeline resumes execution while the remaining words are being written to the cache. Both caches are virtually indexed and physically tagged to allow them to be accessed in the same clock that the address is translated.

An optional Enhanced JTAG (EJTAG) block allows for single-stepping of the processor as well as instruction and data virtual address/value breakpoints. Additionally, real-time tracing of instruction program counter, data address, and data values can be supported.

Figure 1 shows a block diagram of the 4KEc core. The core is divided into *required* and *optional* blocks as shown.




Figure 1 4KEc Core Block Diagram

## Features

- 5-stage pipeline
- 32-bit Address and Data Paths
- MIPS32-Compatible Instruction Set
  - Multiply-Accumulate and Multiply-Subtract Instructions (MADD, MADDU, MSUB, MSUBU)
  - Targeted Multiply Instruction (MUL)
  - Zero/One Detect Instructions (CLZ, CLO)
  - Wait Instruction (WAIT)
  - Conditional Move Instructions (MOVZ, MOVN)
  - Prefetch Instruction (PREF)
- MIPS32 Enhanced Architecture (Release 2) Features
  - Vectored interrupts and support for external interrupt controller
  - Programmable exception vector base
  - Atomic interrupt enable/disable
  - GPR shadow registers (optionally, one or three additional shadows can be added to minimize latency for interrupt handlers)
  - Bit field manipulation instructions
  - Improved virtual memory support (smaller page sizes and hooks for more extensive page table manipulation)
- MIPS16e™ Code Compression
  - 16 bit encodings of 32 bit instructions to improve code density
  - Special PC-relative instructions for efficient loading of addresses and constants
  - SAVE & RESTORE macro instructions for setting up and tearing down stack frames within subroutines
  - Improved support for handling 8 and 16 bit datatypes
- Programmable Cache Sizes
  - Individually configurable instruction and data caches
  - Sizes from 0 - 64KB
  - Direct Mapped, 2-, 3-, or 4-Way Set Associative
  - Loads block only until critical word is available
  - Write-back and write-through support
  - 16-byte cache line size
  - Virtually indexed, physically tagged
  - Cache line locking support
  - Non-blocking prefetches
- Scratchpad RAM Support
  - Can optionally replace 1 way of the I- and/or D-cache with a fast scratchpad RAM
  - Independent external pin interfaces for I- and D-scratchpads
  - 20 index address bits allow access of arrays up to 1MB
  - Interface allows back-stalling the core
- MIPS32 Privileged Resource Architecture
  - Count/Compare registers for real-time timer interrupts
  - I and D watch registers for SW breakpoints
- Programmable Memory Management Unit
  - 16 dual-entry JTLB with variable page size
  - 4-entry ITLB
  - 4-entry DTLB
- Simple Bus Interface Unit (BIU)
  - All I/O's fully registered
  - Separate unidirectional 32-bit address and data buses
  - Two 16-byte collapsing write buffers
  - Designed to allow easy conversion to other bus protocols
- CorExtend™ User Defined Instruction Set Extensions (available in 4KEc Pro™ core)
  - Allows user to define and add instructions to the core at build time
  - Maintains full MIPS32 compatibility
  - Supported by industry standard development tools
  - Single or multi-cycle instructions
  - Separately licensed; a core with this feature is known as the 4KEc Pro™ core
- Multiply/Divide Unit
  - Maximum issue rate of one 32x16 multiply per clock
  - Maximum issue rate of one 32x32 multiply every other clock
  - Early-in iterative divide. Minimum 11 and maximum 34 clock latency (dividend (*rs*) sign extension-dependent)
- Coprocessor 2 interface
  - 32 bit interface to an external coprocessor
- Power Control
  - Minimum frequency: 0 MHz
  - Power-down mode (triggered by WAIT instruction)
  - Support for software-controlled clock divider
  - Support for extensive use of local gated clocks
- EJTAG Debug
  - Support for single stepping
  - Virtual instruction and data address/value breakpoints
  - PC and data tracing
  - TAP controller is chainable for multi-CPU debug
  - Cross-CPU breakpoint support
- Testability
  - Full scan design achieves test coverage in excess of 99% (dependent on library and configuration options)
  - Optional memory BIST for internal SRAM arrays

## Architecture Overview

The 4KEc core contains both required and optional blocks. Required blocks are the lightly shaded areas of the block diagram in [Figure 1](#) and must be implemented to remain

MIPS-compliant. Optional blocks can be added to the 4KEc core based on the needs of the implementation.

The required blocks are as follows:

- Execution Unit
- Multiply/Divide Unit (MDU)
- System Control Coprocessor (CP0)
- Memory Management Unit (MMU)
- Transition Lookaside Buffer (TLB)
- Cache Controllers
- Bus Interface Unit (BIU)
- Power Management

Optional blocks include:

- Instruction Cache
- Data Cache
- Scratchpad RAM interface
- Coprocessor 2 interface
- CorExtend™ User Defined Instruction (UDI) support
- MIPS16e support
- Enhanced JTAG (EJTAG) Controller

The section entitled "4KEc Core Required Logic Blocks" on page 4 discusses the required blocks. The section entitled "4KEc Core Optional Logic Blocks" on page 15 discusses the optional blocks.

## Pipeline Flow

The 4KEc core implements a 5-stage pipeline with performance similar to the R3000® pipeline. The pipeline allows the processor to achieve high frequency while minimizing device complexity, reducing both cost and power consumption.

The 4KEc core pipeline consists of five stages:

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

The 4KEc core implements a bypass mechanism that allows the result of an operation to be forwarded directly to

the instruction that needs it without having to write the result to the register and then read it back.

Figure 2 shows a timing diagram of the 4KEc core pipeline.

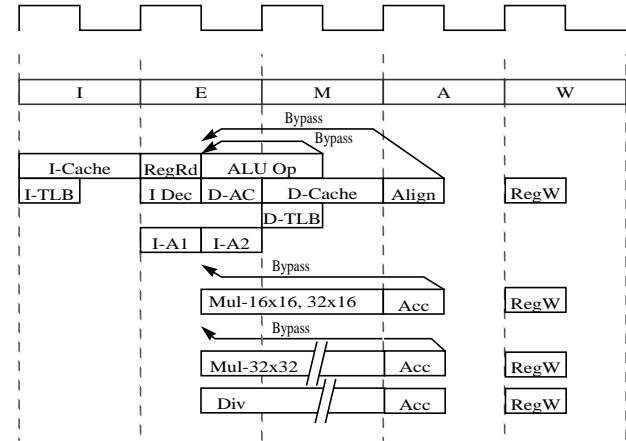


Figure 2 4KEc Core Pipeline

### I Stage: Instruction Fetch

During the Instruction fetch stage:

- An instruction is fetched from instruction cache.
- MIPS16e instructions are expanded into MIPS32-like instructions

### E Stage: Execution

During the Execution stage:

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

### M Stage: Memory Fetch

During the Memory fetch stage:

- The arithmetic ALU operation completes.
- The data cache access and the data virtual-to-physical address translation are performed for load and store instructions.

- Data cache look-up is performed and a hit/miss determination is made.
- A 16x16 or 32x16 multiply calculation completes.
- A 32x32 multiply operation stalls the MDU pipeline for one clock in the M stage.
- A divide operation stalls the MDU pipeline for a maximum of 34 clocks in the M stage. Early-in sign extension detection on the dividend will skip 7, 15, or 23 stall clocks.

#### A Stage: Align

During the Align stage:

- Load data is aligned to its word boundary.
- A 16x16 or 32x16 multiply operation performs the carry-propagate-add. The actual register writeback is performed in the W stage.
- A MUL operation makes the result available for writeback. The actual register writeback is performed in the W stage.

#### W Stage: Writeback

During the Writeback stage:

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

## 4KEc Core Required Logic Blocks

The 4KEc core consists of the following required logic blocks, shown in [Figure 1](#). These logic blocks are defined in the following subsections:

- Execution Unit
- Multiply/Divide Unit (MDU)
- System Control Coprocessor (CP0)
- Memory Management Unit (MMU)
- Transition Lookaside Buffer (TLB)
- Cache Controller
- Bus Interface Unit (BIU)
- Power Management

### Execution Unit

The 4KEc core execution unit implements a load/store architecture with single-cycle ALU operations (logical,

shift, add, subtract) and an autonomous multiply/divide unit. The 4KEc core contains thirty-two 32-bit general-purpose registers used for integer operations and address calculation. Optionally, one or three additional register file shadow sets (each containing thirty-two registers) can be added to minimize context switching overhead during interrupt/exception processing. The register file consists of two read ports and one write port and is fully bypassed to minimize operation latency in the pipeline.

The execution unit includes:

- 32-bit adder used for calculating the data address
- Address unit for calculating the next instruction address
- Logic for branch determination and branch target address calculation
- Load aligner
- Bypass multiplexers used to avoid stalls when executing instructions streams where data producing instructions are followed closely by consumers of their results
- Leading Zero/One detect unit for implementing the CLZ and CLO instructions
- Arithmetic Logic Unit (ALU) for performing bitwise logical operations
- Shifter & Store Aligner

### Multiply/Divide Unit (MDU)

The 4KEc core includes a multiply/divide unit (MDU) that contains a separate pipeline for multiply and divide operations. This pipeline operates in parallel with the integer unit (IU) pipeline and does not stall when the IU pipeline stalls. This setup allows long-running MDU operations, such as a divide, to be partially masked by system stalls and/or other integer unit instructions.

The MDU consists of a 32x16 booth recoded multiplier, result/accumulation registers (HI and LO), a divide state machine, and the necessary multiplexers and control logic. The first number shown ('32' of 32x16) represents the *rs* operand. The second number ('16' of 32x16) represents the *rt* operand. The 4KEc core only checks the value of the latter (*rt*) operand to determine how many times the operation must pass through the multiplier. The 16x16 and 32x16 operations pass through the multiplier once. A 32x32 operation passes through the multiplier twice.

The MDU supports execution of one 16x16 or 32x16 multiply operation every clock cycle; 32x32 multiply

operations can be issued every other clock cycle. Appropriate interlocks are implemented to stall the issuance of back-to-back 32x32 multiply operations. The multiply operand size is automatically determined by logic built into the MDU.

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

Table 1 lists the repeat rate (peak issue rate of cycles until the operation can be reissued) and latency (number of cycles until a result is available) for the 4KEc core multiply and divide instructions. The approximate latency and repeat rates are listed in terms of pipeline clocks. For a more detailed discussion of latencies and repeat rates, refer to Chapter 2 of the *MIPS32 4KE™ Processor Core Family Software User's Manual*.

Table 1 4KEc Core High-Performance Integer Multiply/Divide Unit Latencies and Repeat Rates

Opcode	Operand Size (mul <i>rt</i> ) (div <i>rs</i> )	Latency	Repeat Rate
MULT/MULTU, MADD/MADDU, MSUB/MSUBU	16 bits	1	1
	32 bits	2	2
MUL	16 bits	2	1
	32 bits	3	2
DIV/DIVU	8 bits	12	11
	16 bits	19	18
	24 bits	26	25
	32 bits	33	32

The MIPS architecture defines that the result of a multiply or divide operation be placed in the HI and LO registers. Using the Move-From-HI (MFHI) and Move-From-LO (MFLO) instructions, these values can be transferred to the general-purpose register file.

In addition to the HI/LO targeted operations, the MIPS32 architecture also defines a multiply instruction, MUL, which places the least significant results in the primary

register file instead of the HI/LO register pair. By avoiding the explicit MFLO instruction, required when using the LO register, and by supporting multiple destination registers, the throughput of multiply-intensive operations is increased.

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

### System Control Coprocessor (CP0)

In the MIPS architecture, CP0 is responsible for the virtual-to-physical address translation and cache protocols, the exception control system, the processor's diagnostics capability, the operating modes (kernel, user, and debug), and whether interrupts are enabled or disabled. Configuration information, such as cache size and set associativity, is also available by accessing the CP0 registers, listed in Table 2.

Table 2 Coprocessor 0 Registers in Numerical Order

Register Number	Register Name	Function
0	Index <sup>3</sup>	Index into the TLB array.
1	Random <sup>3</sup>	Randomly generated index into the TLB array.
2	EntryLo0 <sup>3</sup>	Low-order portion of the TLB entry for odd-numbered virtual pages.
3	EntryLo1 <sup>3</sup>	Low-order portion of the TLB entry for odd-numbered virtual pages.
4	Context <sup>1</sup>	Pointer to page table entry in memory.
4	Context-Config <sup>1</sup>	Controls the layout of the Context register.
5	PageMask <sup>3</sup>	Control for variable page sizes in TLB entries.
5	PageGrain <sup>3</sup>	Controls the layout of the EntryLo, PageMask and EntryHi registers.

Table 2 Coprocessor 0 Registers in Numerical Order

Register Number	Register Name	Function
6	Wired <sup>3</sup>	Controls the number of fixed (“wired”) TLB entries.
7	HWREna	Enables access via the RDHWR instruction to selected hardware registers.
8	BadVAddr <sup>1</sup>	Reports the address for the most recent address-related exception.
9	Count <sup>1</sup>	Processor cycle count.
10	EntryHi <sup>3</sup>	High-order portion of the TLB entry.
11	Compare <sup>1</sup>	Timer interrupt control.
12	Status <sup>1</sup>	Processor status and control.
12	IntCtl <sup>1</sup>	Interrupt system status and control.
12	SRSCtl <sup>1</sup>	Shadow register set status and control.
12	SRSMap <sup>1</sup>	Provides mapping from vectored interrupt to a shadow set.
13	Cause <sup>1</sup>	Cause of last general exception.
14	EPC <sup>1</sup>	Program counter at last exception.
15	PRId	Processor identification and revision.
15	EBASE	Exception vector base register.
16	Config	Configuration register.
16	Config1	Configuration register 1.
16	Config2	Configuration register 2.
16	Config3	Configuration register 3.
17	LLAddr	Load linked address.
18	WatchLo <sup>1</sup>	Low-order watchpoint address.
19	WatchHi <sup>1</sup>	High-order watchpoint address.
20-22	Reserved	Reserved in the 4KEc core.
23	Debug <sup>2</sup>	Debug control and exception status.
23	Trace Control <sup>2</sup>	PC/Data trace control register.
23	Trace Control2 <sup>2</sup>	Additional PC/Data trace control.

Table 2 Coprocessor 0 Registers in Numerical Order

Register Number	Register Name	Function
23	User Trace Data <sup>2</sup>	User Trace control register.
23	TraceBPC <sup>2</sup>	Trace breakpoint control.
24	DEPC <sup>2</sup>	Program counter at last debug exception.
25	Reserved	Reserved in the 4KEc core.
26	ErrCtl	Used for software testing of cache arrays.
27	Reserved	Reserved in the 4KEc core.
28	TagLo/ DataLo	Low-order portion of cache tag interface.
29	Reserved	Reserved in the 4KEc core.
30	ErrorEPC <sup>1</sup>	Program counter at last error.
31	DESAVE <sup>2</sup>	Debug handler scratchpad register.
1. Registers used in exception processing. 2. Registers used during debug. 3. Registers used in memory management.		

Coprocessor 0 also contains the logic for identifying and managing exceptions. Exceptions can be caused by a variety of sources, including boundary cases in data, external events, or program errors. Table 3 shows the exception types in order of priority.

Table 3 4KEc Core Exception Types

Exception	Description
Reset	Assertion of <i>SI_ColdReset</i> or <i>SI_Reset</i> signals.
DSS	EJTAG Debug Single Step.
DINT	EJTAG Debug Interrupt. Caused by the assertion of the external <i>EJ_DINT</i> input, or by setting the <i>EjtagBrk</i> bit in the ECR register.
NMI	Assertion of <i>EB_NMI</i> signal.
Machine Check	TLB write that conflicts with an existing entry.
Interrupt	Assertion of unmasked hardware or software interrupt signal.

Table 3 4KEc Core Exception Types (Continued)

Exception	Description
Deferred Watch	Deferred Watch (unmasked by K DM->!(K DM) transition).
DIB	EJTAG debug hardware instruction break matched.
WATCH	A reference to an address in one of the watch registers (fetch).
AdEL	Fetch address alignment error. Fetch reference to protected address.
TLBL	Fetch TLB miss.
TLBL	Fetch TLB hit to page with V=0.
IBE	Instruction fetch bus error.
DBp	EJTAG Breakpoint (execution of SDBBP instruction).
Sys	Execution of SYSCALL instruction.
Bp	Execution of BREAK instruction.
RI	Execution of a Reserved Instruction.
CpU	Execution of a coprocessor instruction for a coprocessor that is not enabled.
Ov	Execution of an arithmetic instruction that overflowed.
Tr	Execution of a trap (when trap condition is true).
DDBL / DDBS	EJTAG Data Address Break (address only) or EJTAG Data Value Break on Store (address+value).
WATCH	A reference to an address in one of the watch registers (data).
AdEL	Load address alignment error. Load reference to protected address.
AdES	Store address alignment error. Store to protected address.
TLBL	Load TLB miss.
TLBL	Load TLB hit to page with V=0.
TLBS	Store TLB miss.
TLBS	Store TLB hit to page with V=0.
TLB Mod	Store to TLB page with D=0.
DBE	Load or store bus error.

Table 3 4KEc Core Exception Types (Continued)

Exception	Description
DDBL	EJTAG data hardware breakpoint matched in load data compare.

### Interrupt Handling

The 4KEc core includes support for six hardware interrupt pins, two software interrupts, and a timer interrupt. These interrupts can be used in any of three interrupt modes, as defined by Release 2 of the MIPS32 Architecture:

- Interrupt compatibility mode, which acts identically to that in an implementation of Release 1 of the Architecture.
- Vectored Interrupt (VI) mode, which adds the ability to prioritize and vector interrupts to a handler dedicated to that interrupt, and to assign a GPR shadow set for use during interrupt processing. The presence of this mode is denoted by the VInt bit in the *Config3* register. This mode is architecturally optional; but it is always present on the 4KEc core, so the VInt bit will always read as a 1 for the 4KEc core.
- External Interrupt Controller (EIC) mode, which redefines the way in which interrupts are handled to provide full support for an external interrupt controller handling prioritization and vectoring of interrupts. This presence of this mode denoted by the VEIC bit in the *Config3* register. Again, this mode is architecturally optional. On the 4KEc core, the VEIC bit is set externally by the static input, *SI\_EICPresent*, to allow system logic to indicate the presence of an external interrupt controller.

The reset state of the processor is to interrupt compatibility mode such that a processor supporting Release 2 of the Architecture, like the 4KEc core, is fully compatible with implementations of Release 1 of the Architecture.

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

### GPR Shadow Registers

Release 2 of the MIPS32 Architecture optionally removes the need to save and restore GPRs on entry to high priority interrupts or exceptions, and to provide specified processor

modes with the same capability. This is done by introducing multiple copies of the GPRs, called *shadow sets*, and allowing privileged software to associate a shadow set with entry to kernel mode via an interrupt vector or exception. The normal GPRs are logically considered shadow set zero.

The number of GPR shadow sets is a build-time option on the 4KEc core. Although Release 2 of the Architecture defines a maximum of 16 shadow sets, the core allows one (the normal GPRs), two, or four shadow sets. The highest number actually implemented is indicated by the `SRSCtlHSS` field. If this field is zero, only the normal GPRs are implemented.


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

If the processor is operating in VI interrupt mode, binding of a vectored interrupt to a shadow set is done by writing to the `SRSMap` register. If the processor is operating in EIC interrupt mode, the binding of the interrupt to a specific shadow set is provided by the external interrupt controller, and is configured in an implementation-dependent way. Binding of an exception or non-vectored interrupt to a shadow set is done by writing to the `ESS` field of the `SRSCtl` register. When an exception or interrupt occurs, the value of `SRSCtlCSS` is copied to `SRSCtlPSS`, and `SRSCtlCSS` is set to the value taken from the appropriate source. On an `ERET`, the value of `SRSCtlPSS` is copied back into `SRSCtlCSS` to restore the shadow set of the mode to which control returns.

## Modes of Operation

The 4KEc core supports three modes of operation: user mode, kernel mode, and debug mode. User mode is most often used for applications programs. Kernel mode is typically used for handling exceptions and operating system kernel functions, including CP0 management and I/O device accesses. An additional Debug mode is used

during system bring-up and software development. Refer to the EJTAG section for more information on debug mode.



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

Figure 3 4KEc Core Virtual Address Map

## Memory Management Unit (MMU)

The 4KEc core contains a fully functional MMU that interfaces between the execution unit and the cache controller. Although the 4KEc core implements a 32-bit architecture, the MMU is modeled after that found in the 64-bit R4000 family, as defined by the MIPS32 Privileged Resource Architecture (PRA).

The 4KEc core implements a TLB-based MMU. The TLB consists of three translation buffers: a 16 dual-entry fully associative Joint TLB (JTLB), a 4-entry fully associative Instruction TLB (ITLB) and a 4-entry fully associative data TLB (DTLB).

When an instruction address is calculated, the virtual address is compared to the contents of the 4-entry ITLB. If



the address is not found in the ITLB, the JTLB is accessed. If the entry is found in the JTLB, that entry is then written into the ITLB. If the address is not found in the JTLB, a TLB refill exception is taken.

When a data address is calculated, the virtual address is compared to both the 4-entry DTLB and the JTLB. If the address is not found in the DTLB, but is found in the JTLB, that address is immediately written to the DTLB. If the address is not found in the JTLB, a TLB refill exception is taken.

Figure 4 shows how the ITLB, DTLB, and JTLB are implemented in the 4KEc core.

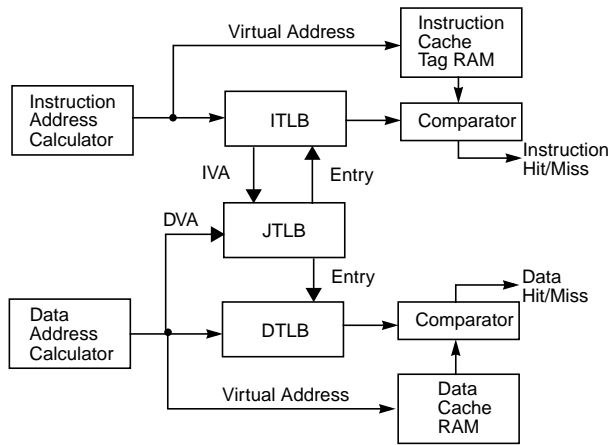


Figure 4 Address Translation During a Cache Access

### Translation Lookaside Buffer (TLB)

The TLB consists of three address translation buffers:

- 16 dual-entry fully associative Joint TLB (JTLB)
- 4-entry fully associative Instruction TLB (ITLB)
- 4-entry fully associative Data TLB (DTLB)

### Joint TLB (JTLB)

The 4KEc core implements a 16-dual-entry, fully associative JTLB that maps 32 virtual pages to their corresponding physical addresses. The purpose of the TLB is to translate virtual addresses and their corresponding ASIDs into a physical memory address. The translation is performed by comparing the upper bits of the virtual address (along with the ASID) against each of the entries in the *tag* portion of the joint TLB structure.

The JTLB is organized as 16 pairs of even and odd entries containing pages that range in size from 4-Kbytes (or 1-

Kbyte) to 256-Mbytes into the 4-Gbyte physical address space. By default, the minimum page size is normally 4-Kbytes on the 4KEc core; as a build time option, it is possible to specify a minimum page size of 1-Kbyte.

The JTLB is organized in page pairs to minimize the overall size. Each *tag* entry corresponds to 2 data entries: an even page entry and an odd page entry. The highest order virtual address bit not participating in the tag comparison is used to determine which of the data entries is used. Since page size can vary on a page-pair basis, the determination of which address bits participate in the comparison and which bit is used to make the even-odd determination is decided dynamically during the TLB look-up.

### Instruction TLB (ITLB)

The ITLB is a small 4-entry, fully associative TLB dedicated to performing translations for the instruction stream. The ITLB only maps minimum sized pages/subpages. The minimum page size is either 1-Kbyte or 4-Kbyte, depending on the *PageGrain* and *Config3* registers.

The ITLB is managed by hardware and is transparent to software. The larger JTLB is used as a backing store for the ITLB. If a fetch address cannot be translated by the ITLB, the JTLB is used to attempt to translate it in the following clock cycle. If successful, the translation information is copied into the ITLB for future use. There is a two cycle ITLB miss penalty.

### Data TLB (DTLB)

The DTLB is a small 4-entry, fully associative TLB dedicated to performing translations for loads and stores. Similar to the ITLB, the DTLB only maps either 1-Kbyte or 4-Kbyte pages/subpages depending on the *PageGrain* and *Config3* registers.

The DTLB is managed by hardware and is transparent to software. The larger JTLB is used as a backing store for the DTLB. The JTLB is looked up in parallel with the DTLB to minimize the DTLB miss penalty. If the JTLB translation is successful, the translation information is copied into the DTLB for future use. There is a one cycle DTLB miss penalty.

### Virtual-to-Physical Address Translation

Converting a virtual address to a physical address begins by comparing the virtual address from the processor with the virtual addresses in the TLB; there is a match when the

virtual page number (VPN) of the address is the same as the VPN field of the entry, and either:

- The Global (*G*) bit of the TLB entry is set, or
- The ASID field of the virtual address is the same as the ASID field of the TLB entry.

This match is referred to as a TLB *hit*. If there is no match, a TLB *miss* exception is taken by the processor and software is allowed to refill the TLB from a page table of virtual/physical addresses in memory.

Figure 5 shows a flow diagram of the address translation process for two different page sizes.

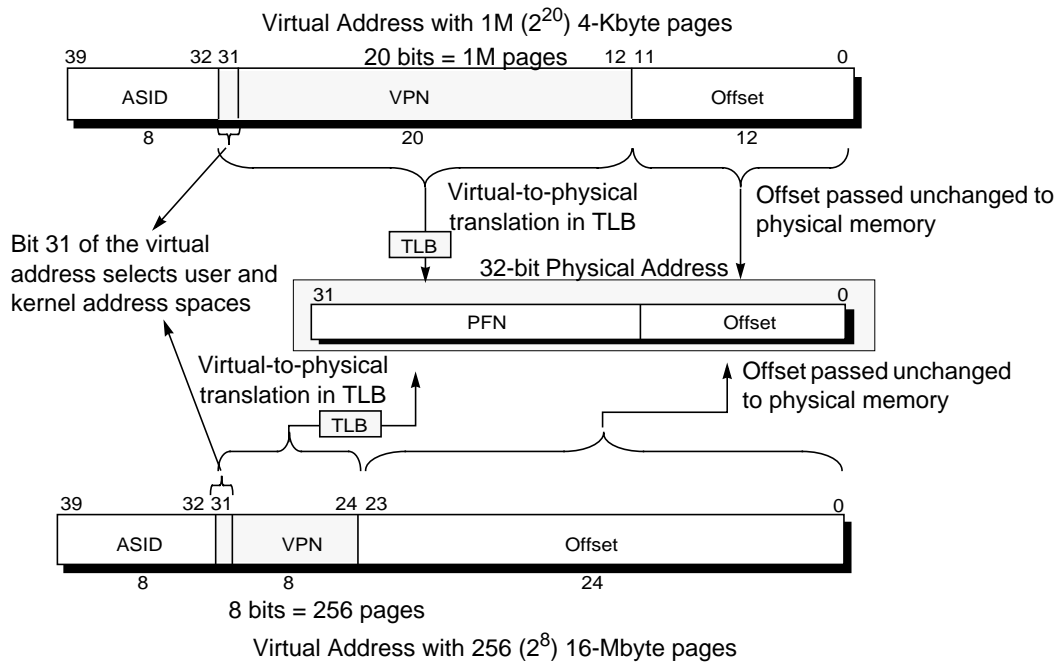


Figure 5 32-bit Virtual Address Translation

The top portion of Figure 5 shows a virtual address for a 4-Kbyte page size. The width of the *Offset* in Figure 5 is defined by the page size. The remaining 20 bits of the address represent the virtual page number (VPN), and index the 1M-entry page table.

The bottom portion of Figure 5 shows the virtual address for a 16-Mbyte page size. The remaining 8 bits of the address represent the VPN, and index the 256-entry page table.

In this figure, the virtual address is extended with an 8-bit address space identifier (ASID), which reduces the frequency of TLB flushing during a context switch. This 8-bit ASID contains the number assigned to that process and is stored in the CP0 *EntryHi* register.

### Hits, Misses, and Multiple Matches

Each JTLB entry contains a tag portion and a data portion. If a match is found, the upper bits of the virtual address are

replaced with the page frame number (PFN) stored in the corresponding entry in the data array of the joint TLB (JTLB). The granularity of JTLB mappings is defined in terms of TLB *pages*. The 4KEc core's JTLB supports pages of different sizes ranging from 1 KB to 256 MB in powers of 4.

If no match occurs (TLB miss), an exception is taken and software refills the TLB from the page table resident in memory. Software can write over a selected TLB entry or use a hardware mechanism to write into a random entry.

The 4KEc core implements a TLB write compare mechanism to ensure that multiple TLB matches do not occur. On the TLB write operation, the write value is compared with all other entries in the TLB. If a match occurs, the 4KEc core takes a machine check exception, sets the TS bit in the CP0 *Status* register, and aborts the write operation.

Table 4 shows the address bits used for even/odd bank selection depending on page size and the relationship between the legal values in the mask register and the selected page size.

Table 4 Mask and Page Size Values

Pagemask[28:11]	Page Size	Even/Odd Bank Select Bit
00000000000000000000	1KB (if present)	VAddr[10]
00000000000000000001	4KB	VAddr[12]
00000000000000000111	16KB	VAddr[14]
000000000000011111	64KB	VAddr[16]
000000000001111111	256KB	VAddr[18]
000000000111111111	1MB	VAddr[20]
000000111111111111	4MB	VAddr[22]
000011111111111111	16MB	VAddr[24]
001111111111111111	64MB	VAddr[26]
111111111111111111	256MB	VAddr[28]

### TLB Tag and Data Formats

Figure 6 shows the format of a TLB *tag* entry. The entry is divided into the follow fields:

- Global process indicator
- Address space identifier
- Virtual page number
- Compressed page mask

Setting the global process indicator (G bit) indicates that the entry is global to all processes and/or threads in the system. In this case, the 8-bit address space identifier (ASID) value is ignored since the entry is not relative to a specific thread or process.

The ASID helps to reduce the frequency of TLB flushing on a context switch. The existence of the ASID allows multiple processes to exist in both the TLB and instruction caches. The current ASID value is stored in the *EntryHi* register and is compared to the ASID value of each entry. Figure 6 and Table 5 show the TLB tag entry format.




Figure 6 TLB Tag Entry Format

Table 5 TLB Tag Entry Fields

Field Name	Description
G	Global Bit. When set, indicates that this entry is global to all processes and/or threads and thus disables inclusion of the ASID in the comparison.
ASID[7:0]	Address Space Identifier. Identifies with which process or thread this TLB entry is associated.
VPN2[31:25], VPN2[24:11]	Virtual Page Number divided by 2. This field contains the upper bits of the virtual page number. Because it represents a pair of TLB pages, it is divided by 2. Bits 31:25 are always included in the TLB lookup comparison. Bits 24:11 are included depending on the page size.
CMASK[7:0]	Compressed page mask value. This field is a compressed version of the page mask. It defines the page size by masking the appropriate VPN2 bits from being involved comparison. It is also used to determine which address bit is used to make the even-odd page determination.

Figure 7 and Table 6 show the TLB data array entry format.




Figure 7 TLB Data Array Entry Format

Table 6 TLB Data Array Entry Fields

Field Name	Description																		
PFN([31:12] or [29:10])	<p>Physical Frame Number. Defines the upper bits of the physical address.</p> <p>The [29:10] range illustrates, that if 1Kbytes page granularity is enabled in the PageGrain register, the PFN is shifted to the right, before being appended to the untranslated part of the virtual address. In this mode the upper two physical address bits are not covered by PFN but forced to zero.</p> <p>For page sizes larger than the minimum configured page size, only a subset of these bits is actually used.</p>																		
C[2:0]	<p>Cacheability. Contains an encoded value of the cacheability attributes and determines whether the page should be placed in the cache or not. The field is encoded as follows:</p> <table border="1" style="width: 100%;"> <thead> <tr> <th>CS[2:0]</th> <th>Coherency Attribute</th> </tr> </thead> <tbody> <tr> <td>000*</td> <td>Cacheable, noncoherent, write through, no write allocate.</td> </tr> <tr> <td>001*</td> <td>Cacheable, noncoherent, write through, write allocate</td> </tr> <tr> <td>010</td> <td>Uncached</td> </tr> <tr> <td>011</td> <td>Cacheable, noncoherent, write back, write allocate</td> </tr> <tr> <td>100*</td> <td>Maps to entry 011b.</td> </tr> <tr> <td>101*</td> <td>Maps to entry 011b.</td> </tr> <tr> <td>110*</td> <td>Maps to entry 011b.</td> </tr> <tr> <td>111*</td> <td>Maps to entry 010b.</td> </tr> </tbody> </table> <p>*Values 2 and 3 are the required MIPS32 mappings for uncached and cacheable references; other values may have different meanings in other MIPS32 processors.</p>	CS[2:0]	Coherency Attribute	000*	Cacheable, noncoherent, write through, no write allocate.	001*	Cacheable, noncoherent, write through, write allocate	010	Uncached	011	Cacheable, noncoherent, write back, write allocate	100*	Maps to entry 011b.	101*	Maps to entry 011b.	110*	Maps to entry 011b.	111*	Maps to entry 010b.
CS[2:0]	Coherency Attribute																		
000*	Cacheable, noncoherent, write through, no write allocate.																		
001*	Cacheable, noncoherent, write through, write allocate																		
010	Uncached																		
011	Cacheable, noncoherent, write back, write allocate																		
100*	Maps to entry 011b.																		
101*	Maps to entry 011b.																		
110*	Maps to entry 011b.																		
111*	Maps to entry 010b.																		
D	<p>“Dirty” or write-enable bit. Indicates that the page has been written and/or is writable. If this bit is set, stores to the page are permitted. If the bit is cleared, stores to the page cause a TLB Modified exception.</p>																		

Table 6 TLB Data Array Entry Fields (Continued)

Field Name	Description
V	<p>Valid bit. Indicates that the TLB entry, and thus the virtual page mapping, are valid. If this bit is set, accesses to the page are permitted. If the bit is cleared, accesses to the page cause a TLB Invalid exception.</p>

### Page Sizes and Replacement Algorithm

To assist in controlling both the amount of mapped space and the replacement characteristics of various memory regions, the 4KEc core provides two mechanisms. First, the page size can be configured, on a per-entry basis, to map a page size of 1Kbyte to 256Mbytes (in multiples of 4). The CP0 *PageMask* register is loaded with the mapping page size, which is then entered into the TLB when a new entry is written. Thus, operating systems can provide special purpose maps. For example, a typical frame buffer can be memory mapped with only one TLB entry.

The second mechanism controls the replacement algorithm when a TLB miss occurs. To select a TLB entry to be written with a new mapping, the 4KEc core provides a random replacement algorithm. However, the processor also provides a mechanism where a programmable number of mappings can be locked into the TLB via the CP0 *Wired* register, thus avoiding random replacement.

### Cache Controllers

The 4KEc core instruction and data cache controllers support caches of various sizes, organizations, and set-associativity. For example, the data cache can be 2 Kbytes in size and 2-way set associative, while the instruction cache can be 8 Kbytes in size and 4-way set associative. Each cache can each be accessed in a single processor cycle. In addition, each cache has its own 32-bit data path and both caches can be accessed in the same pipeline clock cycle. Refer to the section entitled "4KEc Core Optional Logic Blocks" on page 15 for more information on instruction and data cache organization.

The cache controllers also have built-in support for replacing one way of the cache with a scratchpad RAM. See the section entitled "Scratchpad RAM" on page 16 for more information on scratchpad RAMs.

## Bus Interface (BIU)

The Bus Interface Unit (BIU) controls the external interface signals. Additionally, it contains the implementation of the 32-byte collapsing write buffer. The purpose of this buffer is to store and combine write transactions before issuing them at the external interface. When using the write-through cache policy, the write buffer significantly reduces the number of write transactions on the external interface and reduces the amount of stalling in the core due to issuance of multiple writes in a short period of time. When using a write-back cache policy, the write buffer gathers the 4 words of dirty line writebacks.

The write buffer is organized as two 16-byte buffers. Each buffer contains data from a single 16-byte aligned block of memory. One buffer contains the data currently being transferred on the external interface, while the other buffer contains accumulating data from the core. Data from the accumulation buffer is transferred to the external interface buffer under one of these conditions:

- When a store is attempted from the core to a different 16-byte block than is currently being accumulated
- SYNC Instruction
- Store to an invalid merge pattern
- Any load or store to uncached memory
- A load to the line being merged
- A complete 16B block has been gathered

Note that if the data in the external interface buffer has not been written out to memory, the core is stalled until the memory write completes. After completion of the memory write, accumulated buffer data can be written to the external interface buffer.

## Merge Control

The 4KEc core implements two 16-byte collapsing write buffers that allow byte, halfword, or word writes from the core to be accumulated in the buffer into a 16-byte value before bursting the data onto the bus in word format. Note that writes to uncached areas are never merged.

The 4KEc core provides two options for merge pattern control:

- No merge
- Full merge

In *No Merge* mode, writes to a different word within the same line are accumulated in the buffer. Writes to the same word cause the previous word to be driven onto the bus.

In *Full Merge* mode, all combinations of writes to the same line are collected in the buffer. Any pattern of byte enables is possible.

## SimpleBE Mode

To aid in attaching the 4KEc core to structures which cannot easily handle arbitrary byte enable patterns, there is a mode that generates only “simple” byte enables. Only byte enables representing naturally aligned byte, half, and word transactions will be generated. Legal byte enable patterns are shown in Table 7.

Table 7 Valid SimpleBE Byte Enable Patterns

EB_BE[3:0]
0001
0010
0100
1000
0011
1100
1111

The only case where a read can generate “non-simple” byte enables is on an uncached tri-byte load (LWL/LWR). In SimpleBE mode, such reads will be converted into a word read on the external interface.

Writes with non-simple byte enable patterns can arise when a sequence of stores is processed by the merging write buffer, or from uncached tri-byte stores (SWL/SWR). In SimpleBE mode, these stores will be broken into two separate write transactions, one with a valid halfword and a second with a single valid byte. This splitting is independent of the merge pattern control in the write buffer.

## Hardware Reset

For historical reasons within the MIPS architecture, the 4KEc core has two types of reset input signals: *SI\_Reset* and *SI\_ColdReset*.

Functionally, these two signals are ORed together within the core and then used to initialize critical hardware state.

Both reset signals can be asserted either synchronously or asynchronously to the core clock, *SI\_ClkIn*, and will trigger a Reset exception. The reset signals are active high, and must be asserted for a minimum of 5 *SI\_ClkIn* cycles. The falling edge triggers the Reset exception. The primary difference between the two reset signals is that *SI\_Reset* sets a bit in the Status register; this bit could be used by software to distinguish between the two reset signals, if desired. The reset behavior is summarized in Table 8.

Table 8 4KEc Reset Types

SI_Reset	SI_ColdReset	Action
0	0	Normal Operation, no reset.
1	0	Reset exception; sets <i>Status.SR</i> bit.
X	1	Reset exception.

One (or both) of the reset signals must be asserted at power-on or whenever hardware initialization of the core is desired. A power-on reset typically occurs when the machine is first turned on. A hard reset usually occurs when the machine is already on and the system is rebooted.

In debug mode, EJTAG can request that a soft reset (via the *SI\_Reset* pin) be masked. It is system dependent whether this functionality is supported. In normal mode, the *SI\_Reset* pin cannot be masked. The *SI\_ColdReset* pin is never masked.

## Power Management

The 4KEc core offers a number of power management features, including low-power design, active power management, and power-down modes of operation. The core is a static design that supports slowing or halting the clocks, which reduces system power consumption during idle periods.

The 4KEc core provides two mechanisms for system-level low power support:

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

### Register-Controlled Power Management

The RP bit in the CP0 Status register provides a software mechanism for placing the system into a low power state. The state of the RP bit is available externally via the *SI\_RP*

signal. The external agent then decides whether to place the device in a low power mode, such as reducing the system clock frequency.

Three additional bits, *Status\_EXL*, *Status\_ERL*, and *Debug\_DM* support the power management function by allowing the user to change the power state if an exception or error occurs while the 4KEc core is in a low power state. Depending on what type of exception is taken, one of these three bits will be asserted and reflected on the *SI\_EXL*, *SI\_ERL*, or *EJ\_DebugM* outputs. The external agent can look at these signals and determine whether to leave the low power state to service the exception.

The following 4 power-down signals are part of the system interface and change state as the corresponding bits in the CP0 registers are set or cleared:

- The *SI\_RP* signal represents the state of the RP bit (27) in the CP0 Status register.
- The *SI\_EXL* signal represents the state of the EXL bit (1) in the CP0 Status register.
- The *SI\_ERL* signal represents the state of the ERL bit (2) in the CP0 Status register.
- The *EJ\_DebugM* signal represents the state of the DM bit (30) in the CP0 Debug register.

### Instruction-Controlled Power Management

The second mechanism for invoking power-down mode is through execution of the WAIT instruction. When the WAIT instruction is executed, the internal clock is suspended; however, the internal timer and some of the input pins (*SI\_Int[5:0]*, *SI\_NMI*, *SI\_Reset*, and *SI\_ColdReset*) continue to run. Once the CPU is in instruction-controlled power management mode, any interrupt, NMI, or reset condition causes the CPU to exit this mode and resume normal operation.

The 4KEc core asserts the *SI\_Sleep* signal, which is part of the system interface bus, whenever the WAIT instruction is executed. The assertion of *SI\_Sleep* indicates that the clock has stopped and the 4KEc core is waiting for an interrupt.

### Local clock gating

The majority of the power consumed by the 4KEc core is in the clock tree and clocking registers. The core has support for extensive use of local gated-clocks. Power conscious implementors can use these gated clocks to significantly reduce power consumption within the core.

## 4KEc Core Optional Logic Blocks

The 4KEc core contains several optional logic blocks shown in the block diagram in [Figure 1](#).

### Instruction Cache

The instruction cache is an optional on-chip memory block of up to 64 Kbytes. Because the instruction cache is virtually indexed, the virtual-to-physical address translation occurs in parallel with the cache access rather than having to wait for the physical address translation. The tag holds 22 bits of physical address, a valid bit, and a lock bit. The LRU replacement bits (0-6b per set depending on associativity) are stored in a separate array.

The instruction cache block also contains and manages the instruction line fill buffer. Besides accumulating data to be written to the cache, instruction fetches that reference data in the line fill buffer are serviced either by a bypass of that data, or data coming from the external interface. The instruction cache control logic controls the bypass function.

The 4KEc core supports instruction-cache locking. Cache locking allows critical code or data segments to be locked into the cache on a “per-line” basis, enabling the system programmer to maximize the efficiency of the system cache.

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

### Data Cache

The data cache is an optional on-chip memory block of up to 64 Kbytes. This virtually indexed, physically tagged cache is protected. Because the data cache is virtually indexed, the virtual-to-physical address translation occurs in parallel with the cache access. The tag holds 22 bits of physical address, a valid bit, and a lock bit. There is an additional array holding dirty bits and LRU replacement algorithm bits (0-6b depending on associativity) for each set of the cache.

In addition to instruction-cache locking, the 4KEc core also supports a data-cache locking mechanism identical to the instruction cache. Critical data segments are locked into the cache on a “per-line” basis. The locked contents can be

updated on a store hit, but cannot be selected for replacement on a cache miss.

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

### Cache Memory Configuration

The 4KEc core incorporates on-chip instruction and data caches that can each be accessed in a single processor cycle. Each cache has its own 32-bit data path and can be accessed in the same pipeline clock cycle. Table 9 lists the 4KEc core instruction and data cache attributes.

Table 9 4KEc Core Instruction and Data Cache Attributes

Parameter	Instruction	Data
Size	0 - 64 Kbytes	0 - 64 Kbytes
Organization	1 - 4 way set associative	1 - 4 way set associative
Line Size	16 bytes	16 bytes
Read Unit	32 bits	32 bits
Write Policies	na	write-through with write allocate, write-through without write allocate, write-back with write allocate
Miss restart after transfer of	miss word	miss word
Cache Locking	per line	per line

### Cache Protocols

The 4KEc core supports the following cache protocols:

- **Uncached:** Addresses in a memory area indicated as uncached are not read from the cache. Stores to such addresses are written directly to main memory, without changing cache contents.
- **Write-through, no write allocate:** Loads and instruction fetches first search the cache, reading main memory only if the desired data does not reside in the

cache. On data store operations, the cache is first searched to see if the target address is cache resident. If it is resident, the cache contents are updated, and main memory is also written. If the cache look-up misses, only main memory is written.

- **Write-through, write allocate:** Similar to above, but stores missing in the cache will cause a cache refill. The store data is then written to both the cache and main memory
- **Write-back, write allocate:** Stores that miss in the cache will cause a cache refill. Store data, however, is only written to the cache. Caches lines that are written by stores will be marked as dirty. If a dirty line is selected for replacement, the cache line will be written back to main memory.

## Scratchpad RAM

The 4KEc core also supports replacing up to one way of each cache with a scratchpad RAM. Scratchpad RAM is accessed via independent external pin interfaces for instruction and data scratchpads. The external block which connects to a scratchpad interface is user-defined and can consist of a variety of devices. The main requirement is that it must be accessible with timing similar to an internal cache RAM. Normally, this means that an index will be driven one cycle, a tag will be driven the following clock, and the scratchpad must return a hit signal and the data in the second clock. The scratchpad can easily contain a large RAM/ROM or memory-mapped registers. Unlike the fixed single-cycle cache timing, however, the scratchpad interface can also accommodate backstalling the core pipeline if data is not available in a single clock. This backstalling capability can be useful for operations which require multi-cycle latency. It can also be used to enable arbitration of external accesses to a shared scratchpad memory.

The core's functional interface to a scratchpad RAM is slightly different than to a regular cache RAM. Additional index bits allow access to a larger array, 1MB of scratchpad RAM versus 4KB for a cache way. These bits come from the virtual address, so on a 4KEc core care must be taken to avoid virtual aliasing. The core does not automatically refill the scratchpad way and will not select it for replacement on cache misses. Additionally, stores that hit in the scratchpad will not generate writes to main memory.

## MIPS16e Application Specific Extension

The 4KEc core has optional support for the MIPS16e ASE. This ASE improves code density through the use of 16-bit encodings of MIPS32 instructions plus some MIPS16e-specific instructions. PC relative loads allow quick access to constants. Save/Restore macro instructions provide for single instruction stack frame setup/teardown for efficient subroutine entry/exit. Sign- and zero-extend instructions improve handling of 8-bit and 16-bit datatypes.

## Coprocessor 2 Interface

The 4KEc core can be configured to have an interface for an on-chip coprocessor. This coprocessor can be tightly coupled to the processor core, allowing high performance solutions integrating a graphics accelerator or DSP, for example.

The coprocessor interface is extensible and standardized on MIPS cores, allowing for design reuse. The 4KEc core supports a subset of the full coprocessor interface standard: 32b data transfer, no Coprocessor 1 support, single issue, in-order data transfer to coprocessor, one out-of-order data transfer from coprocessor.

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

## CorExtend User Defined Instruction Extensions

The optional CorExtend User Defined Instruction (UDI) block enables the implementation of a small number of application-specific instructions that are tightly coupled to the core's execution unit. The interface to the UDI block is internal and not defined externally on the 4KEc Pro core.

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

## EJTAG Debug Support

The 4KEc core provides for an optional Enhanced JTAG (EJTAG) interface for use in the software debug of



application and kernel code. In addition to standard user mode and kernel modes of operation, the 4KEc core provides a Debug mode that is entered after a debug exception (derived from a hardware breakpoint, single-step exception, etc.) is taken and continues until a debug exception return (DERET) instruction is executed. During this time, the processor executes the debug exception handler routine.

Refer to the section called "External Interface Signals" on page 24 for a list of EJTAG interface signals.

The EJTAG interface operates through the Test Access Port (TAP), a serial communication port used for transferring test data in and out of the 4KEc core. In addition to the standard JTAG instructions, special instructions defined in the EJTAG specification define what registers are selected and how they are used.

### Debug Registers

Three debug registers (DEBUG, DEPC, and DESAVE) have been added to the MIPS Coprocessor 0 (CP0) register set. The DEBUG register shows the cause of the debug exception and is used for setting up single-step operations. The DEPC, or Debug Exception Program Counter, register holds the address on which the debug exception was taken. This is used to resume program execution after the debug operation finishes. Finally, the DESAVE, or Debug Exception Save, register enables the saving of general-purpose registers used during execution of the debug exception handler.

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

### EJTAG Hardware Breakpoints

There are several types of simple hardware breakpoints defined in the EJTAG specification. These stop the normal operation of the CPU and force the system into debug mode. There are two types of simple hardware breakpoints implemented in the 4KEc core: Instruction breakpoints and Data breakpoints.

The 4KEc core can be configured with the following breakpoint options:

- No data or instruction breakpoints
- One data and two instruction breakpoints

- Two data and four instruction breakpoints

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

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

### EJTAG Trace

The 4KEc core includes optional support for real-time tracing of instruction addresses, data addresses and data values. The trace information is collected in an on-chip or off-chip memory, for post-capture processing by trace regeneration software.

On-chip trace memory may be configured in size from 0 to 8 MB; it is accessed through the existing EJTAG TAP interface and requires no additional chip pins. Off-chip trace memory is accessed through a special trace probe and can be configured to use 4, 8, or 16 data pins plus a clock.

### Testability

Testability for production testing of the core is supported through the use of internal scan and memory BIST.

### Internal Scan

Full mux-based scan for maximum test coverage is supported, with a configurable number of scan chains. ATPG test coverage can exceed 99%, depending on standard cell libraries and configuration options.

### Memory BIST

Memory BIST for the cache arrays and on-chip trace memory is optional, but can be implemented either through the use of integrated BIST features provided with the core, or inserted with an industry-standard memory BIST CAD tool.

## Integrated Memory BIST

The core provides an integrated memory BIST solution for testing the internal cache SRAMs, using BIST controllers and logic tightly coupled to the cache subsystem. Several parameters associated with the integrated BIST controllers are configurable, including the algorithm (March C+ or IFA-13).

## User-specified Memory BIST

Memory BIST can also be inserted with a CAD tool or other user-specified method. Wrapper modules and signal

buses of configurable width are provided within the core to facilitate this approach.

## Instruction Set

The 4KEc core instruction set complies with the MIPS32 instruction set architecture. Table 10 provides a summary of instructions implemented by the 4KEc core.

Table 10 4KEc Core Instruction Set

Instruction	Description	Function
ADD	Integer Add	$Rd = Rs + Rt$
ADDI	Integer Add Immediate	$Rt = Rs + Immed$
ADDIU	Unsigned Integer Add Immediate	$Rt = Rs +_U Immed$
ADDIUPC	Unsigned Integer Add Immediate to PC (MIPS16 only)	$Rt = PC +_u Immed$
ADDU	Unsigned Integer Add	$Rd = Rs +_U Rt$
AND	Logical AND	$Rd = Rs \& Rt$
ANDI	Logical AND Immediate	$Rt = Rs \& (0_{16}    Immed)$
BC2F	Branch On COP2 Condition False	if COP2Condition(cc) == 0 PC += (int)offset
BC2FL	Branch On COP2 Condition False Likely	if COP2Condition(cc) == 0 PC += (int)offset else Ignore Next Instruction
BC2T	Branch On COP2 Condition True	if COP2Condition(cc) == 1 PC += (int)offset
BC2TL	Branch On COP2 Condition True Likely	if COP2Condition(cc) == 1 PC += (int)offset else Ignore Next Instruction
BEQ	Branch On Equal	if Rs == Rt PC += (int)offset
BEQL	Branch On Equal Likely	if Rs == Rt PC += (int)offset else Ignore Next Instruction
BGEZ	Branch on Greater Than or Equal To Zero	if !Rs[31] PC += (int)offset
BGEZAL	Branch on Greater Than or Equal To Zero And Link	GPR[31] = PC + 8 if !Rs[31] PC += (int)offset

Table 10 4KEc Core Instruction Set (Continued)

Instruction	Description	Function
BGEZALL	Branch on Greater Than or Equal To Zero And Link Likely	GPR[31] = PC + 8 if !Rs[31] PC += (int)offset else Ignore Next Instruction
BGEZL	Branch on Greater Than or Equal To Zero Likely	if !Rs[31] PC += (int)offset else Ignore Next Instruction
BGTZ	Branch on Greater Than Zero	if !Rs[31] && Rs != 0 PC += (int)offset
BGTZL	Branch on Greater Than Zero Likely	if !Rs[31] && Rs != 0 PC += (int)offset else Ignore Next Instruction
BLEZ	Branch on Less Than or Equal to Zero	if Rs[31]    Rs == 0 PC += (int)offset
BLEZL	Branch on Less Than or Equal to Zero Likely	if Rs[31]    Rs == 0 PC += (int)offset else Ignore Next Instruction
BLTZ	Branch on Less Than Zero	if Rs[31] PC += (int)offset
BLTZAL	Branch on Less Than Zero And Link	GPR[31] = PC + 8 if Rs[31] PC += (int)offset
BLTZALL	Branch on Less Than Zero And Link Likely	GPR[31] = PC + 8 if Rs[31] PC += (int)offset else Ignore Next Instruction
BLTZL	Branch on Less Than Zero Likely	if Rs[31] PC += (int)offset else Ignore Next Instruction
BNE	Branch on Not Equal	if Rs != Rt PC += (int)offset
BNEL	Branch on Not Equal Likely	if Rs != Rt PC += (int)offset else Ignore Next Instruction
BREAK	Breakpoint	Break Exception
CACHE	Cache Operation	See Software User's Manual
CFC2	Move Control Word From Coprocessor 2	Rt = CCR[2, n]
CLO	Count Leading Ones	Rd = NumLeadingOnes(Rs)
CLZ	Count Leading Zeroes	Rd = NumLeadingZeroes(Rs)

Table 10 4KEc Core Instruction Set (Continued)

Instruction	Description	Function
COP0	Coprocessor 0 Operation	See Software User's Manual
COP2	Coprocessor 2 Operation	See Coprocessor 2 Description
CTC2	Move Control Word To Coprocessor 2	$CCR[2, n] = Rt$
DERET	Return from Debug Exception	PC = DEPC Exit Debug Mode
DI	Atomically Disable Interrupts	$Rt = Status; Status_{IE} = 0$
DIV	Divide	LO = (int)Rs / (int)Rt HI = (int)Rs % (int)Rt
DIVU	Unsigned Divide	LO = (uns)Rs / (uns)Rt HI = (uns)Rs % (uns)Rt
EHB	Execution Hazard Barrier	Stop instruction execution until execution hazards are cleared
EI	Atomically Enable Interrupts	$Rt = Status; Status_{IE} = 1$
ERET	Return from Exception	if SR[2] PC = ErrorEPC else PC = EPC SR[1] = 0 SR[2] = 0 LL = 0
EXT	Extract Bit Field	$Rt = \text{ExtractField}(Rs, pos, size)$
INS	Insert Bit Field	$Rt = \text{InsertField}(Rs, Rt, pos, size)$
J	Unconditional Jump	PC = PC[31:28]    offset<<2
JAL	Jump and Link	GPR[31] = PC + 8 PC = PC[31:28]    offset<<2
JALR	Jump and Link Register	Rd = PC + 8 PC = Rs
JALR.HB	Jump and Link Register with Hazard Barrier	Like JALR, but also clears execution and instruction hazards
JALRC	Jump and Link Register Compact - do not execute instruction in jump delay slot(MIPS16 only)	Rd = PC + 2 PC = Rs
JR	Jump Register	PC = Rs
JR.HB	Jump Register with Hazard Barrier	Like JR, but also clears execution and instruction hazards
JRC	Jump Register Compact - do not execute instruction in jump delay slot (MIPS16 only)	PC = Rs
LB	Load Byte	$Rt = (\text{byte})\text{Mem}[Rs+\text{offset}]$

Table 10 4KEc Core Instruction Set (Continued)

Instruction	Description	Function
LBU	Unsigned Load Byte	$Rt = (\text{ubyte})\text{Mem}[Rs+\text{offset}]$
LH	Load Halfword	$Rt = (\text{half})\text{Mem}[Rs+\text{offset}]$
LHU	Unsigned Load Halfword	$Rt = (\text{uhalf})\text{Mem}[Rs+\text{offset}]$
LL	Load Linked Word	$Rt = \text{Mem}[Rs+\text{offset}]$ LL = 1 LLAdr = Rs + offset
LUI	Load Upper Immediate	$Rt = \text{immediate} \ll 16$
LW	Load Word	$Rt = \text{Mem}[Rs+\text{offset}]$
LWC2	Load Word To Coprocessor 2	$\text{CPR}[2, n, 0] = \text{Mem}[Rs+\text{offset}]$
LWPC	Load Word, PC relative	$Rt = \text{Mem}[PC+\text{offset}]$
LWL	Load Word Left	See Software User's Manual
LWR	Load Word Right	See Software User's Manual
MADD	Multiply-Add	$\text{HI} \mid \text{LO} += (\text{int})Rs * (\text{int})Rt$
MADDU	Multiply-Add Unsigned	$\text{HI} \mid \text{LO} += (\text{uns})Rs * (\text{uns})Rt$
MFC0	Move From Coprocessor 0	$Rt = \text{CPR}[0, Rd, sel]$
MFC2	Move From Coprocessor 2	$Rt = \text{CPR}[2, Rd, sel]$
MFHC2	Move From High Half of Coprocessor 2	$Rt = \text{CPR}[2, Rd, sel]_{63..32}$
MFHI	Move From HI	$Rd = \text{HI}$
MFLO	Move From LO	$Rd = \text{LO}$
MOVN	Move Conditional on Not Zero	if $Rt \neq 0$ then $Rd = Rs$
MOVZ	Move Conditional on Zero	if $Rt = 0$ then $Rd = Rs$
MSUB	Multiply-Subtract	$\text{HI} \mid \text{LO} -= (\text{int})Rs * (\text{int})Rt$
MSUBU	Multiply-Subtract Unsigned	$\text{HI} \mid \text{LO} -= (\text{uns})Rs * (\text{uns})Rt$
MTC0	Move To Coprocessor 0	$\text{CPR}[0, n, Sel] = Rt$
MTC2	Move To Coprocessor 2	$\text{CPR}[2, n, sel] = Rt$
MTHC2	Move To High Half of Coprocessor 2	$\text{CPR}[2, Rd, sel] = Rt \mid \mid$ $\text{CPR}[2, Rd, sel]_{31..0}$
MTHI	Move To HI	$\text{HI} = Rs$
MTLO	Move To LO	$\text{LO} = Rs$
MUL	Multiply with register write	$\text{HI} \mid \text{LO} = \text{Unpredictable}$ $Rd = ((\text{int})Rs * (\text{int})Rt)_{31..0}$
MULT	Integer Multiply	$\text{HI} \mid \text{LO} = (\text{int})Rs * (\text{int})Rd$

Table 10 4KEc Core Instruction Set (Continued)

Instruction	Description	Function
MULTU	Unsigned Multiply	$HI \mid LO = (uns)Rs * (uns)Rd$
NOR	Logical NOR	$Rd = \sim(Rs \mid Rt)$
OR	Logical OR	$Rd = Rs \mid Rt$
ORI	Logical OR Immediate	$Rt = Rs \mid Immed$
PREF	Prefetch	Load Specified Line into Cache
RDHWR	Read Hardware Register	Allows unprivileged access to registers enabled by HWREna register
RDPGPR	Read GPR from Previous Shadow Set	$Rt = SGPR[SRSCtl_{PSS}, Rd]$
RESTORE	Restore registers and deallocate stack frame (MIPS16 only)	See Software User's Manual
ROTR	Rotate Word Right	$Rd = Rt_{sa-1..0} \parallel Rt_{31..sa}$
ROTRV	Rotate Word Right Variable	$Rd = Rt_{Rs-1..0} \parallel Rt_{31..Rs}$
SAVE	Save registers and allocate stack frame (MIPS16 only)	See Software User's Manual
SB	Store Byte	$(byte)Mem[Rs+offset] = Rt$
SC	Store Conditional Word	if $LL = 1$ $mem[Rs+offset] = Rt$ $Rt = LL$
SDBBP	Software Debug Break Point	Trap to SW Debug Handler
SEB	Sign Extend Byte	$Rd = (byte)Rs$
SEH	Sign Extend Half	$Rd = (half)Rs$
SH	Store Half	$(half)Mem[Rs+offset] = Rt$
SLL	Shift Left Logical	$Rd = Rt \ll sa$
SLLV	Shift Left Logical Variable	$Rd = Rt \ll Rs[4:0]$
SLT	Set on Less Than	if $(int)Rs < (int)Rt$ $Rd = 1$ else $Rd = 0$
SLTI	Set on Less Than Immediate	if $(int)Rs < (int)Immed$ $Rt = 1$ else $Rt = 0$
SLTIU	Set on Less Than Immediate Unsigned	if $(uns)Rs < (uns)Immed$ $Rt = 1$ else $Rt = 0$

Table 10 4KEc Core Instruction Set (Continued)

Instruction	Description	Function
SLTU	Set on Less Than Unsigned	if (uns)Rs < (uns)Immed Rd = 1 else Rd = 0
SRA	Shift Right Arithmetic	Rd = (int)Rt >> sa
SRAV	Shift Right Arithmetic Variable	Rd = (int)Rt >> Rs[4:0]
SRL	Shift Right Logical	Rd = (uns)Rt >> sa
SRLV	Shift Right Logical Variable	Rd = (uns)Rt >> Rs[4:0]
SSNOP	Superscalar Inhibit No Operation	NOP
SUB	Integer Subtract	Rt = (int)Rs - (int)Rd
SUBU	Unsigned Subtract	Rt = (uns)Rs - (uns)Rd
SW	Store Word	Mem[Rs+offset] = Rt
SWC2	Store Word From Coprocessor 2	Mem[Rs+offset] = CPR[2,n,0]
SWL	Store Word Left	See Software User's Manual
SWR	Store Word Right	See Software User's Manual
SYNC	Synchronize	See Software User's Manual
SYSCALL	System Call	SystemCallException
TEQ	Trap if Equal	if Rs == Rt TrapException
TEQI	Trap if Equal Immediate	if Rs == (int)Immed TrapException
TGE	Trap if Greater Than or Equal	if (int)Rs >= (int)Rt TrapException
TGEI	Trap if Greater Than or Equal Immediate	if (int)Rs >= (int)Immed TrapException
TGEIU	Trap if Greater Than or Equal Immediate Unsigned	if (uns)Rs >= (uns)Immed TrapException
TGEU	Trap if Greater Than or Equal Unsigned	if (uns)Rs >= (uns)Rt TrapException
TLBWI	Write Indexed TLB Entry	See Software Users Manual
TLBWR	Write Random TLB Entry	See Software Users Manual
TLBP	Probe TLB for Matching Entry	See Software Users Manual
TLBR	Read Index for TLB Entry	See Software Users Manual
TLT	Trap if Less Than	if (int)Rs < (int)Rt TrapException
TLTI	Trap if Less Than Immediate	if (int)Rs < (int)Immed TrapException

Table 10 4KEc Core Instruction Set (Continued)

Instruction	Description	Function
TLTIU	Trap if Less Than Immediate Unsigned	if (uns)Rs < (uns)Immed TrapException
TLTU	Trap if Less Than Unsigned	if (uns)Rs < (uns)Rt TrapException
TNE	Trap if Not Equal	if Rs != Rt TrapException
TNEI	Trap if Not Equal Immediate	if Rs != (int)Immed TrapException
WAIT	Wait for Interrupts	Stall until interrupt occurs
WRPGPR	Write to GPR in Previous Shadow Set	SGPR[SRSCtlPSS, Rd] = Rt
WSBH	Word Swap Bytes Within HalfWords	Rd = Rt <sub>23..16</sub>    Rt <sub>31..24</sub>    Rt <sub>7..0</sub>    Rt <sub>15..8</sub>
XOR	Exclusive OR	Rd = Rs ^ Rt
XORI	Exclusive OR Immediate	Rt = Rs ^ (uns)Immed
ZEB	Zero extend byte (MIPS16 only)	Rt = (ubyte) Rs
ZEH	Zero extend half (MIPS16 only)	Rt = (uhalf) Rs

## External Interface Signals

This section describes the signal interface of the 4KEc microprocessor core.

The pin direction key for the signal descriptions is shown in Table 11 below.

The 4KEc core signals are listed in Table 12 below. Note that the signals are grouped by logical function, not by expected physical location. All signals, with the exception of *EJ\_TRST\_N*, are active-high signals. *EJ\_DINT* and *SI\_NMI* go through edge-detection logic so that only one exception is taken each time they are asserted.

Table 11 4KEc Core Signal Direction Key

Dir	Description
I	Input to the 4KEc core sampled on the rising edge of the appropriate CLK signal.
O	Output of the 4KEc core, unless otherwise noted, driven at the rising edge of the appropriate CLK signal.
A	Asynchronous inputs that are synchronized by the core.
S	Static input to the 4KEc core. These signals are normally tied to either power or ground and should not change state while <i>SI_ColdReset</i> is deasserted.



Table 12 4KEc Signal Descriptions

Signal Name	Type	Description
<b>System Interface</b>		
<i>Clock Signals:</i>		
SI_ClkIn	I	Clock Input. All inputs and outputs, except a few of the EJTAG signals, are sampled and/or asserted relative to the rising edge of this signal.
SI_ClkOut	O	Reference Clock for the External Bus Interface. This clock signal provides a reference for deskewing any clock insertion delay created by the internal clock buffering in the core.
<i>Reset Signals:</i>		
SI_ColdReset	A	Hard/Cold Reset Signal. Causes a Reset Exception in the core.
SI_NMI	A	Non-Maskable Interrupt. An edge detect is used on this signal. When this signal is sampled asserted (high) one clock after being sampled deasserted, an NMI is posted to the core.
SI_Reset	A	Soft/Warm Reset Signal. Causes a Reset Exception in the core. Sets Status.SR bit (if <i>SI_ColdReset</i> is not asserted), but is otherwise ORed with <i>SI_ColdReset</i> before it is used internally.
<i>Power Management Signals:</i>		
SI_ERL	O	This signal represents the state of the ERL bit (2) in the CP0 Status register and indicates the error level. The core asserts <i>SI_ERL</i> whenever a Reset, Soft Reset, or NMI exception is taken.
SI_EXL	O	This signal represents the state of the EXL bit (1) in the CP0 Status register and indicates the exception level. The core asserts <i>SI_EXL</i> whenever any exception other than a Reset, Soft Reset, NMI, or Debug exception is taken.
SI_RP	O	This signal represents the state of the RP bit (27) in the CP0 Status register. Software can write this bit to indicate that a reduced power mode may be entered.
SI_Sleep	O	This signal is asserted by the core whenever the WAIT instruction is executed. The assertion of this signal indicates that the clock has stopped and that the core is waiting for an interrupt.
<i>Interrupt Signals:</i>		
SI_EICPresent	S	Indicates whether an external interrupt controller is present. Value is visible to software in the <i>Config3_VEIC</i> register field.
SI_EISS[3:0]	I	General purpose register shadow set number to be used when servicing an interrupt in EIC interrupt mode.
SI_IAck	O	Interrupt acknowledge indication for use in external interrupt controller mode. This signal is active for a single <i>SI_ClkIn</i> cycle when an interrupt is taken. When the processor initiates the interrupt exception, it loads the value of the <i>SI_Int[5:0]</i> pins into the <i>Cause_RIPL</i> field (overlaid with <i>Cause_IP7..IP2</i> ), and signals the external interrupt controller to notify it that the current interrupt request is being serviced. This allows the controller to advance to another pending higher-priority interrupt, if desired.

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description
SI_Int[5:0]	I/A	<p>Active high Interrupt pins. These signals are driven by external logic and when asserted indicate an interrupt exception to the core. The interpretation of these signals depends on the interrupt mode in which the core is operating; the interrupt mode is selected by software.</p> <p>The <i>SI_Int</i> signals go through synchronization logic and can be asserted asynchronously to <i>SI_ClkIn</i>. In External Interrupt Controller (EIC) mode, however, the interrupt pins are interpreted as an encoded value, so they must be asserted synchronously to <i>SI_ClkIn</i> to guarantee that all bits are received by the core in a particular cycle.</p> <p>The interrupt pins are level sensitive and should remain asserted until the interrupt has been serviced.</p> <p>In Release 1 Interrupt Compatibility mode:</p> <ul style="list-style-type: none"> <li>All 6 interrupt pins have the same priority as far as the hardware is concerned.</li> <li>Interrupts are non-vectored.</li> </ul> <p>In Vectored Interrupt (VI) mode:</p> <ul style="list-style-type: none"> <li>The <i>SI_Int</i> pins are interpreted as individual hardware interrupt requests.</li> <li>Internally, the core prioritizes the hardware interrupts and chooses an interrupt vector.</li> </ul> <p>In External Interrupt Controller (EIC) mode:</p> <ul style="list-style-type: none"> <li>An external block prioritizes its various interrupt requests and produces a vector number of the highest priority interrupt to be serviced.</li> <li>The vector number is driven on the <i>SI_Int</i> pins, and is treated as a 6-bit encoded value in the range of 0..63.</li> <li>When the core starts the interrupt exception, signaled by the assertion of <i>SI_IAck</i>, it loads the value of the <i>SI_Int[5:0]</i> pins into the <i>Cause_RIPL</i> field (overlaid with <i>Cause_IP7..IP2</i>). The interrupt controller can then signal another interrupt.</li> </ul>
SI_IPL[5:0]	O	Current interrupt priority level from the <i>Status_IPL</i> register field, provided for use by an external interrupt controller. This value is updated whenever <i>SI_IAck</i> is asserted.
SI_IPTI[2:0]	S	Indicates the <i>SI_Int</i> hardware interrupt pin that the timer interrupt pin ( <i>SI_TimerInt</i> ) is combined with external to the core. The value of this bus is visible to software in the <i>IntCtl_IPTI</i> register field.
SI_SWInt[1:0]	O	Software interrupt request. These signals represent the value in the <i>IP[1:0]</i> field of the <i>Cause</i> register. They are provided for use by an external interrupt controller.
SI_TimerInt	O	<p>Timer interrupt indication. This signal is asserted whenever the <i>Count</i> and <i>Compare</i> registers match and is deasserted when the <i>Compare</i> register is written. This hardware pin represents the value of the <i>Cause_TT</i> register field.</p> <p>For Release 1 Interrupt Compatibility mode or Vectored Interrupt mode:</p> <p>In order to generate a timer interrupt, the <i>SI_TimerInt</i> signal needs to be brought back into the 4KEc core on one of the six <i>SI_Int</i> interrupt pins in a system-dependent manner. Traditionally, this has been accomplished by muxing <i>SI_TimerInt</i> with <i>SI_Int[5]</i>. Exposing <i>SI_TimerInt</i> as an output allows more flexibility for the system designer. Timer interrupts can be muxed or ORed into one of the interrupts, as desired in a particular system. The <i>SI_Int</i> hardware interrupt pin with which the <i>SI_TimerInt</i> signal is merged is indicated via the <i>SI_IPTI</i> static input pins.</p> <p>For External Interrupt Controller (EIC) mode:</p> <p>The <i>SI_TimerInt</i> signal is provided to the external interrupt controller, which then prioritizes the timer interrupt with all other interrupt sources, as desired. The controller then encodes the desired interrupt value on the <i>SI_Int</i> pins. Since <i>SI_Int</i> is usually encoded, the <i>SI_IPTI</i> pins are not meaningful in EIC mode.</p>

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description										
<i>Configuration Inputs:</i>												
SI_CPUNum[9:0]	S	Unique identifier to specify an individual core in a multi-processor system. The hardware value specified on these pins is available in the <i>CPUNum</i> field of the <i>EBase</i> register, so it can be used by software to distinguish a particular processor. In a single processor system, this value should be set to zero.										
SI_Endian	S	Indicates the base endianness of the core. <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>EB_Endian</th> <th>Base Endian Mode</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Little Endian</td> </tr> <tr> <td>1</td> <td>Big Endian</td> </tr> </tbody> </table>	EB_Endian	Base Endian Mode	0	Little Endian	1	Big Endian				
EB_Endian	Base Endian Mode											
0	Little Endian											
1	Big Endian											
SI_MergeMode[1:0]	S	The state of these signals determines the merge mode for the 16-byte collapsing write buffer. <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Encoding</th> <th>Merge Mode</th> </tr> </thead> <tbody> <tr> <td>00<sub>2</sub></td> <td>No Merge</td> </tr> <tr> <td>01<sub>2</sub></td> <td>Reserved</td> </tr> <tr> <td>10<sub>2</sub></td> <td>Full Merge</td> </tr> <tr> <td>11<sub>2</sub></td> <td>Reserved</td> </tr> </tbody> </table>	Encoding	Merge Mode	00 <sub>2</sub>	No Merge	01 <sub>2</sub>	Reserved	10 <sub>2</sub>	Full Merge	11 <sub>2</sub>	Reserved
Encoding	Merge Mode											
00 <sub>2</sub>	No Merge											
01 <sub>2</sub>	Reserved											
10 <sub>2</sub>	Full Merge											
11 <sub>2</sub>	Reserved											
SI_SimpleBE[1:0]	S	The state of these signals can constrain the core to only generate certain byte enables on EC™ interface transactions. This eases connection to some existing bus standards. <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>SI_SimpleBE[1:0]</th> <th>Byte Enable Mode</th> </tr> </thead> <tbody> <tr> <td>00<sub>2</sub></td> <td>All BEs allowed</td> </tr> <tr> <td>01<sub>2</sub></td> <td>Naturally aligned bytes, half-words, and words only</td> </tr> <tr> <td>10<sub>2</sub></td> <td>Reserved</td> </tr> <tr> <td>11<sub>2</sub></td> <td>Reserved</td> </tr> </tbody> </table>	SI_SimpleBE[1:0]	Byte Enable Mode	00 <sub>2</sub>	All BEs allowed	01 <sub>2</sub>	Naturally aligned bytes, half-words, and words only	10 <sub>2</sub>	Reserved	11 <sub>2</sub>	Reserved
SI_SimpleBE[1:0]	Byte Enable Mode											
00 <sub>2</sub>	All BEs allowed											
01 <sub>2</sub>	Naturally aligned bytes, half-words, and words only											
10 <sub>2</sub>	Reserved											
11 <sub>2</sub>	Reserved											
<b>External Bus Interface</b>												
EB_ARdy	I	Indicates whether the target is ready for a new address. The core will not complete the address phase of a new bus transaction until the clock cycle after <i>EB_ARdy</i> is sampled asserted.										
EB_AValid	O	When asserted, indicates that the values on the address bus and access types lines are valid, signifying the beginning of a new bus transaction. <i>EB_AValid</i> must always be valid.										
EB_Instr	O	When asserted, indicates that the transaction is an instruction fetch versus a data reference. <i>EB_Instr</i> is only valid when <i>EB_AValid</i> is asserted.										
EB_Write	O	When asserted, indicates that the current transaction is a write. This signal is only valid when <i>EB_AValid</i> is asserted.										
EB_Burst	O	When asserted, indicates that the current transaction is part of a cache fill or a write burst. Note that there is redundant information contained in <i>EB_Burst</i> , <i>EB_BFirst</i> , <i>EB_BLast</i> , and <i>EB_BLen</i> . This is done to simplify the system design—the information can be used in whatever form is easiest.										

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description															
EB_BFirst	O	When asserted, indicates the beginning of the burst. <i>EB_BFirst</i> is always valid.															
EB_BLast	O	When asserted, indicates the end of the burst. <i>EB_BLast</i> is always valid.															
EB_BLen[1:0]	O	Indicates the length of the burst. This signal is only valid when <i>EB_AValid</i> is asserted. <table border="1" data-bbox="760 436 1206 636"> <thead> <tr> <th>EB_BLength[1:0]</th> <th>Burst Length</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>reserved</td> </tr> <tr> <td>1</td> <td>4</td> </tr> <tr> <td>2</td> <td>reserved</td> </tr> <tr> <td>3</td> <td>reserved</td> </tr> </tbody> </table>	EB_BLength[1:0]	Burst Length	0	reserved	1	4	2	reserved	3	reserved					
EB_BLength[1:0]	Burst Length																
0	reserved																
1	4																
2	reserved																
3	reserved																
EB_SBlock	S	Static input which determines burst order. When asserted, sub-block ordering is used. When deasserted, sequential addressing is used.															
EB_BE[3:0]	O	Indicates which bytes of the <i>EB_RData</i> or <i>EB_WData</i> buses are involved in the current transaction. If an <i>EB_BE</i> signal is asserted, the associated byte is being read or written. <i>EB_BE</i> lines are only valid while <i>EB_AValid</i> is asserted. <table border="1" data-bbox="665 850 1304 1073"> <thead> <tr> <th>EB_BE Signal</th> <th>Read Data Bits Sampled</th> <th>Write Data Bits Driven Valid</th> </tr> </thead> <tbody> <tr> <td>EB_BE[0]</td> <td>EB_RData[7:0]</td> <td>EB_WData[7:0]</td> </tr> <tr> <td>EB_BE[1]</td> <td>EB_RData[15:8]</td> <td>EB_WData[15:8]</td> </tr> <tr> <td>EB_BE[2]</td> <td>EB_RData[23:16]</td> <td>EB_WData[23:16]</td> </tr> <tr> <td>EB_BE[3]</td> <td>EB_RData[31:24]</td> <td>EB_WData[31:24]</td> </tr> </tbody> </table>	EB_BE Signal	Read Data Bits Sampled	Write Data Bits Driven Valid	EB_BE[0]	EB_RData[7:0]	EB_WData[7:0]	EB_BE[1]	EB_RData[15:8]	EB_WData[15:8]	EB_BE[2]	EB_RData[23:16]	EB_WData[23:16]	EB_BE[3]	EB_RData[31:24]	EB_WData[31:24]
EB_BE Signal	Read Data Bits Sampled	Write Data Bits Driven Valid															
EB_BE[0]	EB_RData[7:0]	EB_WData[7:0]															
EB_BE[1]	EB_RData[15:8]	EB_WData[15:8]															
EB_BE[2]	EB_RData[23:16]	EB_WData[23:16]															
EB_BE[3]	EB_RData[31:24]	EB_WData[31:24]															
EB_A[35:2]	O	Address lines for external bus. Only valid when <i>EB_AValid</i> is asserted. <i>EB_A[35:32]</i> are tied to 0 in this core.															
EB_WData[31:0]	O	Output data for writes.															
EB_RData[31:0]	I	Input Data for reads.															
EB_RdVal	I	Indicates that the target is driving read data on <i>EB_RData</i> lines. <i>EB_RdVal</i> must always be valid. <i>EB_RdVal</i> may never be sampled asserted until the rising edge after the corresponding <i>EB_ARdy</i> was sampled asserted.															
EB_WDRdy	I	Indicates that the target of a write is ready. The <i>EB_WData</i> lines can change in the next clock cycle. <i>EB_WDRdy</i> will not be sampled until the rising edge where the corresponding <i>EB_ARdy</i> is sampled asserted.															
EB_RBErr	I	Bus error indicator for read transactions. <i>EB_RBErr</i> is sampled on every rising clock edge until an active sampling of <i>EB_RdVal</i> . <i>EB_RBErr</i> sampled with asserted <i>EB_RdVal</i> indicates a bus error during read. <i>EB_RBErr</i> must be deasserted in idle phases.															
EB_WBErr	I	Bus error indicator for write transactions. <i>EB_WBErr</i> is sampled on the rising clock edge following an active sample of <i>EB_WDRdy</i> . <i>EB_WBErr</i> must be deasserted in idle phases.															
EB_EWBE	I	Indicates that any external write buffers are empty. The external write buffers must deassert <i>EB_EWBE</i> in the cycle after the corresponding <i>EB_WDRdy</i> is asserted and keep <i>EB_EWBE</i> deasserted until the external write buffers are empty.															
EB_WWBE	O	When asserted, indicates that the core is waiting for external write buffers to empty.															
<b>Coprocessor Interface</b>																	

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description
Instruction dispatch: These signals are used to transfer an instruction from the 4KEc core to the COP2 coprocessor.		
CP2_ir_0[31:0]	O	Coprocessor Arithmetic and To/From Instruction Word. Valid in the cycle before <i>CP2_as_0</i> , <i>CP2_ts_0</i> or <i>CP2_fs_0</i> is asserted.
CP2_irenable_0	O	Enable Instruction Registering. When deasserted, no instruction strobes will be asserted in the following cycle. When asserted, there <i>may</i> be an instruction strobe asserted in the following cycle. Instruction strobes include <i>CP2_as_0</i> , <i>CP2_ts_0</i> , <i>CP2_fs_0</i> . <b>Note:</b> This is the only late signal in the interface. The intended function is to use this signal as a clock gate condition on the capture latches in the coprocessor for <i>CP2_ir_0[31:0]</i> .
CP2_as_0	O	Coprocessor2 Arithmetic Instruction Strobe. Asserted in the cycle after an arithmetic coprocessor2 instruction is available on <i>CP2_ir_0[31:0]</i> . If <i>CP2_abusy_0</i> was asserted in the previous cycle, this signal will not be asserted. This signal will never be asserted in the same cycle that <i>CP2_ts_0</i> or <i>CP2_fs_0</i> is asserted.
CP2_abusy_0	I	Coprocessor2 Arithmetic Busy. When asserted, a coprocessor2 arithmetic instruction will not be dispatched. <i>CP2_as_0</i> will not be asserted in the cycle after this signal is asserted.
CP2_ts_0	O	Coprocessor2 To Strobe. Asserted in the cycle after a To COP2 Op instruction is available on <i>CP2_ir_0[31:0]</i> . If <i>CP2_tbusy</i> was asserted in the previous cycle, this signal will not be asserted. This signal will never be asserted in the same cycle that <i>CP2_as_0</i> or <i>CP2_fs_0</i> is asserted.
CP2_tbusy_0	I	To Coprocessor2 Busy. When asserted, a To COP2 Op will not be dispatched. <i>CP2_ts_0</i> will not be asserted in the cycle after this signal is asserted.
CP2_fs_0	O	Coprocessor2 From Strobe. Asserted in the cycle after a From COP2 Op instruction is available on <i>CP2_ir_0[31:0]</i> . If <i>CP2_fbusy_0</i> was asserted in the previous cycle, this signal will not be asserted. This signal will never be asserted in the same cycle that <i>CP2_as_0</i> or <i>CP2_ts_0</i> is asserted.
CP2_fbusy_0	I	From Coprocessor2 Busy. When asserted, a From COP2 Op will not be dispatched. <i>CP2_fs_0</i> will not be asserted in the cycle after this signal is asserted.
CP2_endian_0	O	Big Endian Byte Ordering. When asserted, the processor is using big endian byte ordering for the dispatched instruction. When deasserted, the processor is using little-endian byte ordering. Valid the cycle before <i>CP2_as_0</i> , <i>CP2_fs_0</i> or <i>CP2_ts_0</i> is asserted.
CP2_inst32_0	O	MIPS32 Compatibility Mode - Instructions. When asserted, the dispatched instruction is restricted to the MIPS32 subset of instructions. Please refer to the MIPS64 architecture specification for a complete description of MIPS32 compatibility mode. Valid the cycle before <i>CP2_as_0</i> , <i>CP2_fs_0</i> or <i>CP2_ts_0</i> is asserted. <b>Note:</b> The 4KEc core is a MIPS32 core, and will only issue MIPS32 instructions. Thus <i>CP2_inst32_0</i> is tied high.
CP2_kd_mode_0	O	Kernel/Debug Mode. When asserted, the processor is running in kernel or debug mode. Can be used to enable “privileged” coprocessor instructions. Valid the cycle before <i>CP2_as_0</i> , <i>CP2_fs_0</i> or <i>CP2_ts_0</i> is asserted.
To Coprocessor Data: These signals are used when data is sent from the 4KEc core to the COP2 coprocessor, as part of completing a To Coprocessor instruction.		
CP2_tds_0	O	Coprocessor To Data Strobe. Asserted when To COP Op data is available on <i>CP2_tdata_0[31:0]</i> .

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description																		
CP2_torder_0[2:0]	O	<p>Coprocessor To Order. Specifies which outstanding To COP Op the data is for. Valid only when <i>CP2_tds_0</i> is asserted.</p> <table border="1"> <thead> <tr> <th><i>CP2_torder_0</i>[2:0]</th> <th>Order</th> </tr> </thead> <tbody> <tr> <td>000<sub>2</sub></td> <td>Oldest outstanding To COP Op data transfer</td> </tr> <tr> <td>001<sub>2</sub></td> <td>2nd oldest To COP Op data transfer.</td> </tr> <tr> <td>010<sub>2</sub></td> <td>3rd oldest To COP Op data transfer.</td> </tr> <tr> <td>011<sub>2</sub></td> <td>4th oldest To COP Op data transfer.</td> </tr> <tr> <td>100<sub>2</sub></td> <td>5th oldest To COP Op data transfer.</td> </tr> <tr> <td>101<sub>2</sub></td> <td>6th oldest To COP Op data transfer.</td> </tr> <tr> <td>110<sub>2</sub></td> <td>7th oldest To COP Op data transfer.</td> </tr> <tr> <td>111<sub>2</sub></td> <td>8th oldest To COP Op data transfer.</td> </tr> </tbody> </table> <p><b>Note:</b> The 4KEc core will never send Data Out-of-Order, thus <i>CP2_torder_0</i>[2:0] is tied to 000<sub>2</sub>.</p>	<i>CP2_torder_0</i> [2:0]	Order	000 <sub>2</sub>	Oldest outstanding To COP Op data transfer	001 <sub>2</sub>	2nd oldest To COP Op data transfer.	010 <sub>2</sub>	3rd oldest To COP Op data transfer.	011 <sub>2</sub>	4th oldest To COP Op data transfer.	100 <sub>2</sub>	5th oldest To COP Op data transfer.	101 <sub>2</sub>	6th oldest To COP Op data transfer.	110 <sub>2</sub>	7th oldest To COP Op data transfer.	111 <sub>2</sub>	8th oldest To COP Op data transfer.
<i>CP2_torder_0</i> [2:0]	Order																			
000 <sub>2</sub>	Oldest outstanding To COP Op data transfer																			
001 <sub>2</sub>	2nd oldest To COP Op data transfer.																			
010 <sub>2</sub>	3rd oldest To COP Op data transfer.																			
011 <sub>2</sub>	4th oldest To COP Op data transfer.																			
100 <sub>2</sub>	5th oldest To COP Op data transfer.																			
101 <sub>2</sub>	6th oldest To COP Op data transfer.																			
110 <sub>2</sub>	7th oldest To COP Op data transfer.																			
111 <sub>2</sub>	8th oldest To COP Op data transfer.																			
CP2_tordlim_0[2:0]	S	<p>To Coprocessor Data Out-of-Order Limit. This signal forces the integer processor core to limit how much it can reorder To COP Data. The value on this signal corresponds to the maximum allowed value to be used on <i>CP2_torder_0</i>[2:0].</p> <p><b>Note:</b> The 4KEc core will never send Data Out-of-Order, thus <i>CP2_tordlim_0</i>[2:0] is ignored.</p>																		
CP2_tdata_0[31:0]	O	To Coprocessor Data. Data to be transferred to the coprocessor. Valid when <i>CP2_tds_0</i> is asserted.																		
From Coprocessor Data: These signals are used when data is sent to the 4KEc core from the COP2 coprocessor, as part of completing a From Coprocessor instruction.																				
CP2_fds_0	I	Coprocessor From Data Strobe. Asserted when From COP Op data is available on <i>CP2_fdata_0</i> [31:0].																		
CP2_forder_0[2:0]	I	<p>Coprocessor From Order. Specifies which outstanding From COP Op the data is for. Valid only when <i>CP2_fds_0</i> is asserted.</p> <table border="1"> <thead> <tr> <th><i>CP2_forder_0</i>[2:0]</th> <th>Order</th> </tr> </thead> <tbody> <tr> <td>000<sub>2</sub></td> <td>Oldest outstanding From COP Op data transfer</td> </tr> <tr> <td>001<sub>2</sub></td> <td>2nd oldest From COP Op data transfer.</td> </tr> <tr> <td>010<sub>2</sub></td> <td>3rd oldest From COP Op data transfer.</td> </tr> <tr> <td>011<sub>2</sub></td> <td>4th oldest From COP Op data transfer.</td> </tr> <tr> <td>100<sub>2</sub></td> <td>5th oldest From COP Op data transfer.</td> </tr> <tr> <td>101<sub>2</sub></td> <td>6th oldest From COP Op data transfer.</td> </tr> <tr> <td>110<sub>2</sub></td> <td>7th oldest From COP Op data transfer.</td> </tr> <tr> <td>111<sub>2</sub></td> <td>8th oldest From COP Op data transfer.</td> </tr> </tbody> </table> <p><b>Note:</b> Only values 000<sub>2</sub> and 001<sub>2</sub> are allowed see <i>CP2_fordlim_0</i>[2:0] below</p>	<i>CP2_forder_0</i> [2:0]	Order	000 <sub>2</sub>	Oldest outstanding From COP Op data transfer	001 <sub>2</sub>	2nd oldest From COP Op data transfer.	010 <sub>2</sub>	3rd oldest From COP Op data transfer.	011 <sub>2</sub>	4th oldest From COP Op data transfer.	100 <sub>2</sub>	5th oldest From COP Op data transfer.	101 <sub>2</sub>	6th oldest From COP Op data transfer.	110 <sub>2</sub>	7th oldest From COP Op data transfer.	111 <sub>2</sub>	8th oldest From COP Op data transfer.
<i>CP2_forder_0</i> [2:0]	Order																			
000 <sub>2</sub>	Oldest outstanding From COP Op data transfer																			
001 <sub>2</sub>	2nd oldest From COP Op data transfer.																			
010 <sub>2</sub>	3rd oldest From COP Op data transfer.																			
011 <sub>2</sub>	4th oldest From COP Op data transfer.																			
100 <sub>2</sub>	5th oldest From COP Op data transfer.																			
101 <sub>2</sub>	6th oldest From COP Op data transfer.																			
110 <sub>2</sub>	7th oldest From COP Op data transfer.																			
111 <sub>2</sub>	8th oldest From COP Op data transfer.																			

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description												
CP2_fordlim_0[2:0]	O	From Coprocessor Data Out-of-Order Limit. This signal sets the limit on how much the coprocessor can reorder From COP Data. The value on this signal corresponds to the maximum allowed value to be used on <i>CP2_forder_0[2:0]</i> . <b>Note:</b> The 4KEc core can handle one Out-of-Order From Data transfer. <i>CP2_fordlim_0[2:0]</i> is therefore tied to 001 <sub>2</sub> . The core will also never have more than two outstanding From COP instructions issued, which also automatically limits <i>CP2_forder_0[2:0]</i> to 001 <sub>2</sub> .												
CP2_fdata_0[31:0]	I	From Coprocessor Data. Data to be transferred from coprocessor. Valid when <i>CP2_fds_0</i> is asserted.												
Coprocessor Condition Code Check: These signals are used to report the result of a condition code check to the 4KEc core from the COP2 coprocessor. This is only used for BC2 instructions.														
CP2_cccs_0	I	Coprocessor Condition Code Check Strobe. Asserted when coprocessor condition code check bits are available on <i>CP2_ccc_0</i> .												
CP2_ccc_0	I	Coprocessor Conditions Code Check. Valid when <i>CP2_cccs_0</i> is asserted. When asserted, the branch instruction checking the condition code should take the branch. When deasserted, the branch instruction should not branch.												
Coprocessor Exceptions: These signals are used by the COP2 coprocessor to report exception for each instruction.														
CP2_excsc_0	I	Coprocessor Exception Strobe. Asserted when coprocessor exception signalling is available on <i>CP2_exc_0</i> and <i>CP2_exccode_0</i> .												
CP2_exc_0	I	Coprocessor Exception. When asserted, a Coprocessor exception is signaled on <i>CP2_exccode_0[4:0]</i> . Valid when <i>CP2_excsc_0</i> is asserted.												
CP2_exccode_0[4:0]	I	Coprocessor Exception Code. Valid when both <i>CP2_excsc_0</i> and <i>CP2_exc_0</i> are asserted. <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th><i>CP2_exccode_0[4:0]</i></th> <th>Exception</th> </tr> </thead> <tbody> <tr> <td>01010<sub>2</sub></td> <td>(RI) Reserved Instruction Exception</td> </tr> <tr> <td>10000<sub>2</sub></td> <td>(IS1) Available for Coprocessor specific Exception</td> </tr> <tr> <td>10001<sub>2</sub></td> <td>(IS1) Available for Coprocessor specific Exception</td> </tr> <tr> <td>10010<sub>2</sub></td> <td>C2E Exception</td> </tr> <tr> <td>All others</td> <td>Reserved</td> </tr> </tbody> </table>	<i>CP2_exccode_0[4:0]</i>	Exception	01010 <sub>2</sub>	(RI) Reserved Instruction Exception	10000 <sub>2</sub>	(IS1) Available for Coprocessor specific Exception	10001 <sub>2</sub>	(IS1) Available for Coprocessor specific Exception	10010 <sub>2</sub>	C2E Exception	All others	Reserved
<i>CP2_exccode_0[4:0]</i>	Exception													
01010 <sub>2</sub>	(RI) Reserved Instruction Exception													
10000 <sub>2</sub>	(IS1) Available for Coprocessor specific Exception													
10001 <sub>2</sub>	(IS1) Available for Coprocessor specific Exception													
10010 <sub>2</sub>	C2E Exception													
All others	Reserved													
Instruction Nullification: These signals are used by the 4KEc core to signal nullification of each instruction to the COP2 coprocessor.														
CP2_nulls_0	O	Coprocessor Null Strobe. Asserted when a nullification signal is available on <i>CP2_null_0</i> .												
CP2_null_0	O	Nullify Coprocessor Instruction. When deasserted, the 4KEc core is signalling that the instruction is not nullified. When asserted, the 4KEc core is signalling that the instruction is nullified, and no further transactions will take place for this instruction. Valid when <i>CP2_nulls_0</i> is asserted.												
Instruction Killing: These signals are used by the 4KEc core to signal killing of each instruction to the COP2 coprocessor.														
CP2_kills_0	O	Coprocessor Kill Strobe. Asserted when kill signalling is available on <i>CP2_kill_0[1:0]</i> .												

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description										
CP2_kill_0[1:0]	O	Kill Coprocessor Instruction. Valid when <i>CP2_kills_0</i> is asserted.  <table border="1"> <thead> <tr> <th>CP2_kill_0[1:0]</th> <th>Type of Kill</th> </tr> </thead> <tbody> <tr> <td>00<sub>2</sub></td> <td>Instruction is not killed and results can be committed.</td> </tr> <tr> <td>01<sub>2</sub></td> <td></td> </tr> <tr> <td>10<sub>2</sub></td> <td>Instruction is killed. (not due to <i>CP2_exc_0</i>)</td> </tr> <tr> <td>11<sub>2</sub></td> <td>Instruction is killed. (due to <i>CP2_exc_0</i>)</td> </tr> </tbody> </table> If an instruction is killed, no further transactions will take place on the interface for this instruction.	CP2_kill_0[1:0]	Type of Kill	00 <sub>2</sub>	Instruction is not killed and results can be committed.	01 <sub>2</sub>		10 <sub>2</sub>	Instruction is killed. (not due to <i>CP2_exc_0</i> )	11 <sub>2</sub>	Instruction is killed. (due to <i>CP2_exc_0</i> )
			CP2_kill_0[1:0]	Type of Kill								
			00 <sub>2</sub>	Instruction is not killed and results can be committed.								
			01 <sub>2</sub>									
			10 <sub>2</sub>	Instruction is killed. (not due to <i>CP2_exc_0</i> )								
11 <sub>2</sub>	Instruction is killed. (due to <i>CP2_exc_0</i> )											
Miscellaneous COP2 signals:												
CP2_reset	O	Coprocessor Reset. Asserted when a hard or soft reset is performed by the integer unit.										
CP2_present	S	COP2 Present. Must be asserted when COP2 hardware is connected to the Coprocessor 2 Interface.										
CP2_idle	I	Coprocessor Idle. Asserted when the coprocessor logic is idle. Enables the processor to go into sleep mode and shut down the clock. Valid only if <i>CP2_present</i> is asserted.										
<b>EJTAG Interface</b>												
TAP interface. These signals comprise the EJTAG Test Access Port. These signals will not be connected if the core does not implement the TAP controller.												
EJ_TRST_N	I	Active-low Test Reset Input (TRST*) for the EJTAG TAP. At power-up, the assertion of <i>EJ_TRST_N</i> causes the TAP controller to be reset.										
EJ_TCK	I	Test Clock Input (TCK) for the EJTAG TAP.										
EJ_TMS	I	Test Mode Select Input (TMS) for the EJTAG TAP.										
EJ_TDI	I	Test Data Input (TDI) for the EJTAG TAP.										
EJ_TDO	O	Test Data Output (TDO) for the EJTAG TAP.										
EJ_TDOzstate	O	Drive indication for the output of TDO for the EJTAG TAP at chip level: 1: The TDO output at chip level must be in Z-state 0: The TDO output at chip level must be driven to the value of <i>EJ_TDO</i>  IEEE Standard 1149.1-1990 defines TDO as a 3-stated signal. To avoid having a 3-state core output, the 4KEc core outputs this signal to drive an external 3-state buffer.										
<i>Debug Interrupt:</i>												
EJ_DINTsup	S	Value of DINTsup for the Implementation register. When high, this signal indicates that the EJTAG probe can use the DINT signal to interrupt the processor.										
EJ_DINT	I	Debug exception request when this signal is asserted in a CPU clock period after being deasserted in the previous CPU clock period. The request is cleared when debug mode is entered. Requests when in debug mode are ignored.										



Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description
<i>Debug Mode Indication:</i>		
EJ_DebugM	O	Asserted when the core is in Debug Mode. This can be used to bring the core out of a low power mode. In systems with multiple processor cores, this signal can be used to synchronize the cores when debugging.
<i>Device ID bits:</i>		
These inputs provide an identifying number visible to the EJTAG probe. If the EJTAG TAP controller is not implemented, these inputs are not connected. These inputs are always available for soft core customers. On hard cores, the core “hardener” can set these inputs to their own values.		
EJ_ManufID[10:0]	S	Value of the ManufID[10:0] field in the Device ID register. As per IEEE 1149.1-1990 section 11.2, the manufacturer identity code shall be a compressed form of JEDEC standard manufacturer’s identification code in the JEDEC Publications 106, which can be found at: <a href="http://www.jedec.org/">http://www.jedec.org/</a> ManufID[6:0] bits are derived from the last byte of the JEDEC code by discarding the parity bit. ManufID[10:7] bits provide a binary count of the number of bytes in the JEDEC code that contain the continuation character (0x7F). Where the number of continuations characters exceeds 15, these 4 bits contain the modulo-16 count of the number of continuation characters.
EJ_PartNumber[15:0]	S	Value of the PartNumber[15:0] field in the Device ID register.
EJ_Version[3:0]	S	Value of the Version[3:0] field in the Device ID register.
<i>System Implementation Dependent Outputs:</i>		
These signals come from EJTAG control registers. They have no effect on the core, but can be used to give EJTAG debugging software additional control over the system.		
EJ_SRstE	O	Soft Reset Enable. EJTAG can deassert this signal if it wants to mask soft resets. If this signal is deasserted, none, some, or all soft reset sources are masked.
EJ_PerRst	O	Peripheral Reset. EJTAG can assert this signal to request the reset of some or all of the peripheral devices in the system.
EJ_PrRst	O	Processor Reset. EJTAG can assert this signal to request that the core be reset. This can be fed into the <i>SI_Reset</i> signal.
<i>EJTAG Trace Interface</i>		
These signals enable an interface to optional off-chip trace memory. The EJTAG Trace interface connects to the Probe Interface Block (PIB) which in turn connects to the physical off-chip trace pins. Note that if on-chip trace memory is used, access occurs via the EJTAG TAP interface, and this interface is not required.		

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description																		
TC_ClockRatio[2:0]	O	<p>Clock ratio. This is the clock ratio set by software in <i>TCBCONTROLB.CR</i>. The value will be within the boundaries defined by <i>TC_CRMax</i> and <i>TC_CRMin</i>. The table below shows the encoded values for clock ratio.</p> <table border="1"> <thead> <tr> <th>TC_ClockRatio</th> <th>Clock Ratio</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>8:1 (Trace clock is eight times the core clock)</td> </tr> <tr> <td>001</td> <td>4:1 (Trace clock is four times the core clock)</td> </tr> <tr> <td>010</td> <td>2:1 (Trace clock is double the core clock)</td> </tr> <tr> <td>011</td> <td>1:1 (Trace clock is same as the core clock)</td> </tr> <tr> <td>100</td> <td>1:2 (Trace clock is one half the core clock)</td> </tr> <tr> <td>101</td> <td>1:4 (Trace clock is one fourth the core clock)</td> </tr> <tr> <td>110</td> <td>1:6 (Trace clock is one sixth the core clock)</td> </tr> <tr> <td>111</td> <td>1:8 (Trace clock is one eight the core clock)</td> </tr> </tbody> </table>	TC_ClockRatio	Clock Ratio	000	8:1 (Trace clock is eight times the core clock)	001	4:1 (Trace clock is four times the core clock)	010	2:1 (Trace clock is double the core clock)	011	1:1 (Trace clock is same as the core clock)	100	1:2 (Trace clock is one half the core clock)	101	1:4 (Trace clock is one fourth the core clock)	110	1:6 (Trace clock is one sixth the core clock)	111	1:8 (Trace clock is one eight the core clock)
TC_ClockRatio	Clock Ratio																			
000	8:1 (Trace clock is eight times the core clock)																			
001	4:1 (Trace clock is four times the core clock)																			
010	2:1 (Trace clock is double the core clock)																			
011	1:1 (Trace clock is same as the core clock)																			
100	1:2 (Trace clock is one half the core clock)																			
101	1:4 (Trace clock is one fourth the core clock)																			
110	1:6 (Trace clock is one sixth the core clock)																			
111	1:8 (Trace clock is one eight the core clock)																			
TC_CRMax[2:0]	S	Maximum clock ratio supported. This static input sets the CRMax field of the <i>TCBCONFIG</i> register. It defines the capabilities of the Probe Interface Block (PIB) module. This field determines the minimum value of <i>TC_ClockRatio</i> .																		
TC_CRMin[2:0]	S	Minimum clock ratio supported. This input sets the CRMin field of the <i>TCBCONFIG</i> register. It defines the capabilities of the PIB module. This field determines the maximum value of <i>TC_ClockRatio</i> .																		
TC_ProbeWidth[1:0]	S	<p>This static input will set the PW field of the <i>TCBCONFIG</i> register.</p> <p>If this interface is not driving a PIB module, but some chip-level TCB-like module, then this field should be set to 2'b11 (reserved value for PW).</p> <table border="1"> <thead> <tr> <th>TC_ProbeWidth</th> <th>Number physical data pin on PIB</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>4 bits</td> </tr> <tr> <td>01</td> <td>8 bits</td> </tr> <tr> <td>10</td> <td>16 bits</td> </tr> <tr> <td>11</td> <td>Not directly to PIB</td> </tr> </tbody> </table>	TC_ProbeWidth	Number physical data pin on PIB	00	4 bits	01	8 bits	10	16 bits	11	Not directly to PIB								
TC_ProbeWidth	Number physical data pin on PIB																			
00	4 bits																			
01	8 bits																			
10	16 bits																			
11	Not directly to PIB																			
TC_PibPresent	S	Must be asserted when a PIB is attached to the TC Interface. When de-asserted (low) all the other inputs are disregarded.																		
TC_TrEnable	O	Trace Enable, when asserted the PIB must start running its output clock and can expect valid data on all other outputs.																		
TC_Calibrate	O	<p>This signal is asserted when the Cal bit in the <i>TCBCONTROLB</i> register is set.</p> <p>For a simple PIB which only serves one TCB, this pin can be ignored. For a multi-core capable PIB which also uses <i>TC_Valid</i> and <i>TC_Stall</i>, the PIB must start producing the calibration pattern when this signal is asserted.</p>																		

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description														
TC_DataBits[2:0]	I	<p>This input identifies the number of bits picked up by the probe interface module in each “cycle”.</p> <p>If <i>TC_ClockRatio</i> indicates a clock-ratio higher than 1:2, then clock multiplication in the Probe logic is used. The “cycle” is equal to each core clock cycle.</p> <p>If <i>TC_ClockRatio</i> indicates a clock-ratio lower than or equal to 1:2, then “cycle” is (clock-ratio * 2) of the core clock cycle. For example, with a clock ratio of 1:2, a “cycle” is equal to core clock cycle; with a clock ratio of 1:4, a “cycle” is equal to one half of core clock cycle.</p> <p>This input controls the down-shifting amount and frequency of the trace word on <i>TC_Data[63:0]</i>. The bit width and the corresponding <i>TC_DataBits</i> value is shown in the table below.</p> <table border="1" data-bbox="750 659 1208 961"> <thead> <tr> <th><i>TC_DataBits[2:0]</i></th> <th>Probe uses following bits from <i>TC_Data</i> each cycle</th> </tr> </thead> <tbody> <tr> <td>000</td> <td><i>TC_Data[3:0]</i></td> </tr> <tr> <td>001</td> <td><i>TC_Data[7:0]</i></td> </tr> <tr> <td>010</td> <td><i>TC_Data[15:0]</i></td> </tr> <tr> <td>011</td> <td><i>TC_Data[31:0]</i></td> </tr> <tr> <td>100</td> <td><i>TC_Data[63:0]</i></td> </tr> <tr> <td>Others</td> <td>Unused</td> </tr> </tbody> </table> <p>This input might change as the value on <i>TC_ClockRatio[2:0]</i> changes.</p>	<i>TC_DataBits[2:0]</i>	Probe uses following bits from <i>TC_Data</i> each cycle	000	<i>TC_Data[3:0]</i>	001	<i>TC_Data[7:0]</i>	010	<i>TC_Data[15:0]</i>	011	<i>TC_Data[31:0]</i>	100	<i>TC_Data[63:0]</i>	Others	Unused
<i>TC_DataBits[2:0]</i>	Probe uses following bits from <i>TC_Data</i> each cycle															
000	<i>TC_Data[3:0]</i>															
001	<i>TC_Data[7:0]</i>															
010	<i>TC_Data[15:0]</i>															
011	<i>TC_Data[31:0]</i>															
100	<i>TC_Data[63:0]</i>															
Others	Unused															
TC_Valid	O	Asserted when a valid new trace word is started on the <i>TC_Data[63:0]</i> signals. <i>TC_Valid</i> is only asserted when <i>TC_DataBits</i> is 100.														
TC_Stall	I	When asserted, a new <i>TC_Valid</i> in the following cycle is stalled. <i>TC_Valid</i> is still asserted, but the <i>TC_Data</i> value and <i>TC_Valid</i> are held static, until the cycle after <i>TC_Stall</i> is sampled low. <i>TC_Stall</i> is only sampled in the cycle before a new <i>TC_Valid</i> cycle, and only when <i>TC_DataBits</i> is 100, indicating a full word of <i>TC_Data</i> .														
TC_Data[63:0]	O	Trace word data. The value on this 64-bit interface is shifted down as indicated in <i>TC_DataBits[2:0]</i> . In the first cycle where a new trace word is valid on all the bits and <i>TC_DataBits[2:0]</i> is 100, <i>TC_Valid</i> is also asserted. The Probe Interface Block (PIB) will only be connected to [(N-1):0] bits of this output bus. N is the number of bits picked up by the PIB in each core clock cycle. For clock ratios 1:2 and lower, N is equal to the number of physical trace pins (legal values of N are 4, 8, or 16). For higher clock ratios, N is larger than the number of physical trace pins.														
TC_ProbeTrigIn	A	Rising edge trigger input. The source should be the Probe Trigger input. The input is considered asynchronous; i.e., it is double registered in the core.														
TC_ProbeTrigOut	O	Single cycle (relative to the “cycle” defined the description of <i>TC_DataBits</i> ) high strobe, trigger output. The target of this trigger is intended to be the external probe’s trigger output.														
TC_ChipTrigIn	A	Rising edge trigger input. The source should be on-chip. The input is considered asynchronous; i.e., it is double registered in the core.														
TC_ChipTrigOut	O	Single cycle (relative to core clock) high strobe, trigger output. The target of this trigger is intended to be an on-chip unit.														
<b>Performance Monitoring Interface</b>																

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description
These signals can be used to implement performance counters, which can be used to monitor hardware/software performance.		
PM_DCacheHit	O	This signal is asserted whenever there is a data cache hit.
PM_DCacheMiss	O	This signal is asserted whenever there is a data-cache miss.
PM_DTLBHit	O	This signal is asserted whenever there is a hit in the data TLB.
PM_DTLBMiss	O	This signal is asserted whenever there is a miss in the data TLB.
PM_ICacheHit	O	This signal is asserted whenever there is an instruction-cache hit.
PM_ICacheMiss	O	This signal is asserted whenever there is an instruction-cache miss.
PM_InstComplete	O	This signal is asserted each time an instruction completes in the pipeline.
PM_ITLBHit	O	This signal is asserted whenever there is an instruction TLB hit.
PM_ITLBMiss	O	This signal is asserted whenever there is an instruction TLB miss.
PM_JTLBHit	O	This signal is asserted whenever there is a joint TLB hit.
PM_JTLBMiss	O	This signal is asserted whenever there is a joint TLB miss.
PM_WTBMerge	O	This signal is asserted whenever there is a successful merge in the write-through buffer.
PM_WTBNoMerge	O	This signal is asserted whenever a non-merging store is written to the write-through buffer.
<b>ScratchPad RAM interface</b>		
This interface allows a ScratchPad RAM (SPRAM) array to be connected in parallel with the cache arrays, enabling fast access to data. There are independent interfaces for Instruction and Data ScratchPads. Signals related to the Instruction Scratchpad interface are prefixed with "ISP_". Signals related to the Data Scratchpad interface are prefixed with "DSP_". Note: In order to achieve single cycle access, the ScratchPad interface is not registered, unlike the other core interfaces. This requires more careful timing considerations.		
DSP_TagAddr[19:4]	O	Virtual index into the SPRAM used for tag reads and writes.
DSP_TagRdStr	O	Tag Read Strobe - Hit, Stall, TagRdValue use this strobe.
DSP_TagWrStr	O	Tag Write Strobe - If SPRAM tag is software configurable, this signal will indicate when to update the tag value.
DSP_TagCmpValue[23:0]	O	Tag Compare Value - This bus is used for both tag comparison and tag write value. For tag comparison, the bus usage is {PA[31:10], 2'b0} and contains the address to determine hit/miss. For tag writes, the bus contains {PA[31:10], Lock, Valid} from the <i>TagLo</i> register.
DSP_DataAddr[19:2]	O	Virtual index into the SPRAM used for data reads and writes.
DSP_DataWrValue[31:0]	O	Data Write Value - Data value to be written to the data array.
DSP_DataRdStr	O	Data Read Strobe - Indicates that the data array should be read.
DSP_DataWrStr	O	Data Write Strobe - Indicates that the data array should be written.
DSP_DataWrMask[3:0]	O	Data Write Mask - Byte enables for a data write.
DSP_DataRdValue[31:0]	I	Data Read Value - Data value read from the data array.
DSP_TagRdValue[23:0]	I	Tag Read Value - Tag value read from the tag array. Written to <i>TagLo</i> register on a CACHE instruction. Read value maps into these <i>TagLo</i> fields: {PA[31:10], Lock, Valid}

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description
DSP_Hit	I	Hit - Indicates that this read was to an address covered by the SPRAM.
DSP_Stall	I	Stall - Indicates that the read has not yet completed.
DSP_Present	S	Present - Indicates that a SPRAM array is connected to this port.
ISP_Addr[19:2]	O	Virtual index into the SPRAM used for both reads and writes of tag and data.
ISP_RdStr	O	Read Strobe - indicates a read of the tag and data arrays. Hit and Stall signals are also based off of this strobe.
ISP_TagWrStr	O	Tag Write Strobe - If SPRAM tag is software configurable, this signal will indicate when to update the tag value.
ISP_DataTagValue[31:0]	O	Write/Compare Data For data writes, this is the value to be written to the data array. For tag writes the bus contains the {8'b0, PA[31:10], Lock, Valid} from the TagLo register. For tag comparison, the bus has the address to be used for hit/miss determination in the format {8'b0, PA[31:10], Uncacheable, 1'b0}. When high, the Uncacheable bit indicates that the physical address bits (PA[31:10]) are to an uncacheable address; when the Uncacheable bit is low, the physical address is to a cacheable address.
ISP_DataWrStr	O	Data Write Strobe - Indicates that the data array should be written.
ISP_DataRdValue[31:0]	I	Data Read Value - Data value read from the data array.
ISP_TagRdValue[23:0]	I	Tag Read Value - Tag value read from the tag array. Written to <i>TagLo</i> register on a CACHE instruction. Read value maps into these <i>TagLo</i> fields: {PA[31:10], Lock, Valid}
ISP_Hit	I	Hit - Indicates that this read was to an address covered by the SPRAM.
ISP_Stall	I	Stall - Indicates that the read has not yet completed.
ISP_Present	S	Present - Indicates that a SPRAM array is connected to this port.
<b>Integrated Memory BIST Interface</b>		
These signals provide the interface to optional integrated memory BIST capability for testing the SRAM arrays within the core.		
gmbinvoke	I	Enable signal for integrated BIST controllers.
gmbdone	O	Common completion indicator for all integrated BIST sequences.
gmbddfai	O	When high, indicates that the integrated BIST test failed on the data cache data array.
gmbtdfai	O	When high, indicates that the integrated BIST test failed on the data cache tag array.
gmbwdfai	O	When high, indicates that the integrated BIST test failed on the data cache way select array.
gmbdifai	O	When high, indicates that the integrated BIST test failed on the instruction cache data array.
gmbtifai	O	When high, indicates that the integrated BIST test failed on the instruction cache tag array.
gmbwifai	O	When high, indicates that the integrated BIST test failed on the instruction cache way select array.
<b>Scan Test Interface</b>		
These signals provide an interface for testing the core. The use and configuration of these pins are implementation-dependent.		

Table 12 4KEc Signal Descriptions (Continued)

Signal Name	Type	Description
gscanenable	I	This signal should be asserted while scanning vectors into or out of the core. The <i>gscanenable</i> signal must be deasserted during normal operation and during capture clocks in test mode.
gscanmode	I	This signal should be asserted during all scan testing both while scanning and during capture clocks. The <i>gscanmode</i> signal must be deasserted during normal operation.
gscanramwr	I	This signal controls the read and write strobes to the cache SRAM when <i>gscanmode</i> is asserted.
gscanin_X	I	These signal(s) are the inputs to the scan chain(s).
gscanout_X	O	These signal(s) are the outputs from the scan chain(s).
BistIn[n:0]	I	Input to user-specified BIST controller.
BistOut[n:0]	O	Output from user-specified BIST controller.

## EC Interface Transactions

The 4KEc core implements the EC™ interface for its bus transactions. This interface uses a pipelined, in-order protocol with independent address, read data, and write data buses. The following subsections describe the four basic bus transactions: single read, single write, burst read, and burst write.

### Single Read

Figure 8 shows the basic timing relationships of signals during a simple read transaction. During a single read cycle, the 4KEc core drives the address onto *EB\_A[35:2]* and byte enable information onto *EB\_BE[3:0]*. To maximize performance, the EC interface does not define a maximum number of outstanding bus cycles. Instead it provides the *EB\_ARdy* input signal. This signal is driven by external logic and controls the generation of addresses on the bus.

In the 4KEc core, the address is driven whenever it becomes available, regardless of the state of *EB\_ARdy*. However, the 4KEc core always continues to drive the address until the clock after *EB\_ARdy* is sampled asserted. For example, at the rising edge of the clock 2 in Figure 8, the *EB\_ARdy* signal is sampled low, indicating that external logic is not ready to accept the new address. However, the 4KEc core still drives *EB\_A[35:2]* in this clock as shown. On the rising edge of clock 3, the 4KEc core samples *EB\_ARdy* asserted and continues to drive the address until the rising edge of clock 4.

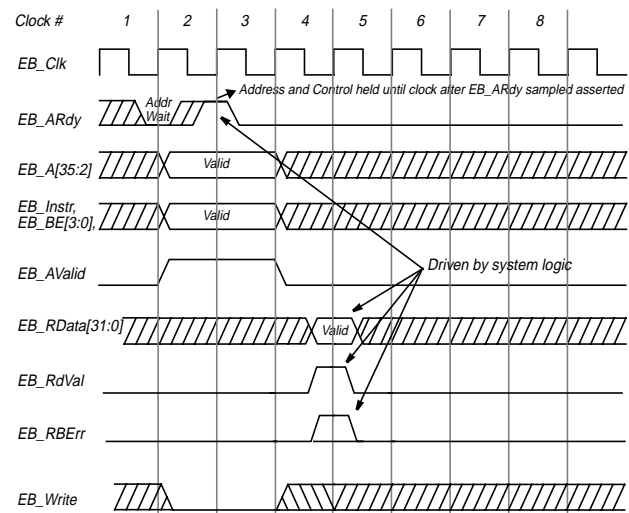


Figure 8 Single Read Transaction Timing Diagram

The *EB\_Instr* signal is only asserted during a single read cycle if there is an instruction fetch from non-cacheable memory space. The *EB\_AValid* signal is driven in each clock that *EB\_A[35:2]* is valid on the bus. The 4KEc core drives *EB\_Write* low to indicate a read transaction.

The *EB\_RData[31:0]* and *EB\_RdVal* signals are first sampled on the rising edge of clock 4, one clock after *EB\_ARdy* is sampled asserted. Data is sampled on every clock thereafter until *EB\_RdVal* is sampled asserted.

If a bus error occurs during the data transaction, external logic asserts *EB\_RBErr* in the same clock as *EB\_RdVal*.

## Single Write

Figure 9 shows a typical write transaction. The 4KEc core drives address and control information onto the  $EB\_A[35:2]$  and  $EB\_BE[3:0]$  signals on the rising edge of clock 2. As in the single read cycle, these signals remain active until the clock edge after the  $EB\_ARdy$  signal is sampled asserted. The 4KEc core asserts the  $EB\_Write$  signal to indicate that a valid write cycle is on the bus and  $EB\_AValid$  to indicate that valid address is on the bus.

The 4KEc core drives write data onto  $EB\_WData[31:0]$  in the same clock as the address and continues to drive data until the clock edge after the  $EB\_WDRdy$  signal is sampled asserted. If a bus error occurs during a write operation, external logic asserts the  $EB\_WBErr$  signal one clock after asserting  $EB\_WDRdy$ .




Figure 9 Single Write Transaction Timing Diagram

## Burst Read

The 4KEc core is capable of generating burst transactions on the bus. A burst transaction is used to transfer multiple data items in one transaction.

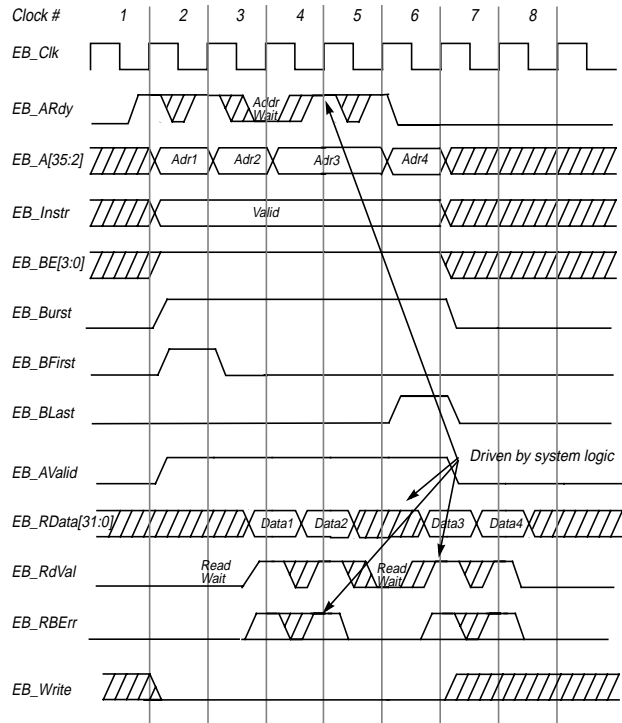


Figure 10 Burst Read Transaction Timing Diagram

Figure 10 shows an example of a burst read transaction. Burst read transactions initiated by the 4KEc core always contain four data transfers in a sequence determined by the critical word (the address that caused the miss) and  $EB\_SBlock$ . In addition, the data requested is always a 16-byte aligned block.

The order of words within this 16-byte block varies depending on which of the words in the block is being requested by the execution unit and the ordering protocol selected. The burst always starts with the word requested by the execution unit and proceeds in either an ascending or descending address order, wrapping when the block boundary is reached. Table 13 and Table 14 show the sequence of address bits 2 and 3.

Table 13 Sequential Ordering Protocols

Starting Address $EB\_A[3:2]$	Address Progression of $EB\_A[3:2]$
00	00, 01, 10, 11
01	01, 10, 11, 00
10	10, 11, 00, 01
11	11, 00, 01, 10

Table 14 Sub-Block Ordering Protocols

Starting Address EB_A[3:2]	Address Progression of EB_A[3:2]
00	00, 01, 10, 11
01	01, 00, 11, 10
10	10, 11, 00, 01
11	11, 10, 01, 00

The 4KEc core drives address and control information onto the *EB\_A[35:2]* and *EB\_BE[3:0]* signals on the rising edge of clock 2. As in the single read cycle, these signals remain active until the clock edge after the *EB\_ARdy* signal is sampled asserted. The 4KEc core continues to drive *EB\_AValid* as long as a valid address is on the bus.

The *EB\_Instr* signal is asserted if the burst read is for an instruction fetch. The *EB\_Burst* signal is asserted while the address is on the bus to indicate that the current address is part of a burst transaction. The 4KEc core asserts the *EB\_BFirst* signal in the same clock as the first address is driven and the *EB\_BLast* signal in the same clock as the last address to indicate the start and end of a burst cycle.

The 4KEc core first samples the *EB\_RData[31:0]* signals two clocks after *EB\_ARdy* is sampled asserted. External logic asserts *EB\_RdVal* to indicate that valid data is on the bus. The 4KEc core latches data internally whenever *EB\_RdVal* is sampled asserted.

Note that on the rising edge of clocks 3 and 6 in Figure 10, the *EB\_RdVal* signal is sampled deasserted, causing wait states in the data return. There is also an address wait state caused by *EB\_ARdy* being sampled deasserted on the rising edge of clock 4. Note that the core holds address 3 on the *EB\_A* bus for an extra clock because of this wait state. External logic asserts the *EB\_RBErr* signal in the same clock as data if a bus error occurs during that data transfer.

### Burst Write

Burst write transactions are used to empty one of the write buffers. A burst write transaction is only performed if the write buffer contains 16 bytes of data associated with the same aligned memory block, otherwise individual write transactions are performed. Figure 11 shows a timing diagram of a burst write transaction. Unlike the read burst, a write burst always begins with *EB\_A[3:2]* equal to 00b.

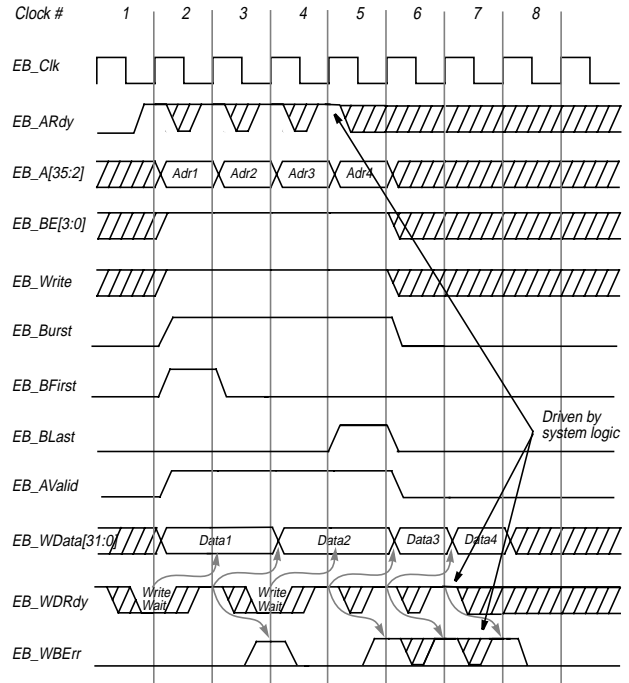


Figure 11 Burst Write Transaction Timing Diagram

The 4KEc core drives address and control information onto the *EB\_A[35:2]* and *EB\_BE[3:0]* signals on the rising edge of clock 2. As in the single read cycle, these signals remain active until the clock edge after the *EB\_ARdy* signal is sampled asserted. The 4KEc core continues to drive *EB\_AValid* as long as a valid address is on the bus.

The 4KEc core asserts the *EB\_Write*, *EB\_Burst*, and *EB\_AValid* signals during the time the address is driven. *EB\_Write* indicates that a write operation is in progress. The assertion of *EB\_Burst* indicates that the current operation is a burst. *EB\_AValid* indicates that valid address is on the bus.

The 4KEc core asserts the *EB\_BFirst* signal in the same clock as address 1 is driven to indicate the start of a burst cycle. In the clock that the last address is driven, the 4KEc core asserts *EB\_BLast* to indicate the end of the burst transaction.

In Figure 11, the first data word (Data1) is driven in clocks 2 and 3 due to the *EB\_WDRdy* signal being sampled deasserted at the rising edge of clock 2, causing a wait state. When *EB\_WDRdy* is sampled asserted on the rising edge of clock 3, the 4KEc core responds by driving the second word (Data2).



External logic drives the *EB\_WBErr* signal one clock after the corresponding assertion of *EB\_WDRdy* if a bus error has occurred as shown by the arrows in Figure 11.

## Revision History

In the left hand page margins of this document you may find vertical change bars to note the location of significant changes to this document since its last release. Significant changes are defined as those which you should take note of as you use the MIPS IP. Changes to correct grammar, spelling errors or similar may or may not be noted with change bars. Change bars will be removed for changes which are more than one revision old.

Please note: Limitations on the authoring tools make it difficult to place change bars on changes to figures. Change bars on figure titles are used to denote a potential change in the figure itself. Certain parts of this document (Instruction set descriptions, EJTAG register definitions) are references to Architecture specifications, and the change bars within these sections indicate alterations since the previous version of the relevant Architecture document.

Revision	Date	Description
02.00	November 8, 2002	<ul style="list-style-type: none"><li>• Added this revision history table.</li><li>• Various updates to describe new MIPS32 Release 2 capabilities, included in version 3.0 or higher core releases.</li></ul>

Copyright © 2001-2002 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 or transferred for the purpose of reexporting in violation of any U.S. or non-U.S. regulation, treaty, Executive Order, law, statute, amendment or supplement thereto.

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, R3000, R4000, R5000 and R10000 are among the registered trademarks of MIPS Technologies, Inc. in the United States and other countries, and MIPS16, MIPS16e, MIPS32, MIPS64, MIPS-3D, MIPS-based, MIPS I, MIPS II, MIPS III, MIPS IV, MIPS V, MIPSsim, SmartMIPS, MIPS Technologies logo, 4K, 4Kc, 4Km, 4Kp, 4KE, 4KEc, 4KEm, 4KEp, 4KS, 4KSc, 4KSd, M4K, 5K, 5Kc, 5Kf, 20Kc, 25Kf, ASMACRO, ATLAS, At the Core of the User Experience., BusBridge, CoreFPGA, CoreLV, EC, JALGO, MALTA, MDMX, MGB, PDtrace, Pipeline, Pro, Pro Series, SEAD, SEAD-2, SOC-it and YAMON are among the trademarks of MIPS Technologies, Inc.

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

Template: D1.06, Build with Conditional Tags: 2B