MIPS R10000 Microprocessor User's Manual

Version 1.0

Copyright © 1995 MIPS Technologies, Inc.

ALL RIGHTS RESERVED

U.S. GOVERNMENT RESTRICTED RIGHTS LEGEND

Use, duplication or disclosure by the Government is subject to restrictions as set forth in FAR 52.227.19(c)(2) or subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and/or in similar or successor clauses in the FAR, or the DOD or NASA FAR Supplement. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94039-7311.

RISCompiler, RISC/os, R2000, R6000, R4000, R4400, and R10000 are trademarks of MIPS Technologies, Inc. MIPS and R3000 are registered trademarks of MIPS Technologies, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

MIPS Technologies, Inc.

2011 North Shoreline

Mountain View, California 94039-7311

http://www.mips.com

Acknowledgments

This book represents a consortium of efforts, and is principally derived from material provided by **Randy Martin**, **Yung-Chin Chen**, and **Ken Yeager**.

Thanks also to Randy for his many painstaking reviews of this manual.

Also providing invaluable service were the following:

Shabbir Latif, for once again running point between Engineering and Publications, answering questions, and presenting tutorials to clarify the complicated details of the R10000 processor operations.

Charlie Price, for use of his rejuvenated *MIPS-4 Instruction Set Architecture.*

Steve Proffitt, for both his technical assistance, and helping handle the multitude of niggling details involved in getting this manual printed.

The following also provided technical help in innumerable ways: **Arun Mehta**, **Tim Layman**, **Greg Shippen**, **Yeffi Van Atta**, **John Brennan**, **Len Widra**, **Roy Johnson**, **Hector Sucar**, **Hong-Men Su**, **Mazin Khurshid**, **Steve Whitney**, **Doug Yanagawa** (chip illustrations and socket pinouts), **Mike Gupta**, **Steven Peltier**, **Rob Conrad**, **Hai Nguyen**, **Bill Voegtli**, and **Sharad Mehrotra** at the University of Illinois.

Remediating a prior deficiency, thanks to **Tom McReynolds**.

In Production and Creative, thanks to **Melissa Miller** for her design of the cover; **Yen Nguyen**, for handling the printing; both **Kay Maitz** and **Beth Fraker** for resolving various design issues; and **Michael Ritchie** for tracking progress.

Joe Heinrich
June, 1995
Mt. View, California

About This Manual

This manual describes the MIPS R10000 RISC microprocessor (also referred to as the *processor* in this book).

Glossary

Certain specialized terms used in this book are defined in the *Glossary* at the end of this manual.

Stylistic Conventions

A brief note on some of the stylistic conventions used in this book: bits, fields, and registers of interest from a software perspective are italicized (such as the *BE* bit in the *Config* register).

Signal names of more importance from a hardware point of view are rendered in bold (such as **Reset***). The **asterisk** appended to the signal name (as in **Reset***) indicates the signal is low-active.

A **range of bits** uses a colon as a separator; for instance, **(15:0)** represents the 16-bit range that runs from bit 0, inclusive, through bit 15. In some places an ellipsis **(15...0)** or partial ellipsis **(15..0)** may used in place of a colon for visibility**.**

Unfamiliar terms presented for the first time are printed in **bold letters**, and are followed as closely as possible by a definition or description.

Getting MIPS Documents On-Line

The information in this manual, and other MIPS-related product information, is also available over the Word Wide Web at:

http://www.mips.com/

Requests can also be e-mailed to **webmaster@mti.sgi.com**.

MIPS documents are also available for retrieval on-line, through the file transport protocol (FTP). To retrieve them, follow the steps below. The text you are to type is shown in **`Courier Bold`** font; the computer's responses are in shown in `Courier Regular` font.

1. First, place yourself in the directory on your system within which you want to store the retrieved files. Do this by typing:

   `cd <directory_you_want_file_to_be_in>`

2. Access the MIPS document server, **sgigate**, through FTP by typing:

   ```
   ftp sgigate.sgi.com
   ```

3. The server tells you when you are connected for FTP by responding:

   ```
   Connected to sgigate.sgi.com.
   ```

4. Next (after some announcements) the server asks you to log in by requesting a name and then a password.

   ```
   Name (sgigate.sgi.com:<login_name>):
   ```

5. Login by typing **anonymous** for your name and your **electronic mail address** for your password.

   ```
   Name (sgigate.sgi.com:<login_name>): anonymous
   331 Guest login ok, type your name as password.
   Password: your_email_address
   ```

6. The system indicates you have logged in by supplying an FTP prompt:

   ```
   ftp>
   ```

7. Go to the pub/doc directory by typing:

   ```
   ftp> cd pub/doc
   ```

8. You can take a look at the contents of the **doc** directory by listing them:

   ```
   ftp> ls
   ```

9. You will find several subdirectories, such as R4200, R8000, and R10000. When you find the subdirectory you want, **cd** into that subdirectory and retrieve the file you want by typing:

   ```
   get <filename>
   ```

10. This copies the file from **sgigate** back to your system.
11. When you have retrieved all the files you want, exit from **ftp** by typing **quit**:

    ```
    ftp> quit
    ```

12. If the file was encoded for transmission, you must decode it, after retrieval, by typing:

    ```
    uudecode <filename>
    ```

13. If the file was compressed for transmission, you must uncompress it, after retrieval, by typing:

    ```
    uncompress <filename>
    ```

14. If you tarred the file, type:

    ```
    tar xvof <filename>
    ```

Contents

About This Manual

Glossary v
Stylistic Conventions v
Getting MIPS Documents On-Line v

1 Introduction to the R10000 Processor

MIPS Instruction Set Architecture (ISA) 2
What is a Superscalar Processor? 3
Pipeline and Superpipeline Architecture 3
Superscalar Architecture 3
What is an R10000 Microprocessor? 4
R10000 Superscalar Pipeline 5
Instruction Queues 6
Execution Pipelines 6
64-bit Integer ALU Pipeline 6
Load/Store Pipeline 7
64-bit Floating-Point Pipeline 7
Functional Units 9
Primary Instruction Cache (I-cache) 9
Primary Data Cache (D-cache) 9
Instruction Decode And Rename Unit 10
Branch Unit 10
External Interfaces 10
Instruction Queues 11
Integer Queue 11
Floating-Point Queue 11
Address Queue 12
Program Order and Dependencies 13
Instruction Dependencies 13
Execution Order and Stalling 13
Branch Prediction and Speculative Execution 14
Resolving Operand Dependencies 14
Resolving Exception Dependencies 14
Strong Ordering 15
An Example of Strong Ordering 16
R10000 Pipelines 17
Stage 1 17
Stage 2 17
Stage 3 18
Stages 4-6 18
Floating-Point Multiplier (3-stage Pipeline) 18
Floating-Point Divide and Square-Root Units 18
Floating-Point Adder (3-stage Pipeline) 18
Integer ALU1 (1-stage Pipeline) 18
Integer ALU2 (1-stage Pipeline) 18
Address Calculation and Translation in the TLB 19

Implications of R10000 Microarchitecture on Software....20
Superscalar Instruction Issue....20
Speculative Execution....21
Nonblocking Caches....21
R10000-Specific CPU Instructions....22
PREF....22
LL/SC....23
SYNC....24
Performance....24
User Instruction Latency and Repeat Rate....25
Other Performance Issues....27
Cache Performance....27

2
System Configurations

Uniprocessor Systems....30
Multiprocessor Systems....31
Multiprocessor Systems Using Dedicated External Agents....31
Multiprocessor Systems Using a Cluster Bus....32

3
Interface Signal Descriptions

Power Interface Signals....34
Secondary Cache Interface Signals....35
System Interface Signals....36
Test Interface Signals....38

4 Cache Organization and Coherency

Primary Instruction Cache 40
Primary Data Cache 42
Secondary Cache....45
Cache Algorithms....47
Descriptions of the Cache Algorithms 48
Uncached 48
Cacheable Noncoherent 48
Cacheable Coherent Exclusive 48
Cacheable Coherent Exclusive on Write....48
Uncached Accelerated 48
Relationship Between Cached and Uncached Operations 49
Cache Algorithms and Processor Requests 50
Cache Block Ownership 51

5 Secondary Cache Interface

Tag and Data Arrays....54
Secondary Cache Interface Frequencies....54
Secondary Cache Indexing....55
Indexing the Data Array....55
Indexing the Tag Array 56
Secondary Cache Way Prediction Table 57
Secondary Cache Tag....59
SCTag(25:4), Physical Tag 59
SCTag(3:2), PIdx 59
SCTag(1:0), Cache Block State 60
Read Sequences....60
4-Word Read Sequence....61
8-Word Read Sequence....62
16 or 32-Word Read Sequence....63
Tag Read Sequence....64
Write Sequences....65
4-Word Write Sequence....65
8-Word Write Sequence....66
16 or 32-Word Write Sequence....67
Tag Write Sequence....68

6 System Interface Operations

Request and Response Cycles 70
System Interface Frequencies 70
Register-to-Register Operation....70
System Interface Signals....71
Master and Slave States....71
Connecting to an External Agent....71
Cluster Bus 72
System Interface Connections 73
Uniprocessor System 73
Multiprocessor System Using Dedicated External Agents 74
Multiprocessor System Using the Cluster Bus....75
System Interface Requests and Responses 76
Processor Requests....76
External Responses 77
External Requests....77
Processor Responses 77
Outstanding Requests and Request Numbers....77
Request and Response Relationship....78
System Interface Buffers....79
Incoming Buffer....79
Cached Request Buffer 80
Outgoing Buffer....80
Uncached Buffer....81
Cluster Request Buffer....82
System Interface Flow Control....83
Processor Write and Eliminate Request Flow Control 83
Processor Read and Upgrade Request Flow Control....83
Processor Coherency Data Response Flow Control....83
External Request Flow Control 83
External Data Response Flow Control 83
System Interface Block Data Ordering....84
External Block Data Responses 84
Processor Coherency Data Responses....84
Processor Block Write Requests....84
System Interface Bus Encoding....85
SysCmd[11:0] Encoding 85
SysCmd[11] Encoding 85
SysCmd[10:0] Address Cycle Encoding 85
SysCmd[10:0] Data Cycle Encoding....89
SysCmd[11:0] Map....91

SysAD[63:0] Encoding 92
SysAD[63:0] Address Cycle Encoding 92
SysAD[63:60] 92
SysAD[59:58] Uncached Attribute 92
SysAD[57] 93
SysAD[56:40] 93
SysAD[39:0] 93
SysAD[63:0] Data Cycle Encoding 94
SysState[2:0] Encoding 94
SysResp[4:0] Encoding 95
Interrupts 95
Hardware Interrupts 95
Software Interrupts 96
Timer Interrupt 96
Nonmaskable Interrupt 96
Protocol Abbreviations 97
System Interface Arbitration 98
System Interface Arbitration Rules 99
Uniprocessor System 100
Multiprocessor System Using Cluster Bus 101
System Interface Request and Response Protocol 102
Processor Request Protocol 102
Processor Block Read Request Protocol 102
Processor Double/Single/Partial-Word Read Request Protocol 104
Processor Block Write Request Protocol 106
Processor Double/Single/Partial-Word Write Request Protocol 108
Processor Upgrade Request Protocol 110
Processor Eliminate Request Protocol 112
Processor Request Flow Control Protocol 114
External Response Protocol 116
External Block Data Response Protocol 116
External Double/Single/Partial-Word Data Response Protocol 118
External Completion Response Protocol 119
External Request Protocol 121
External Intervention Request Protocol 121
External Allocate Request Number Request Protocol 123
External Invalidate Request Protocol 124
External Interrupt Request Protocol 125
Processor Response Protocol 126
Processor Coherency State Response Protocol 126
Processor Coherency Data Response Protocol 127

System Interface Coherency 129
External Intervention Shared Request 129
External Intervention Exclusive Request 129
External Invalidate Request 129
External Coherency Request Action 130
Coherency Conflicts 131
Internal Coherency Conflicts 131
External Coherency Conflicts 132
External Coherency Request Latency 133
SysGblPerf* Signal 135
Cluster Bus Operation 136
Support for I/O 139
Support for External Duplicate Tags 139
Support for a Directory-Based Coherency Protocol 140
Support for Uncached Attribute 140
Support for Hardware Emulation 141

7
Clock Signals

System Interface Clock and Internal Processor Clock Domains 144
Secondary Cache Clock 144
Phase-Locked-Loop 145

8
Initialization

Initialization of Logical Registers 148
Power-On Reset Sequence 148
Cold Reset Sequence 150
Soft Reset Sequence 152
Mode Bits 153

9

Error Protection and Handling

Correctable Errors 158
Uncorrectable Errors 158
Propagation of Uncorrectable Errors 159
Cache Error Exception 160
CP0 CacheErr Register EW Bit 161
CP0 Status Register DE Bit 161
CACHE Instruction 161
Error Protection Schemes Used by R10000 162
- Parity 162
- Sparse Encoding 162
- ECC 162

Primary Instruction Cache Error Protection and Handling 163
- Error Protection 163
- Error Handling 163

Primary Data Cache Error Protection and Handling 164
- Error Protection 164
- Error Handling 164

Secondary Cache Error Protection and Handling 165
- Error Protection 165
- Error Handling 165
  - Data Array 165
    - Data Array in Correction Mode 165
    - Data Array in Noncorrection Mode 167
  - Tag Array 167

System Interface Error Protection and Handling 168
- Error Protection 168
- Error Handling 168
  - SysCmd(11:0) Bus 168
  - SysAD(63:0) Bus 169
    - Processor in Master State 169
    - Processor in Slave State 169
    - Correctable Error Detected 169
    - Uncorrectable Error Detected 169
  - SysState(2:0) Bus 170
  - SysResp(4:0) Bus 170
- Protocol Observation 171

10 CACHE Instructions

Notes on CACHE Instruction Operations 174
Virtual Address 174
Physical Address 174
CP0 Not Usable 174
TLB Refill and TLB Invalid Exceptions on CacheOps 175
Hit Operation Accesses 175
Watch Exception 175
Address Error Exception 175
Write Back 175
Invalidation 176
CE Bit 176
CH Bit 176
Serial Operation of CACHE Instructions 176
Instructions Not Supported 176
Op Field Encoding 177
Index Invalidate (I) 178
Index WriteBack Invalidate (D) 178
Index WriteBack Invalidate (S) 179
Index Load Tag (I) 180
Index Load Tag (D) 180
Index Load Tag (S) 181
Index Store Tag (I) 181
Index Store Tag (D) 182
Index Store Tag (S) 182
Hit Invalidate (I) 183
Hit Invalidate (D) 183
Hit Invalidate (S) 184
Cache Barrier 184
Hit Writeback Invalidate (D) 185
Hit WriteBack Invalidate (S) 186
Index Load Data (I) 187
Index Load Data (D) 187
Index Load Data (S) 187
Index Store Data (I) 188
Index Store Data (D) 188
Index Store Data (S) 188

11 JTAG Interface Operation

Test Access Port (TAP) 190
TAP Controller (Input) 190
Instruction Register 191
Bypass Register 191
Boundary Scan Register 192

12 Electrical Specifications

DC Electrical Specification 196
DC Power Supply Levels 196
DCOk and Power Supply Sequencing 197
Maximum Operating Conditions 197
Input Signal Level Sensing 198
Mode Definitions 198
Vref[SC,Sys] 198
Unused Inputs 198
DC Input/Output Specifications 199
AC Electrical Specification 200
Maximum Operating Conditions 200
Test Specification 200
Secondary Cache and System Interface Timing 200
Enable/Output Delay, Setup, Hold Time 201
Asynchronous Inputs 201
Signal Integrity Issues 202
Reference Voltage 202
Power Supply Regulation 202
Maximum Input Voltage Levels 202
Decoupling Capacitance 203
Output I-V Curves 203
Switching and Slew Rate Characteristics 203

13 Packaging

R10000 Single-Chip Package, 599CLGA 206
Mechanical Characteristics 206
Electrical Characteristics 207
Thermal Characteristics 208
Assembly Drawings and Pinout List 208
599CLGA Pinout 210

14
Coprocessor 0

Index Register (0) 221
Random Register (1) 222
EntryLo0 (2), and EntryLo1 (3) Registers 223
Context (4) 224
PageMask Register (5) 225
Wired Register (6) 226
BadVAddr Register (8) 227
Count and Compare Registers (9 and 11) 228
EntryHi Register (10) 229
Status Register (12) 230
- Status Register Fields 232
- Diagnostic Status Field 233
- Coprocessor Accessibility 235

Cause Register (13) 236
Exception Program Counter (14) 238
Processor Revision Identifier (PRId) Register (15) 239
Config Register (16) 240
Load Linked Address (LLAddr) Register (17) 241
WatchLo (18) and WatchHi (19) Registers 242
XContext Register (20) 243
FrameMask Register (21) 244
Diagnostic Register (22) 245
Performance Counter Registers (25) 247
ECC Register (26) 251
CacheErr Register (27) 252
- CacheErr Register Format for Primary Instruction Cache Errors 252
- CacheErr Register Format for Primary Data Cache Errors 253
- CacheErr Register Format for Secondary Cache Errors 253
- CacheErr Register Format for System Interface Errors 254

TagLo (28) and TagHi (29) Registers 255
- CacheOp is Index Load/Store Tag 256
  - Primary Instruction Cache Operation 256
  - Primary Data Cache Operation 257
  - Secondary Cache Operation 258
- CacheOp is Index Load/Store Data 259
  - Primary Instruction Cache Operation 259
  - Primary Data Cache Operation 260
  - Secondary Cache Operation 260

ErrorEPC Register (30) 261
CP0 Instructions 262
Hazards 262
Branch on Coprocessor 0 262
CP0 Move Instructions 263
CACHE Instruction 264
DMFC0 Instruction 267
DMTC0 Instruction 268
ERET Instruction 269
MFC0 Instruction 270
Move To/From the Performance Counter 271
MTC0 Instruction 272
TLBP Instruction 273
TLBR Instruction 274
TLBWI Instruction 275
TLBWR Instruction 276

15
Floating-Point Unit

Floating Point Unit Operations 278
Floating-Point Unit Control 279
Floating-Point General Registers (FGRs) 279
32- and 64-Bit Operations 280
Load and Store Operations 281
Floating-Point Control Registers 284
Floating-Point Implementation and Revision Register 284
Floating-Point Status Register (FSR) 285
Bit Descriptions of the FSR 286
Loading the FSR 287
FPU Instructions 288
CVT.L.fmt 288
Moves and Conditional Moves 289
CFC1/CTC1 289

16 Memory Management

Processor Modes....292
Processor Operating Modes....292
Addressing Modes....293
Virtual Address Space....293
User Mode Operations....294
32-bit User Mode (useg)....295
64-bit User Mode (xuseg)....295
Supervisor Mode Operations....296
32-bit Supervisor Mode, User Space (suseg)....296
32-bit Supervisor Mode, Supervisor Space (sseg)....297
64-bit Supervisor Mode, User Space (xsuseg)....297
64-bit Supervisor Mode, Current Supervisor Space (xsseg)....297
64-bit Supervisor Mode, Separate Supervisor Space (csseg)....297
Kernel Mode Operations....298
32-bit Kernel Mode, User Space (kuseg)....299
32-bit Kernel Mode, Kernel Space 0 (kseg0)....299
32-bit Kernel Mode, Kernel Space 1 (kseg1)....299
32-bit Kernel Mode, Supervisor Space (ksseg)....299
32-bit Kernel Mode, Kernel Space 3 (kseg3)....299
64-bit Kernel Mode, User Space (xkuseg)....300
64-bit Kernel Mode, Current Supervisor Space (xksseg)....300
64-bit Kernel Mode, Physical Spaces (xkphys)....300
64-bit Kernel Mode, Kernel Space (xkseg)....302
64-bit Kernel Mode, Compatibility Spaces (ckseg1:0, cksseg, ckseg3)....302
Address Space Access Privilege Differences Between the R4400 and R1000....302
Virtual Address Translation....304
Cache Pages....304
Cache Page Size Encodings....304
Using the TLB....305
Cache Algorithm Field....305
Format of a TLB Entry....305
Address Translation....306
Address Space Identification (ASID)....306
Global Processes (G)....306
Avoiding TLB Conflict....306

17 CPU Exceptions

Causing and Returning from an Exception 308
Exception Vector Locations 308
TLB Refill Vector Selection 309
Priority of Exceptions 311
Cold Reset Exception 312
Soft Reset Exception 313
NMI Exception 315
Address Error Exception 316
TLB Exceptions 317
TLB Refill Exception 318
Cause 318
Processing 318
Servicing 318
TLB Invalid Exception 319
Cause 319
Processing 319
Servicing 319
TLB Modified Exception 320
Cause 320
Processing 320
Servicing 320
Cache Error Exception 321
Virtual Coherency Exception 321
Bus Error Exception 322
Integer Overflow Exception 323
Trap Exception 324
System Call Exception 325
Breakpoint Exception 326
Reserved Instruction Exception 327
Coprocessor Unusable Exception 328
Floating-Point Exception 329
Watch Exception 330
Interrupt Exception 331
MIPSIV Instructions 332
COP0 Instructions 333
COP1 Instructions 333
COP2 Instructions 333

18 Cache Test Mode

Interface Signals....336
System Interface Clock Divisor....336
Entering Cache Test Mode....337
Exit Sequence....338
SysAD(63:0) Encoding....339
Cache Test Mode Protocol....340
Normal Write Protocol....340
Auto-Increment Write Protocol....341
Normal Read Protocol....342
Auto-Increment Read Protocol....343

A Glossary

Superscalar Processor....A-2
Pipeline....A-2
Pipeline Latency....A-2
Pipeline Repeat Rate....A-2
Out-of-Order Execution....A-2
Dynamic Scheduling....A-3
Instruction Fetch, Decode, Issue, Execution, Completion, and Graduation....A-3
Active List....A-3
Free List and Busy Registers....A-4
Register Renaming....A-4
Nonblocking Loads and Stores....A-5
Speculative Branching....A-6
Logical and Physical Registers....A-7
Register Files....A-7
ANDES Architecture....A-7

Introduction to the R10000 Processor

1

This user's manual describes the R10000 superscalar microprocessor for the system designer, paying special attention to the external interface and the transfer protocols.

This chapter describes the following:

- MIPS ISA
- what makes a generic superscalar microprocessor
- specifics of the R10000 superscalar microprocessor
- implementation-specific CPU instructions

1.1 MIPS Instruction Set Architecture (ISA)

MIPS has defined an instruction set architecture (ISA), implemented in the following sets of CPU designs:

- MIPS I, implemented in the R2000 and R3000
- MIPS II, implemented in the R6000
- MIPS III, implemented in the R4400
- MIPS IV, implemented in the R8000 and R10000

The original MIPS I CPU ISA has been extended forward three times, as shown in Figure 1-1; each extension is backward compatible. The ISA extensions are inclusive; each new architecture level (or version) includes the former levels.†

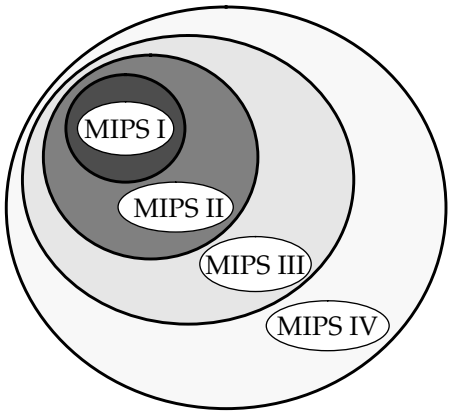

Figure 1-1 MIPS ISA with Extensions

The practical result is that a processor implementing MIPS IV is also able to run MIPS I, MIPS II, or MIPS III binary programs without change.

† For more ISA information, please refer to the *MIPS IV Instruction Set Architecture*, published by MIPS Technologies, and written by Charles Price. Contact information is provided both in the *Preface*, and inside the front cover, of this manual.

1.2 What is a Superscalar Processor?

A superscalar processor is one that can fetch, execute and complete more than one instruction in parallel.

Pipeline and Superpipeline Architecture

Previous MIPS processors had linear pipeline architectures; an example of such a linear pipeline is the R4400 superpipeline, shown in Figure 1-2. In the R4400 superpipeline architecture, an instruction is executed each cycle of the pipeline clock (PCycle), or each *pipe stage*.

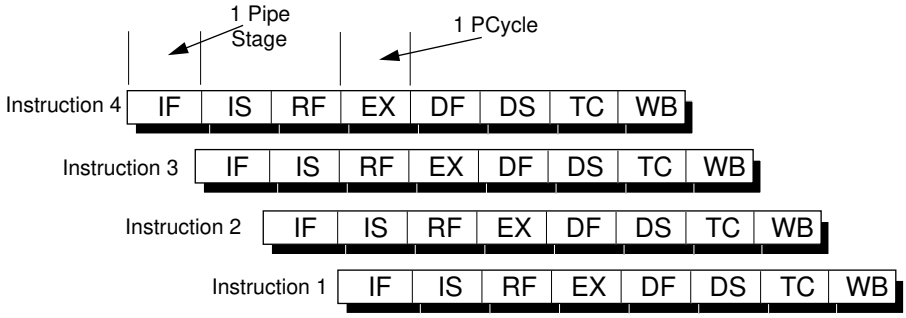

Figure 1-2 R4400 Pipeline

Superscalar Architecture

The structure of 4-way superscalar pipeline is shown in Figure 1-3. At each stage, four instructions are handled in parallel. Note that there is only one EX stage for integers.

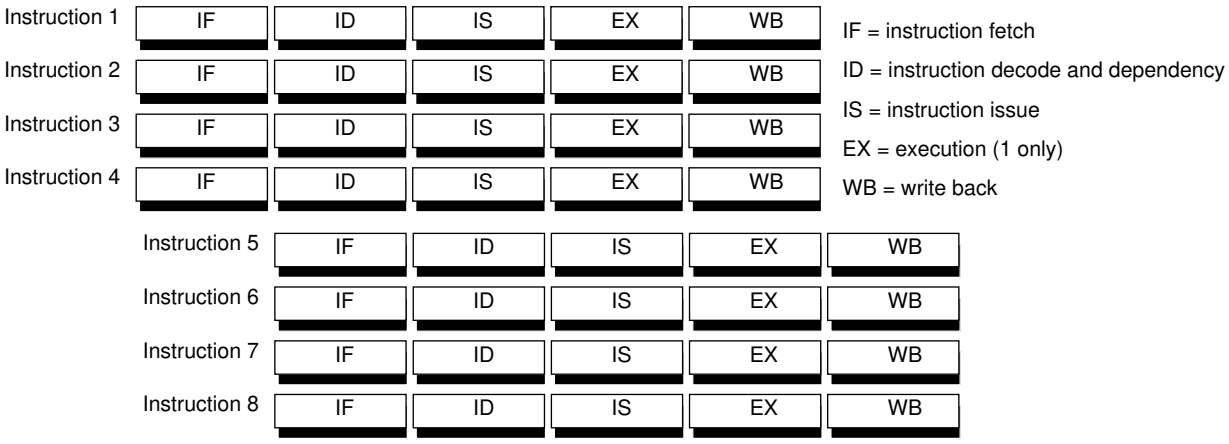

Figure 1-3 4-Way Superscalar Pipeline

1.3 What is an R10000 Microprocessor?

The R10000 processor is a single-chip superscalar RISC microprocessor that is a follow-on to the MIPS RISC processor family that includes, chronologically, the R2000, R3000, R6000, R4400, and R8000.

The R10000 processor uses the MIPS ANDES architecture, or *Architecture with Non-sequential Dynamic Execution Scheduling*.

The R10000 processor has the following major features (terms in **bold** are defined in the Glossary):

- it implements the 64-bit MIPS IV instruction set architecture (ISA)
- it can decode four instructions each pipeline cycle, appending them to one of three *instruction queues*
- it has five *execution pipelines* connected to separate internal integer and floating-point execution (or *functional*) units
- it uses **dynamic instruction scheduling** and **out-of-order execution**
- it uses speculative instruction issue (also termed "**speculative branching**")
- it uses a precise exception model (exceptions can be traced back to the instruction that caused them)
- it uses **non-blocking caches**
- it has separate on-chip 32-Kbyte primary instruction and data caches
- it has individually-optimized secondary cache and System interface ports
- it has an internal controller for the external secondary cache
- it has an internal System interface controller with multiprocessor support

The R10000 processor is implemented using 0.35-micron CMOS VLSI circuitry on a single 17 mm-by-18 mm chip that contains about 6.8 million transistors, including about 4.4 million transistors in its primary caches.

R10000 Superscalar Pipeline

The R10000 superscalar processor fetches and decodes four instructions in parallel each cycle (or pipeline stage). Each pipeline includes stages for fetching (stage 1 in Figure 1-4), decoding (stage 2) issuing instructions (stage 3), reading register operands (stage 3), executing instructions (stages 4 through 6), and storing results (stage 7).

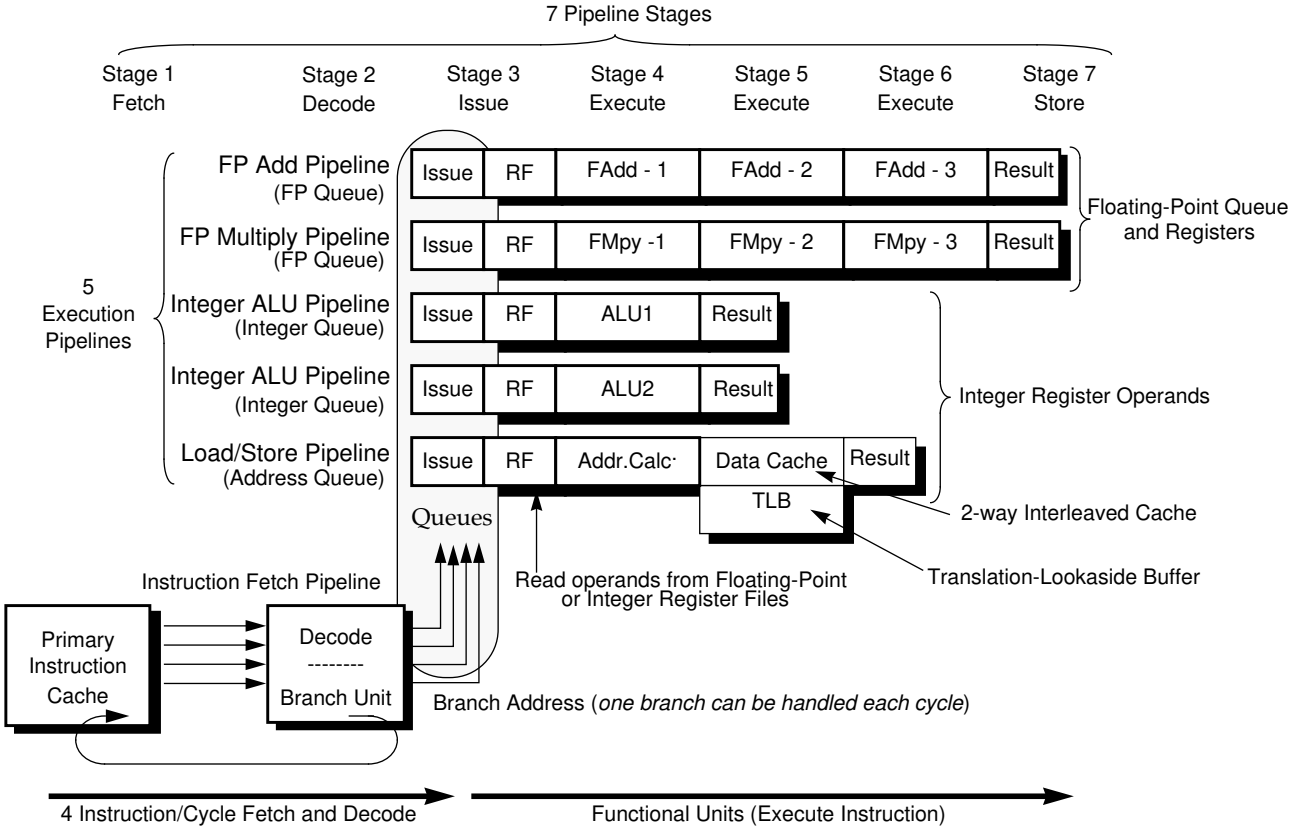

Figure 1-4 Superscalar Pipeline Architecture in the R10000

Instruction Queues

As shown in Figure 1-4, each instruction decoded in stage 2 is appended to one of three instruction *queues*:

- integer queue
- address queue
- floating-point queue

Execution Pipelines

The three instruction queues can issue (see the Glossary for a definition of *issue*) one new instruction per cycle to each of the five execution pipelines:

- the integer queue issues instructions to the two integer ALU pipelines
- the address queue issues one instruction to the Load/Store Unit pipeline
- the floating-point queue issues instructions to the floating-point adder and multiplier pipelines

A sixth pipeline, the fetch pipeline, reads and decodes instructions from the instruction cache.

64-bit Integer ALU Pipeline

The 64-bit integer pipeline has the following characteristics:

- it has a 16-entry integer instruction queue that dynamically issues instructions
- it has a 64-bit 64-location integer physical register file, with seven read and three write ports (32 logical registers; see *register renaming* in the Glossary)
- it has two 64-bit arithmetic logic units:
  - ALU1 contains an arithmetic-logic unit, shifter, and integer branch comparator
  - ALU2 contains an arithmetic-logic unit, integer multiplier, and divider

Load/Store Pipeline

The load/store pipeline has the following characteristics:

- it has a 16-entry address queue that dynamically issues instructions, and uses the integer register file for base and index registers
- it has a 16-entry address stack for use by non-blocking loads and stores
- it has a 44-bit virtual address calculation unit
- it has a 64-entry fully associative **Translation-Lookaside Buffer** (TLB), which converts virtual addresses to physical addresses, using a 40-bit physical address. Each entry maps two pages, with sizes ranging from 4 Kbytes to 16 Mbytes, in powers of 4.

64-bit Floating-Point Pipeline

The 64-bit floating-point pipeline has the following characteristics:

- it has a 16-entry instruction queue, with dynamic issue
- it has a 64-bit 64-location floating-point physical register file, with five read and three write ports (32 logical registers)
- it has a 64-bit parallel multiply unit (3-cycle pipeline with 2-cycle latency) which also performs move instructions
- it has a 64-bit add unit (3-cycle pipeline with 2-cycle latency) which handles addition, subtraction, and miscellaneous floating-point operations
- it has separate 64-bit divide and square-root units which can operate concurrently (these units share their issue and completion logic with the floating-point multiplier)

A block diagram of the processor and its interfaces is shown in Figure 1-5, followed by a description of its major logical blocks.

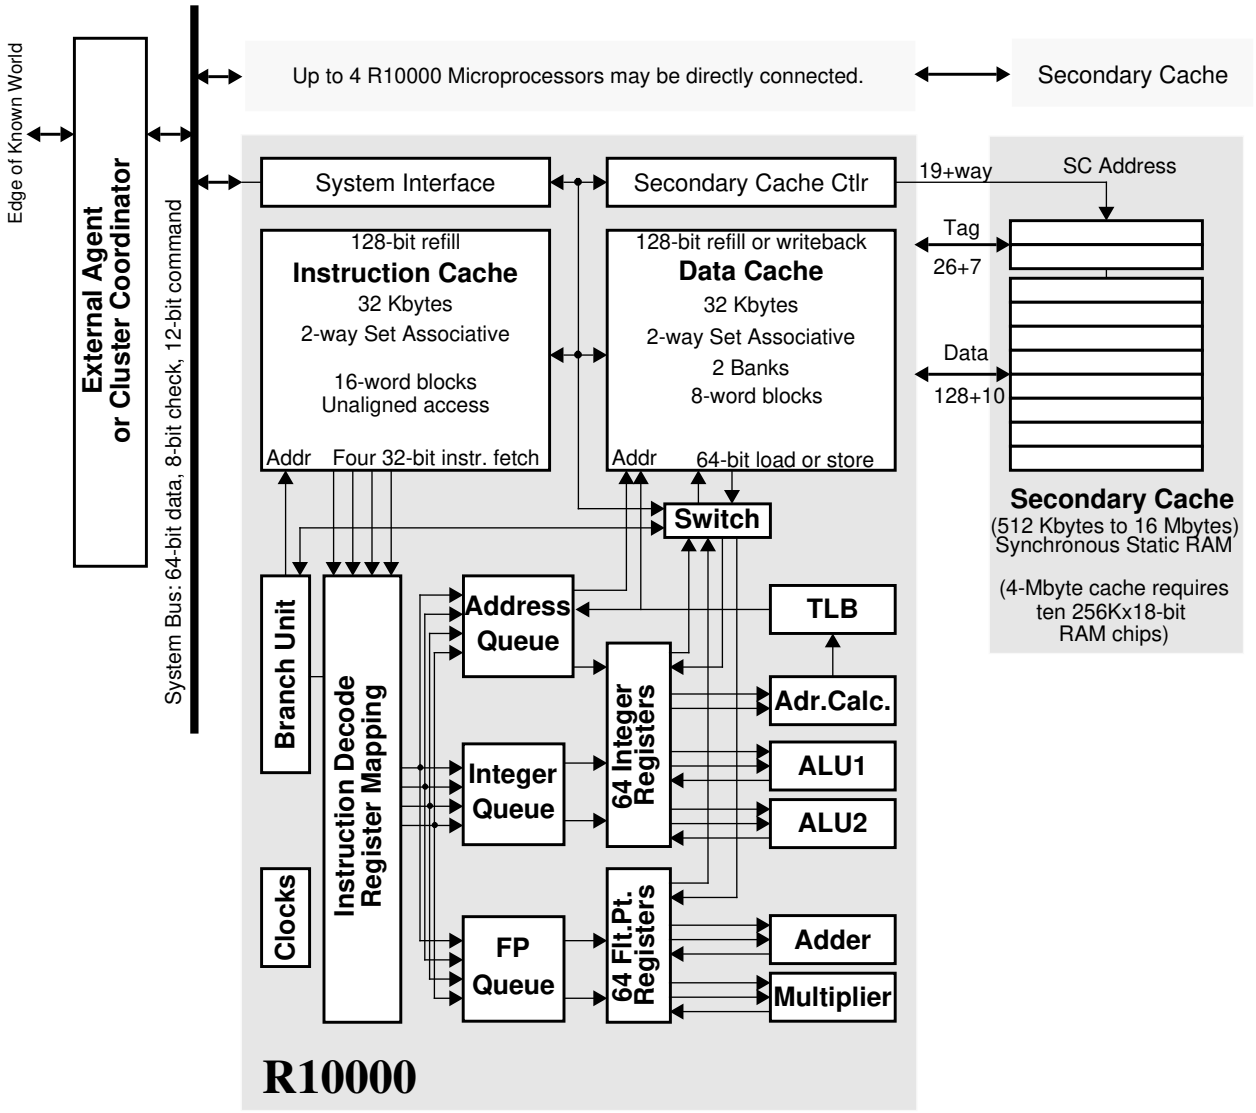

Figure 1-5 Block Diagram of the R10000 Processor

Functional Units

The five execution pipelines allow overlapped instruction execution by issuing instructions to the following five functional units:

- two integer ALUs (ALU1 and ALU2)
- the Load/Store unit (address calculate)
- the floating-point adder
- the floating-point multiplier

There are also three "iterative" units to compute more complex results:

- Integer multiply and divide operations are performed by an Integer Multiply/Divide execution unit; these instructions are issued to ALU2. ALU2 remains busy for the duration of the divide.
- Floating-point divides are performed by the Divide execution unit; these instructions are issued to the floating-point multiplier.
- Floating-point square root are performed by the Square-root execution unit; these instructions are issued to the floating-point multiplier.

Primary Instruction Cache (I-cache)

The primary instruction cache has the following characteristics:

- it contains 32 Kbytes, organized into 16-word blocks, is 2-way set associative, using a **least-recently used** (LRU) replacement algorithm
- it reads four consecutive instructions per cycle, beginning on any word boundary within a cache block, but cannot fetch across a block boundary.
- its instructions are predecoded, its fields are rearranged, and a 4-bit unit select code is appended
- it checks parity on each word
- it permits non-blocking instruction fetch

Primary Data Cache (D-cache)

The primary data cache has the following characteristics:

- it has two interleaved arrays (two 16 Kbyte ways)
- it contains 32 Kbytes, organized into 8-word blocks, is 2-way set associative, using an LRU replacement algorithm.
- it handles 64-bit load/store operations
- it handles 128-bit refill or write-back operations
- it permits non-blocking loads and stores
- it checks parity on each byte

Instruction Decode And Rename Unit

The instruction decode and rename unit has the following characteristics:

- it processes 4 instructions in parallel
- it replaces logical register numbers with physical register numbers (register renaming)
  - it maps integer registers into a 33-word-by-6-bit mapping table that has 4 write and 12 read ports
  - it maps floating-point registers into a 32-word-by-6-bit mapping table that has 4 write and 16 read ports
- it has a 32-entry active list of all instructions within the pipeline.

Branch Unit

The branch unit has the following characteristics:

- it allows one branch per cycle
- conditional branches can be executed speculatively, up to 4-deep
- it has a 44-bit adder to compute branch addresses
- it has a 4-quadword branch-resume buffer, used for reversing mispredicted speculatively-taken branches
- the Branch Link Quadword contains four instructions following a subroutine call, for rapid use when returning from leaf subroutines
- it has program trace RAM that stores the program counter for each instruction in the pipeline

External Interfaces

The external interfaces have the following characteristics:

- a 64-bit System interface allows direct-connection for 2-way to 4-way multiprocessor systems. 8-bit ECC Error Check and Correction is made on address and data transfers.
- a secondary cache interface with 128-bit data path and tag fields. 9-bit ECC Error Check and Correction is made on data quadwords, 7-bit ECC is made on tag words. It allows connection to an external secondary cache that can range from 512 Kbytes to 16 Mbytes, using external static RAMs. The secondary cache can be organized into either 16- or 32-word blocks, and is 2-way set associative.

Bit definitions are given in Chapter 3.

1.4 Instruction Queues

The processor keeps decoded instructions in three instruction queues, which dynamically issue instructions to the execution units. The queues allow the processor to fetch instructions at its maximum rate, without stalling because of instruction conflicts or dependencies.

Each queue uses instruction tags to keep track of the instruction in each execution pipeline stage. These tags set a *Done* bit in the active list as each instruction is completed.

Integer Queue

The integer queue issues instructions to the two integer arithmetic units: ALU1 and ALU2.

The integer queue contains 16 instruction entries. Up to four instructions may be written during each cycle; newly-decoded integer instructions are written into empty entries in no particular order. Instructions remain in this queue only until they have been issued to an ALU.

Branch and shift instructions can be issued only to ALU1. Integer multiply and divide instructions can be issued only to ALU2. Other integer instructions can be issued to either ALU.

The integer queue controls six dedicated ports to the integer register file: two operand read ports and a destination write port for each ALU.

Floating-Point Queue

The floating-point queue issues instructions to the floating-point multiplier and the floating-point adder.

The floating-point queue contains 16 instruction entries. Up to four instructions may be written during each cycle; newly-decoded floating-point instructions are written into empty entries in random order. Instructions remain in this queue only until they have been issued to a floating-point execution unit.

The floating-point queue controls six dedicated ports to the floating-point register file: two operand read ports and a destination port for each execution unit.

The floating-point queue uses the multiplier's issue port to issue instructions to the square-root and divide units. These instructions also share the multiplier's register ports.

The floating-point queue contains simple sequencing logic for multiple-pass instructions such as Multiply-Add. These instructions require one pass through the multiplier, then one pass through the adder.

Address Queue

The address queue issues instructions to the load/store unit.

The address queue contains 16 instruction entries. Unlike the other two queues, the address queue is organized as a circular **First-In First-Out** (FIFO) buffer. A newly decoded load/store instruction is written into the next available sequential empty entry; up to four instructions may be written during each cycle.

The FIFO order maintains the program's original instruction sequence so that memory address dependencies may be easily computed.

Instructions remain in this queue until they have graduated; they cannot be deleted immediately after being issued, since the load/store unit may not be able to complete the operation immediately.

The address queue contains more complex control logic than the other queues. An issued instruction may fail to complete because of a memory dependency, a cache miss, or a resource conflict; in these cases, the queue must continue to reissue the instruction until it is completed.

The address queue has three issue ports:

- First, it issues each instruction once to the address calculation unit. This unit uses a 2-stage pipeline to compute the instruction's memory address and to translate it in the TLB. Addresses are stored in the address stack and in the queue's dependency logic. This port controls two dedicated read ports to the integer register file. If the cache is available, it is accessed at the same time as the TLB. A tag check can be performed even if the data array is busy.
- Second, the address queue can re-issue accesses to the data cache. The queue allocates usage of the four sections of the cache, which consist of the tag and data sections of the two cache banks. Load and store instructions begin with a tag check cycle, which checks to see if the desired address is already in cache. If it is not, a refill operation is initiated, and this instruction waits until it has completed. Load instructions also read and align a doubleword value from the data array. This access may be either concurrent to or subsequent to the tag check. If the data is present and no dependencies exist, the instruction is marked *done* in the queue.
- Third, the address queue can issue store instructions to the data cache. A store instruction may not modify the data cache until it graduates. Only one store can graduate per cycle, but it may be anywhere within the four oldest instructions, if all previous instructions are already completed.

The access and store ports share four register file ports (integer read and write, floating-point read and write). These shared ports are also used for Jump and Link and Jump Register instructions, and for move instructions between the integer and register files.

1.5 Program Order and Dependencies

From a programmer's perspective, instructions appear to execute sequentially, since they are fetched and graduated in program order (the order they are presented to the processor by software). When an instruction stores a new value in its destination register, that new value is immediately available for use by subsequent instructions.

Internal to the processor, however, instructions are executed dynamically, and some results may not be available for many cycles; yet the hardware must behave as if each instruction is executed sequentially.

This section describes various conditions and dependencies that can arise from them in pipeline operation, including:

- instruction dependencies
- execution order and stalling
- branch prediction and speculative execution
- resolving operand dependencies
- resolving exception dependencies

Instruction Dependencies

Each instruction depends on all previous instructions which produced its operands, because it cannot begin execution until those operands become valid. These dependencies determine the order in which instructions can be executed.

Execution Order and Stalling

The actual execution order depends on the processor's organization; in a typical pipelined processor, instructions are executed only in program order. That is, the next sequential instruction may begin execution during the next cycle, if all of its operands are valid. Otherwise, the pipeline stalls until the operands do become valid.

Since instructions execute in order, stalls usually delay all subsequent instructions.

A clever compiler can improve performance by re-arranging instructions to reduce the frequency of these stall cycles.

- In an *in-order superscalar processor*, several consecutive instructions may begin execution simultaneously, if all their operands are valid, but the processor stalls at any instruction whose operands are still busy.
- In an *out-of-order superscalar processor*, such as the R10000, instructions are decoded and stored in queues. Each instruction is eligible to begin execution as soon as its operands become valid, independent of the original instruction sequence. In effect, the hardware rearranges instructions to keep its execution units busy. This process is called *dynamic issuing*.

Branch Prediction and Speculative Execution

Although one or more instructions may begin execution during each cycle, each instruction takes several (or many) cycles to complete. Thus, when a branch instruction is decoded, its branch condition may not yet be known. However, the R10000 processor can *predict* whether the branch is taken, and then continue decoding and executing subsequent instructions along the predicted path.

When a branch prediction is wrong, the processor must back up to the original branch and take the other path. This technique is called *speculative execution*. Whenever the processor discovers a mispredicted branch, it aborts all speculatively-executed instructions and restores to the processor to the state it held before the branch.

Resolving Operand Dependencies

Operands include registers, memory, and condition bits. Each operand type has its own dependency logic. In the R10000 processor, dependencies are resolved in the following manner:

- register dependencies are resolved by using *register renaming* and the associative comparator circuitry in the queues
- memory dependencies are resolved in the Load/Store Unit
- condition bit dependencies are resolved in the active list and instruction queues

Resolving Exception Dependencies

In addition to operand dependencies, each instruction is implicitly dependent upon any previous instruction that generates an exception. Exceptions are caused whenever an instruction cannot be properly completed, and are usually due to either an untranslated virtual address or an erroneous operand.

The processor design implements *precise exceptions*, by:

- identifying the instruction which caused the exception
- preventing the exception-causing instruction from graduating
- aborting all subsequent instructions

Thus, all register values remain the same as if instructions were executed singly. Effectively, all previous instructions are completed, but the faulting instruction and all subsequent instructions do not modify any values.

Strong Ordering

A multiprocessor system that exhibits the same behavior as a uniprocessor system in a multiprogramming environment is said to be *strongly ordered*.

The R10000 processor behaves as if strong ordering is implemented, although it does not actually execute all memory operations in strict program order.

In the R10000 processor, store operations remain pending until the store instruction is ready to graduate. Thus, stores are executed in program order, and memory values are precise following any exception.

For improved performance however, cached load operations my occur in any order, subject to memory dependencies on pending store instructions. To maintain the appearance of strong ordering, the processor detects whenever the reordering of a cached load might alter the operation of the program, backs up, and then re-executes the affected load instructions. Specifically, whenever a primary data cache block is invalidated due to an external coherency request, its index is compared with all outstanding load instructions. If there is a match and the load has been completed, the load is prevented from graduating. When it is ready to graduate, the entire pipeline is flushed, and the processor is restored to the state it had before the load was decoded.

An uncached or uncached accelerated load or store instruction is executed when the instruction is ready to graduate. This guarantees strong ordering for uncached accesses.

Since the R10000 processor behaves as if it implemented strong ordering, a suitable system design allows the processor to be used to create a shared-memory multiprocessor system with strong ordering.

An Example of Strong Ordering

Given that locations X and Y have no particular relationship—that is, they are not in the same cache block—an example of strong ordering is as follows:

- Processor A performs a store to location X and later executes a load from location Y.
- Processor B performs a store to location Y and later executes a load from location X.

The two processors are running asynchronously, and the order of the above two sequences is unknown.

For the system to be strongly ordered, either processor A must load the new value of Y, or processor B must load the new value of X, or both processors A and B must load the new values of Y and X, respectively, under all conditions.

If processors A and B both load old values of Y and X, respectively, under any conditions, the system is not strongly ordered.

New Value		Strongly Ordered
Processor A	**Processor B**	
No	No	No
Yes	No	Yes
No	Yes	Yes
Yes	Yes	Yes

1.6 R10000 Pipelines

This section describes the stages of the superscalar pipeline.

Instructions are processed in six partially-independent pipelines, as shown in Figure 1-4. The Fetch pipeline reads instructions from the instruction cache†, decodes them, renames their registers, and places them in three instruction queues. The instruction queues contain integer, address calculate, and floating-point instructions. From these queues, instructions are dynamically issued to the five pipelined execution units.

Stage 1

In stage 1, the processor fetches four instructions each cycle, independent of their alignment in the instruction cache — except that the processor cannot fetch across a 16-word cache block boundary. These words are then aligned in the 4-word *Instruction* register.

If any instructions were left from the previous decode cycle, they are merged with new words from the instruction cache to fill the *Instruction* register.

Stage 2

In stage 2, the four instructions in the *Instruction* register are decoded and renamed. (Renaming determines any dependencies between instructions and provides precise exception handling.) When renamed, the *logical* registers referenced in an instruction are mapped to *physical* registers. Integer and floating-point registers are renamed independently.

A logical register is mapped to a new physical register whenever that logical register is the destination of an instruction. Thus, when an instruction places a new value in a logical register, that logical register is renamed (mapped) to a new physical register, while its previous value is retained in the old physical register.

As each instruction is renamed, its logical register numbers are compared to determine if any dependencies exist between the four instructions decoded during this cycle. After the physical register numbers become known, the Physical Register Busy table indicates whether or not each operand is valid. The renamed instructions are loaded into integer or floating-point instruction queues.

Only one branch instruction can be executed during stage 2. If the instruction register contains a second branch instruction, this branch is not decoded until the next cycle.

The branch unit determines the next address for the Program Counter; if a branch is taken and then reversed, the branch resume cache provides the instructions to be decoded during the next cycle.

† The processor checks only the instruction cache during an instruction fetch; it does not check the data cache.

Stage 3

In stage 3, decoded instructions are written into the queues. Stage 3 is also the start of each of the five execution pipelines.

Stages 4-6

In stages 4 through 6, instructions are executed in the various functional units. These units and their execution process are described below.

Floating-Point Multiplier (3-stage Pipeline)

Single- or double-precision multiply and conditional move operations are executed in this unit with a 2-cycle latency and a 1-cycle repeat rate. The multiplication is completed during the first two cycles; the third cycle is used to pack and transfer the result.

Floating-Point Divide and Square-Root Units

Single- or double-precision division and square-root operations can be executed in parallel by separate units. These units share their issue and completion logic with the floating-point multiplier.

Floating-Point Adder (3-stage Pipeline)

Single- or double-precision add, subtract, compare, or convert operations are executed with a 2-cycle latency and a 1-cycle repeat rate. Although a final result is not calculated until the third pipeline stage, internal bypass paths set a 2-cycle latency for dependent add or multiply instructions.

Integer ALU1 (1-stage Pipeline)

Integer add, subtract, shift, and logic operations are executed with a 1-cycle latency and a 1-cycle repeat rate. This ALU also verifies predictions made for branches that are conditional on integer register values.

Integer ALU2 (1-stage Pipeline)

Integer add, subtract, and logic operations are executed with a 1-cycle latency and a 1-cycle repeat rate. Integer multiply and divide operations take more than one cycle.

Address Calculation and Translation in the TLB

A single memory address can be calculated every cycle for use by either an integer or floating-point load or store instruction. Address calculation and load operations can be calculated out of program order.

The calculated address is translated from a 44-bit virtual address into a 40-bit physical address using a translation-lookaside buffer. The TLB contains 64 entries, each of which can translate two pages. Each entry can select a page size ranging from 4 Kbytes to 16 Mbytes, inclusive, in increments of 4, as shown in Figure 1-6.

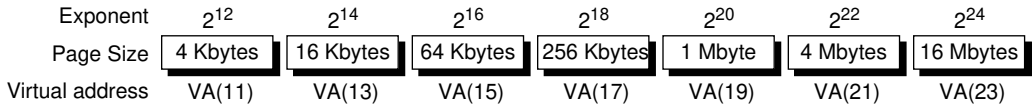

Figure 1-6 TLB Page Sizes

Load instructions have a 2-cycle latency if the addressed data is already within the data cache.

Store instructions do not modify the data cache or memory until they graduate.

1.7 Implications of R10000 Microarchitecture on Software

The R10000 processor implements the MIPS architecture by using the following techniques to improve throughput:

- *superscalar instruction issue*
- *speculative execution*
- *non-blocking caches*

These microarchitectural techniques have special implications for compilation and code scheduling.

Superscalar Instruction Issue

The R10000 processor has parallel functional units, allowing up to four instructions to be fetched and up to five instructions to be issued or completed each cycle. An ideal code stream would match the fetch bandwidth of the processor with a mix of independent instructions to keep the functional units as busy as possible.

To create this ideal mix, every cycle the hardware would select one instruction from each of the columns below. (Floating-point divide, floating-point square root, integer multiply and integer divide cannot be started on each cycle.) The processor can look ahead in the code, so the mix should be kept close to the ideal described below.

Column A	Column B	Column C	Column D	Column E
FPadd	FP mul	FPload	add/sub	add/sub
	FPdiv	FPstore	shift	mul
	FPsqrt	load	branch	div
		store	logical	logical

Data dependencies are detected in hardware, but limit the degree of parallelism that can be achieved. Compilers can intermix instructions from independent code streams.

Speculative Execution

Speculative execution increases parallelism by fetching, issuing, and completing instructions even in the presence of unresolved conditional branches and possible exceptions. Following are some suggestions for increasing program efficiency:

- Compilers should reduce the number of branches as much as possible
- "Jump Register" instructions should be avoided.
- Aggressive use of the new integer and floating point conditional move instructions is recommended.
- Branch prediction rates may be improved by organizing code so that each branch goes the same direction most of the time, since a branch that is taken 50% of the time has higher average cost than one taken 90% of the time. The MIPS IV conditional move instructions may be effective in improving performance by replacing unpredictable branches.

Nonblocking Caches

As processor speed increases, the processor's data latency and bandwidth requirements rise more rapidly than the latency and bandwidth of cost-effective main memory systems. The memory hierarchy of the R10000 processor tries to minimize this effect by using large set-associative caches and higher bandwidth cache refills to reduce the cost of loads, stores, and instruction fetches. Unlike the R4400, the R10000 processor does not stall on data cache misses, instead defers execution of any dependent instructions until the data has been returned and continues to execute independent instructions (including other memory operations that may miss in the cache). Although the R10000 allows a number of outstanding primary and secondary cache misses, compilers should organize code and data to reduce cache misses. When cache misses are inevitable, the data reference should be scheduled as early as possible so that the data can be fetched in parallel with other unrelated operations.

As a further antidote to cache miss stalls, the R10000 processor supports prefetch instructions, which serve as hints to the processor to move data from memory into the secondary and primary caches when possible. Because prefetches do not cause dependency stalls or memory management exceptions, they can be scheduled as soon as the data address can be computed, without affecting exception semantics. Indiscriminate use of prefetch instructions can slow program execution because of the instruction-issue overhead, but selective use of prefetches based on compiler miss prediction can yield significant performance improvement for dense matrix computations.

1.8 R10000-Specific CPU Instructions

This section describes the processor-specific implementations of the following instructions:

- PREF
- LL/SC
- SYNC

Chapter 14 describes the CP0-specific instructions, and Chapter 15 describes the FPU-specific instructions.

PREF

In the R1000 processor, the Prefetch instruction, PREF, attempts to fetch data into the secondary and primary data caches. The action taken by a Prefetch instruction is controlled by the instruction hint field, as decoded in Table 1-1.

Table 1-1 PREF Instruction Hint Field

Hint Value	Name of Hint	Action Taken
0	Load	Prefetch data into cache LRU way
1	Store	Prefetch data into cache LRU way
2-3		*undefined*
4	load_streamed	Prefetch data into cache way 0
5	store_streamed	Prefetch data into cache way 0
6	load_retained	Prefetch data into cache way 1
7	store_retained	Prefetch data into cache way 1
8-31		*undefined*

For a “store” Prefetch, an *Exclusive* copy of the cache block must be obtained, in order that it may be written.

LL/SC

Load Linked and Store Conditional instructions are used together to implement a memory semaphore. Each LL/SC sequence has three sections:

1. The LL loads a word from memory.
2. A short sequence of instructions checks or modifies this word. This sequence must not contain any of the events listed below, or the Store Conditional will fail:
   - exception
   - execution of ERET
   - load instruction
   - store instruction
   - SYNC instruction
   - CACHE instruction
   - PREF instruction
   - external intervention exclusive or invalidate to the secondary cache block containing the linked address
3. The SC stores a new value into the memory word, unless the new value has been modified. If the word has not been modified, the store succeeds and a 1 is stored in the destination register. Otherwise the Store Conditional fails, memory is not modified, and a 0 is loaded into the destination register. Since the instruction format has only a single field to select a data register (*rt*), this destination register is the same as the register which was stored.

Load Linked and Store Conditional instructions (LL, LLD, SC, and SCD) do not implicitly perform SYNC operations in the R10000 processor.

SYNC

The SYNC instruction is implemented in a "lightweight" manner: after decoding a SYNC instruction, the processor continues to fetch and decode further instructions. It is allowed to issue load and store instructions speculatively and out-of-order, following a SYNC.

The R10000 processor only allows a SYNC instruction to graduate when the following conditions are met:

- all previous instructions have been successfully completed
- the uncached buffer does not contain any uncached stores
- the address cycle of a processor double/single/partial-word write request resulting from an uncached store was not issued to the System interface in any of the prior three **SysClk** cycles
- the **SysGblPerf*** signal is asserted

A SYNC instruction is not prevented from graduating if the uncached buffer contains any uncached accelerated stores.

1.9 Performance

As it executes programs, the R10000 superscalar processor performs many operations in parallel. Instructions can also be executed out of order. Together, these two facts greatly improve performance, but they also make it difficult to predict the time required to execute any section of a program, since it often depends on the instruction mix and the critical dependencies between instructions.

The processor has five largely independent execution units, each of which are individualized for a specific class of instructions. Any one of these units may limit processor performance, even as the other units sit idle. If this occurs, instructions which use the idle units can be added to the program without adding any appreciable delay.

User Instruction Latency and Repeat Rate

Table 1-2 shows the latencies and repeat rates for all user instructions executed in ALU1, ALU2, Load/Store, Floating-Point Add and Floating-Point Multiply functional units (definitions of *latency* and *repeat rate* are given in the Glossary). Kernel instructions are not included, nor are control instructions not issued to these execution units.

Table 1-2 Latencies and Repeat Rates for User Instructions

Instruction Type	Execution Unit	Latency	Repeat Rate	Comment
Integer Instructions				
Add/Sub/Logical/Set	ALU 1/2	1	1	
MF/MT HI/LO	ALU 1/2	1	1	
Shift/LUI	ALU 1	1	1	
Cond. Branch Evaluation	ALU 1	1	1	
Cond. Move	ALU 1	1	1	
MULT	ALU 2	5/6	6	Latency relative to Lo/Hi
MULTU	ALU 2	6/7	7	Latency relative to Lo/Hi
DMULT	ALU 2	9/10	10	Latency relative to Lo/Hi
DMULTU	ALU 2	10/11	11	Latency relative to Lo/Hi
DIV/DIVU	ALU 2	34/35	35	Latency relative to Lo/Hi
DDIV/DDIVU	ALU 2	66/67	67	Latency relative to Lo/Hi
Load (not include loads to CP1)	Load/Store	2	1	Assuming cache hit
Store	Load/Store	-	1	Assuming cache hit
Floating-Point Instructions				
MTC1/DMTC1	ALU 1	3	1	
Add/Sub/Abs/Neg/Round/ Trunc/Ceil/Floor/C.cond	FADD	2	1	
CVT.S.W/CVT.S.L	FADD	4	2	Repeat rate is on average
CVT (others)	FADD	2	1	
Mul	FMPY	2	1	
MFC1/DMFC1	FMPY	2	1	
Cond. Move/Move	FMPY	2	1	
DIV.S/RECIP.S	FMPY	12	14	
DIV.D/RECIP.D	FMPY	19	21	
SQRT.S	FMPY	18	20	
SQRT.D	FMPY	33	35	
RSQRT.S	FMPY	30	20	
RSQRT.D	FMPY	52	35	
MADD	FADD+FMPY	2/4	1	Latency is 2 only if the result is used as the operand specified by *fr* of another MADD
LWC1/LDC1/LWXC1/LDXC1	LoadStore	3	1	Assuming cache hit

Please note the following about Table 1-2:

- For integer instructions, conditional trap evaluation takes a single cycle, like conditional branches.
- Branches and conditional moves are not conditionally issued.
- The repeat rate above for Load/Store does not include Load Link and Store Conditional.

- Prefetch instruction is not included here.
- The latency for multiplication and division depends upon the next instruction.
- An instruction using register *Lo* can be issued one cycle earlier than one using *Hi*.
- For floating-point instructions, CP1 branches are evaluated in the Graduation Unit.
- Up to four branches can be evaluated at one cycle.†
- CTC1 and CFC1 are not included in this table.
- The repeat pattern for the CVT.S.(W/L) is "I I x x I I x x ..."; the repeat rate given here, 2, is the average.
- The latency for MADD instructions is 2 cycles if the result is used as the operand specified by *fr* of the second MADD instruction.
- Load Linked and Store Conditional instructions (LL, LLD, SC, and SCD) do not implicitly perform SYNC operations in the R10000. Any of the following events that occur between a Load Linked and a Store Conditional will cause the Store Conditional to fail: an exception; execution of an ERET, a load, a store, a SYNC, a CacheOp, a prefetch, or an external intervention/invalidation on the block containing the linked address. Instruction cache misses do not cause the Store Conditional to fail.

For more information about implementations of the LL, SC, and SYNC instructions, please see the section titled, R10000-Specific CPU Instructions, in this chapter.

† Only one branch can be decoded at any particular cycle. Since each conditional branch is predicted, the real direction of each branch must be "evaluated." For example,

```
beq  r2,r3,L1
nop
```

A comparison of r2 and r3 is made to determine whether the branch is taken or not. If the branch prediction is correct, the branch instruction is graduated. Otherwise, the processor must back out of the instruction stream decoded after this branch, and inform the IFetch to fetch the correct instructions. The evaluation is made in the ALU for integer branches and in the Graduation Unit for floating-point branches. A single integer branch can be evaluated during any cycle, but there may be up to 4 condition codes waiting to be evaluated for floating-point branches. Once the condition code is evaluated, all dependant FP branches can be evaluated during the same cycle.

Other Performance Issues

Table 1-2 shows execution times within the functional units only. Performance may also be affected by instruction fetch times, and especially by the execution of conditional branches.

In an effort to keep the execution units busy, the processor predicts branches and speculatively executes instructions along the predicted path. When the branch is predicted correctly, this significantly improves performance: for typical programs, branch prediction is 85% to 90% correct. When a branch is mispredicted, the processor must discard instructions which were speculatively fetched and executed. Usually, this effort uses resources which otherwise would have been idle, however in some cases speculative instructions can delay previous instructions.

Cache Performance

The execution of load and store instructions can greatly affect performance. These instructions are executed quickly if the required memory block is contained in the primary data cache, otherwise there are significant delays for accessing the secondary cache or main memory. Out-of-order execution and non-blocking caches reduce the performance loss due to these delays, however.

The latency and repeat rates for accessing the secondary cache are summarized in Table 1-3. These rates depend on the ratio of the secondary cache's clock to the processor's internal pipeline clock. The best performance is achieved when the clock rates are equal; slower external clocks add to latency and repeat times.

The primary data cache contains 8-word blocks, which are refilled using 2-cycle transfers from the quadword-wide secondary cache. Latency runs to the time in which the processor can use the addressed data.

The primary instruction cache contains 16-word blocks, which are refilled using 4-cycle transfers.

Table 1-3 Latency and Repeat Rates for Secondary Cache Reads

SCClkDiv Mode	Latency‡ (PClk Cycles)	Repeat Rate* (PClk Cycles)
1	6	2 (data cache) 4 (instruction cache)
1.5	8-10†	3 (data cache) 6 (instruction cache)
2	9-12†	4 (data cache) 8 (instruction cache)

‡ Assumes the cache way was correctly predicted, and there are no conflicting requests.

* Repeat rate = **PClk** cycles needed to transfer 2 quadwords (data cache) or 4 quadwords (instruction cache). Rate is valid for bursts of 2 to 3 cache misses; if more than three cache misses in a row, there can be a 1-cycle "bubble."

† Clock synchronization causes variability.

The processor mitigates access delays to the secondary cache in the following ways:

- The processor can execute up to 16 load and store instructions speculatively and out-of-order, using non-blocking primary and secondary caches. That is, it looks ahead in its instruction stream to find load and store instructions which can be executed early; if the addressed data blocks are not in the primary cache, the processor initiates cache refills as soon as possible.
- If a speculatively executed load initiates a cache refill, the refill is completed even if the load instruction is aborted. It is likely the data will be referenced again.
- The data cache is interleaved between two banks, each of which contains independent tag and data arrays. These four sections can be allocated separately to achieve high utilization. Five separate circuits compete for cache bandwidth (address calculate, tag check, load unit, store unit, external interface.)
- The external interface gives priority to its refill and interrogate operations. The processor can execute tag checks, data reads for load instructions, or data writes for store instructions. When the primary cache is refilled, any required data can be streamed directly to waiting load instructions.
- The external interface can handle up to four non-blocking memory accesses to secondary cache and main memory.

Main memory typically has much longer latencies and lower bandwidth than the secondary cache, which make it difficult for the processor to mitigate their effect. Since main memory accesses are non-blocking, delays can be reduced by overlapping the latency of several operations. However, although the first part of the latency may be concealed, the processor cannot look far enough ahead to hide the entire latency.

Programmers may use pre-fetch instructions to load data into the caches before it is needed, greatly reducing main memory delays for programs which access memory in a predictable sequence.

System Configurations

2

The R10000 processor provides the capability for a wide range of computer systems; this chapter describes some of the uni- and multiprocessor alternatives.

2.1 Uniprocessor Systems

In a typical uniprocessor system, the System interface of the R10000 processor connects in a point-to-point fashion with an external agent. Such a system is shown in Figure 2-1. The external agent is typically an ASIC that provides a gateway to the memory and I/O subsystems; in fact, this ASIC may incorporate the memory controller itself.

If hardware I/O coherency is desired, the external agent may use the multiprocessor primitives provided by the processor to maintain cache coherency for interventions and invalidations. External duplicate tags can be used by the external agent to filter external coherency requests.

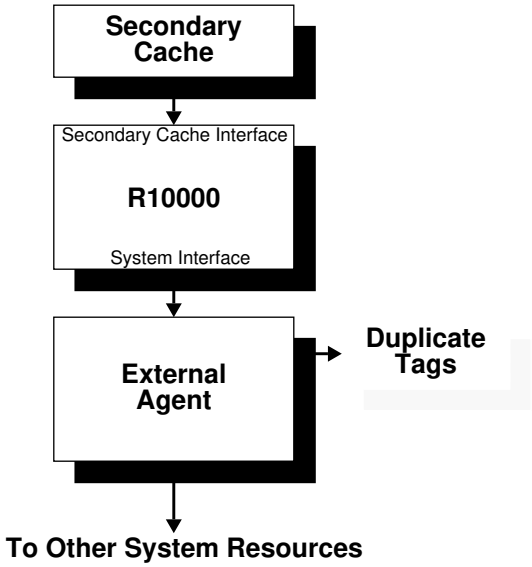

Figure 2-1 Uniprocessor System Organization

2.2 Multiprocessor Systems

Two types of multiprocessor systems can be implemented with R10000 processor:

- a dedicated external agent interfaces with each R10000 processor
- up to four R10000 processors and an external agent reside on a cluster bus

Multiprocessor Systems Using Dedicated External Agents

A multiprocessor system may be created with R10000 processors by providing a dedicated external agent for each processor; such a system is shown in Figure 2-2. The external agent provides a path between the processor System interface and some type of coherent interconnect. In such a system, the processor provides support for three coherency schemes:

- snoopy-based
- snoopy-based with external duplicate tags and control
- directory-based with external directory structure and control

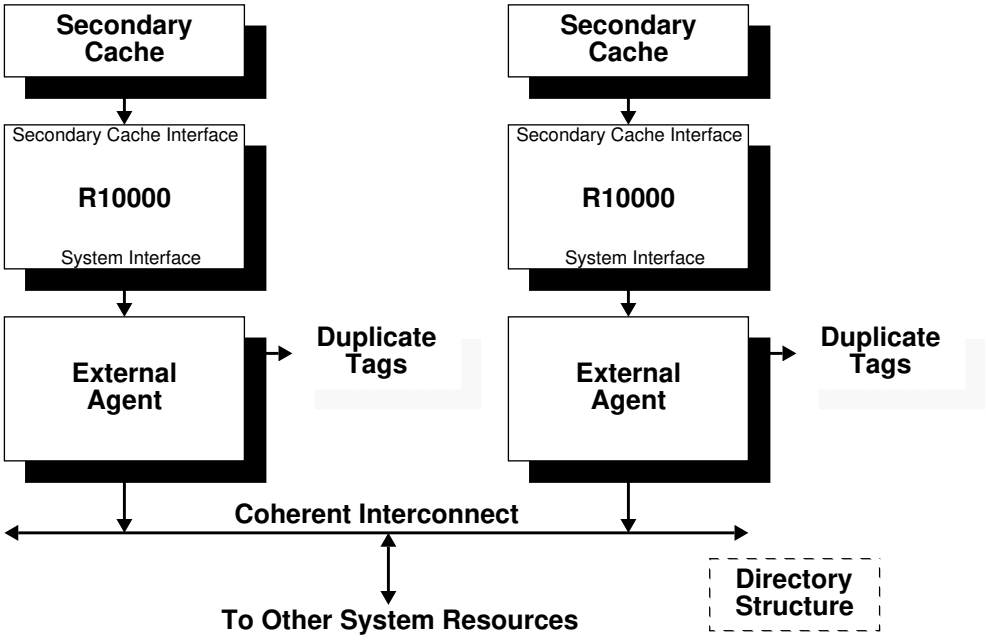

Figure 2-2 Multiprocessor System Organization using Dedicated External Agents

Multiprocessor Systems Using a Cluster Bus

A multiprocessor system may be created with R10000 processors by using a cluster bus configuration. Such a system is shown in Figure 2-3. A cluster bus is created by attaching the System interfaces of up to four R10000 processors with an external agent (the *cluster coordinator*). The cluster coordinator is responsible for managing the flow of data within the cluster.

This organization can reduce the number of ASICs and the pin count needed for a small multiprocessor systems.

The cluster bus protocol supports three coherency schemes:

- snoopy-based
- snoopy-based with external duplicate tags and control
- directory-based with external directory structure and control

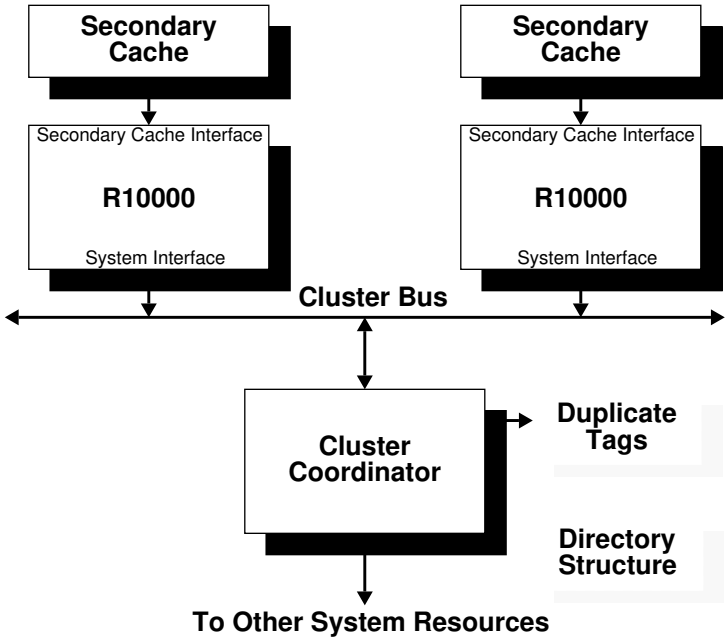

Figure 2-3 Multiprocessor System Organization Using the Cluster Bus

Interface Signal Descriptions

3

This chapter gives a list and description of the interface signals.

The R10000 interface signals may be divided into the following groups:

- Power interface
- Secondary Cache interface
- System interface
- Test interface

The following sections present a summary of the external interface signals for each of these groups. An asterisk (*) indicates signals that are asserted as a logical 0.

3.1 Power Interface Signals

Table 3-1 presents the R10000 processor power interface signals.

Table 3-1 Power Interface Signals

Signal Name	Description	Type
Vcc	**Vcc** core **Vcc** for the core circuits.	Input
VccQSC	**Vcc** output driver secondary cache **Vcc** for the secondary cache interface output drivers.	Input
VccQSys	**Vcc** output driver system **Vcc** for the System interface output drivers.	Input
VrefSC	Voltage reference secondary cache Voltage reference for the secondary cache interface input receivers.	Input
VrefSys	Voltage reference system Voltage reference for the System interface input receivers.	Input
VrefByp	Voltage reference bypass This pin must be tied to **VrefSys** or **Vss**.	Input
Vss	**Vss** **Vss** for the core circuits and output drivers.	Input
VccPa	**Vcc** PLL analog **Vcc** for the PLL analog circuits.	Input
VssPa	**Vss** PLL analog **Vss** for the PLL analog circuits.	Input
VccPd	**Vcc** PLL digital **Vcc** for the PLL digital circuits.	Input
VssPd	**Vss** PLL digital **Vss** for the PLL digital circuits.	Input
DCOk	DC voltages are OK The external agent asserts these two signals when **Vcc**, **VccQ[SC,Sys]**, **Vref[SC,Sys]**, **Vcc[Pa,Pd]**, and **SysClk** are stable.	Input

3.2 Secondary Cache Interface Signals

Table 3-2 presents the R10000 processor secondary cache interface signals.

Table 3-2 Secondary Cache Interface Signals

Signal Name	Description	Type
SSRAM* Clock Signals		
SCClk(5:0) SCClk*(5:0)	Secondary cache clock Duplicated complementary secondary cache clock outputs.	Output
SSRAM Address Signals		
SCAAddr(18:0) SCBAddr(18:0)	Secondary cache address bus Duplicated complementary 19-bit bus which specifies the set address of the secondary cache data and tag SSRAM that is to be accessed.	Output
SCTagLSBAddr	Secondary cache tag LSB address Signal that specifies the least significant bit of the address for the secondary cache tag SSRAM.	Output
SSRAM Data Signals		
SCADWay SCBDWay	Secondary cache data way Duplicated signal that indicates the way of the secondary cache data SSRAM that is to be accessed.	Output
SCData(127:0)	Secondary cache data bus 128-bit bus to read/write cache data from/to secondary cache data SSRAM.	Bidirectional
SCDataChk(9:0)	Secondary cache data check bus A 10-bit bus used to read/write ECC and even parity from/to the secondary cache data SSRAM.	Bidirectional
SCADOE* SCBDOE*	Secondary cache data output enable Duplicated signal that enables the outputs of the secondary cache data SSRAM.	Output
SCADWr* SCBDWr*	Secondary cache data write enable Duplicated signal that enables writing the secondary cache data SSRAM.	Output
SCADCS* SCBDCS*	Secondary cache data chip select Duplicated signal that enables the secondary cache data SSRAM.	Output
SSRAM Tag Signals		
SCTWay	Secondary cache tag way Signal indicating the way of the secondary cache tag SSRAM to be accessed.	Output
SCTag(25:0)	Secondary cache tag bus A 26-bit bus to read/write cache tags from/to the secondary cache tag SSRAM.	Bidirectional
SCTagChk(6:0)	Secondary cache tag check bus A 7-bit bus used to read/write ECC from/to the secondary cache tag SSRAM.	Bidirectional
SCTOE*	Secondary cache tag output enable A signal that enables the outputs of the secondary cache tag SSRAM.	Output
SCTWr*	Secondary cache tag write enable A signal that enables writing the secondary cache tag SSRAM.	Output
SCTCS*	Secondary cache tag chip select A signal which enables the secondary cache tag SSRAM.	Output

* All cache static RAM (SRAM) are synchronous SRAM (SSRAM).

3.3 System Interface Signals

Table 3-3 presents the R10000 processor System interface signals.

Table 3-3 System Interface Signals

Signal Name	Description	Type
System Clock Signals		
SysClk SysClk*	System clock Complementary system clock input.	Input
SysClkRet SysClkRet*	System clock return Complementary system clock return output used for termination of the system clock.	Output
System Arbitration Signals		
SysReq*	System request The processor asserts this signal when it wants to perform a processor request and it is not already master of the System interface.	Output
SysGnt*	System grant The external agent asserts this signal to grant mastership of the System interface to the processor.	Input
SysRel*	System release The master of the System interface asserts this signal for one **SysClk** cycle to indicate that it will relinquish mastership of the System interface in the following **SysClk** cycle.	Bidirectional
System Flow Control Signals		
SysRdRdy*	System read ready The external agent asserts this signal to indicate that it can accept processor read and upgrade requests.	Input
SysWrRdy*	System write ready The external agent asserts this signal to indicate that it can accept processor write and eliminate requests.	Input
System Address/Data Bus Signals		
SysCmd(11:0)	System command A 12-bit bus for transferring commands between processor and the external agent.	Bidirectional
SysCmdPar	System command bus parity Odd parity for the system command bus.	Bidirectional
SysAD(63:0)	System address/data bus A 64-bit bus for transferring addresses and data between R10000 and the external agent.	Bidirectional
SysADChk(7:0)	System address/data check bus An 8-bit ECC bus for the system address/data bus.	Bidirectional
SysVal*	System valid The master of the System interface asserts this signal when it is driving valid information on the system command and system address/data buses.	Bidirectional

Table 3-3 (cont.) System Interface Signals

Signal Name	Description	Type
System State Bus Signals		
SysState(2:0)	System state bus A 3-bit bus used for issuing processor coherency state responses and also additional status indications.	Output
SysStatePar	System state bus parity Odd parity for the system state bus.	Output
SysStateVal*	System state bus valid The processor asserts this signal for one **SysClk** cycle when issuing a processor coherency state response on the system state bus.	Output
System Response Bus Signals		
SysResp(4:0)	System response bus A 5-bit bus used by the external agent for issuing external completion responses.	Input
SysRespPar	System response bus parity Odd parity for the system response bus.	Input
SysRespVal*	System response bus valid The external agent asserts this signal for one **SysClk** cycle when issuing an external completion response on the system response bus.	Input
System Miscellaneous Signals		
SysReset*	System reset The external agent asserts this signal to reset the processor.	Input
SysNMI*	System non-maskable interrupt The external agent asserts this signal to indicate a non-maskable interrupt.	Input
SysCorErr*	System correctable error The processor asserts this signal for one **SysClk** cycle when a correctable error is detected and corrected.	Output
SysUncErr*	System uncorrectable error The processor asserts this signal for one **SysClk** cycle when an uncorrectable tag error is detected.	Output
SysGblPerf*	System globally performed The external agent asserts this signal to indicate that all processor requests have been globally performed with respect to all external agents.	Input
SysCyc*	System cycle The external agent may use this signal to define a virtual System interface clock in a hardware emulation environment.	Input

3.4 Test Interface Signals

Table 3-4 presents the R10000 processor test interface signals.

Table 3-4 Test Interface Signals

Signal Name	Description	Type
JTAG Signals		
JTDI	JTAG serial data input Serial data input.	Input
JTDO	JTAG serial data output Serial data output.	Output
JTCK	JTAG clock Clock input.	Input
JTMS	JTAG mode select Mode select input.	Input
Miscellaneous Test Signals		
TCA	Testability control A (for manufacturing test only) This signal must be tied to **Vss**, through a 100 ohm resistor.	Input
TCB	Testability control B (for manufacturing test only) This signal must be tied to **Vss**, through a 100 ohm resistor.	Input
PLLDis	PLL disable (for manufacturing test only) This signal must be tied to **Vss**.	Input
PLLRC	PLL Control Node (for manufacturing test only) There must be no connection made to this signal.	
PLLSpare(1:4)	These four pins must be tied to **Vss**.	
Spare(1,3)	These two pins must be tied to **Vss**, through a 100 ohm resistor.	
TriState	Tristate Control The system asserts this signal to tristate all outputs and input/output pads except for **SCClk**, **SCCLK***, and **JTDO**.	Input
SelDVCO	Select differential VCO (for manufacturing test only) This signal must be tied to **Vss.**	Input

Cache Organization and Coherency

4

The processor implements a two-level cache structure consisting of separate primary instruction and data caches and a joint secondary cache.

Each cache is two-way set associative and uses a **write back protocol**; that is, two cache blocks are assigned to each set (as shown in Figure 4-1), and a cache store writes data into the cache instead of writing it directly to memory. Some time later this data is independently written to memory.

A write-invalidate cache coherency protocol (described later in this chapter) is supported through a set of cache states and external coherency requests.

4.1 Primary Instruction Cache

The processor has an on-chip 32-Kbyte primary instruction cache (also referred to simply as the *instruction cache*), which is a subset of the secondary cache. Organization of the instruction cache is shown in Figure 4-1.

The instruction cache has a fixed block size of 16 words and is two-way set associative with a **least-recently-used** (LRU) replacement algorithm.†

The instruction cache is indexed with a virtual address and tagged with a physical address.

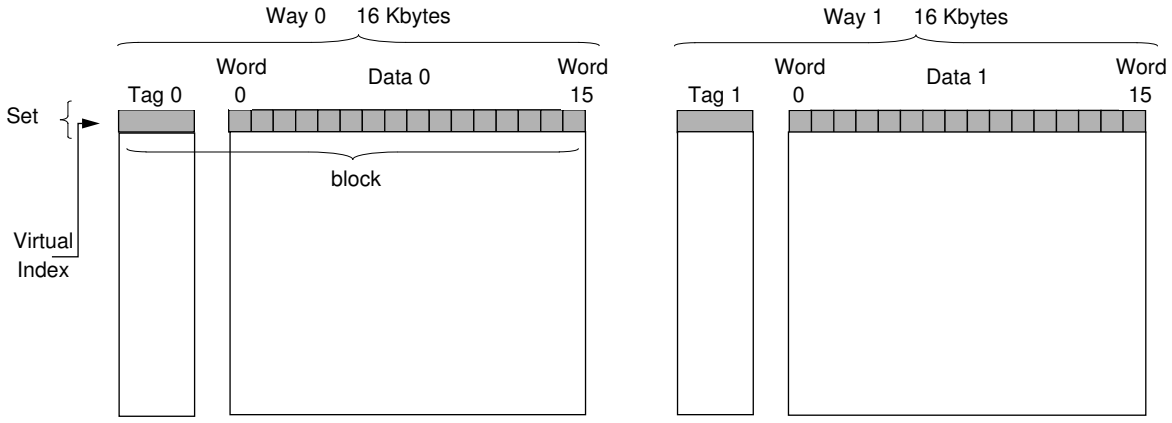

Figure 4-1 Organization of Primary Instruction Cache

Each instruction cache block is in one of the following two states:

- *Invalid*
- *Valid*

† The precise implementation of the LRU algorithm is affected by the speculative execution of instructions.

An instruction cache block can be changed from one state to the other as a result of any one of the following events:

- a primary instruction cache read miss
- subset property enforcement
- any of various CACHE instructions
- external intervention exclusive and invalidate requests

These events are illustrated in Figure 4-2, which shows the primary instruction cache state diagram.

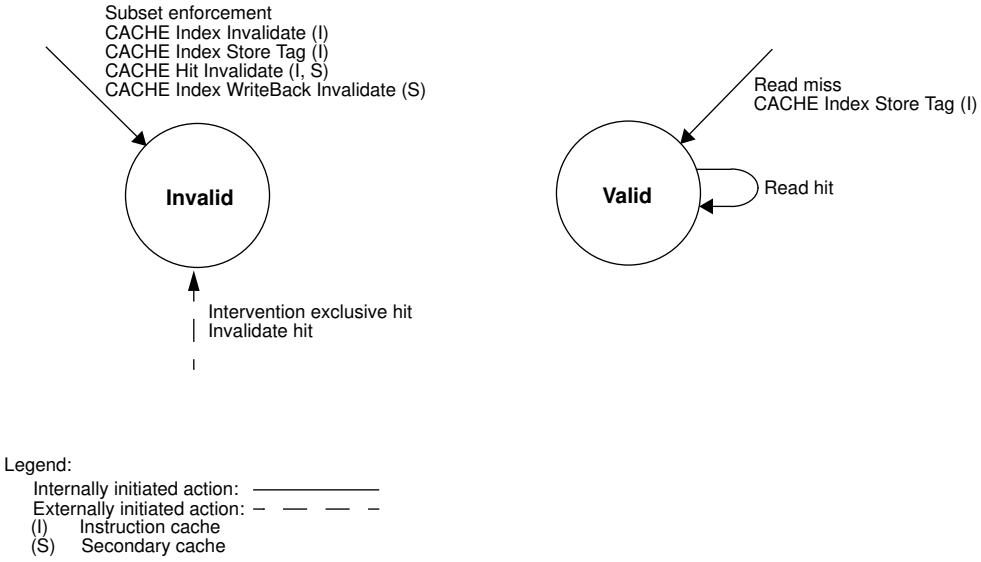

Figure 4-2 Primary Instruction Cache State Diagram

4.2 Primary Data Cache

The processor has an on-chip 32-Kbyte primary data cache (also referred to simply as the *data cache*), which is a subset of the secondary cache. The data cache uses a fixed block size of 8 words and is two-way set associative (that is, two cache blocks are assigned to each set, as shown in Figure 4-3) with an LRU replacement algorithm.†

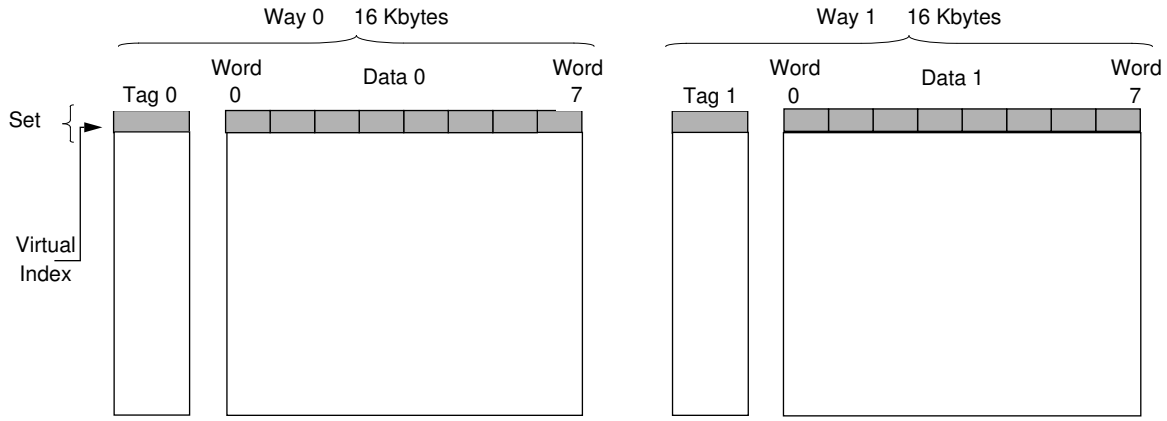

Figure 4-3 Organization of Primary Data Cache

The data cache uses a write back protocol, which means a cache store writes data into the cache instead of writing it directly to memory. Sometime later this data is independently written to memory, as shown in Figure 4-4.

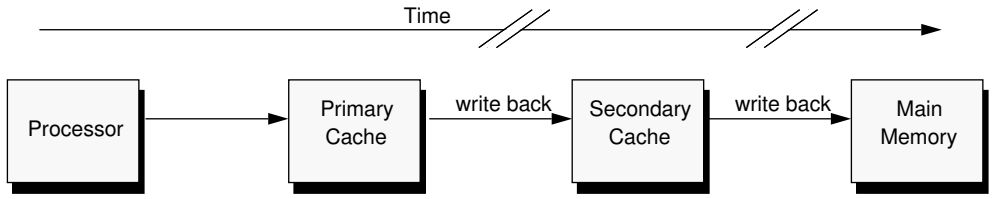

Figure 4-4 Write Back Protocol

Write back from the primary data cache goes to the secondary cache, and write back from the secondary cache goes to main memory, through the system interface. The primary data cache is written back to the secondary cache before the secondary cache is written back to the system interface.

† The precise implementation of the LRU algorithm is affected by the speculative execution of instructions.

The data cache is indexed with a virtual address and tagged with a physical address. Each primary cache block is in one of the following four states:

- *Invalid*
- *CleanExclusive*
- *DirtyExclusive*
- *Shared*

A primary data cache block is said to be *Inconsistent* when the data in the primary cache has been modified from the corresponding data in the secondary cache. The primary data cache is maintained as a subset of the secondary cache where the state of a block in the primary data cache always matches the state of the corresponding block in the secondary cache.

A data cache block can be changed from one state to another as a result of any one of the following events:

- primary data cache read/write miss
- primary data cache write hit
- subset enforcement
- a CACHE instruction
- external intervention shared request
- intervention exclusive request
- invalidate request

These events are illustrated in Figure 4-5, which shows the primary data cache state diagram.

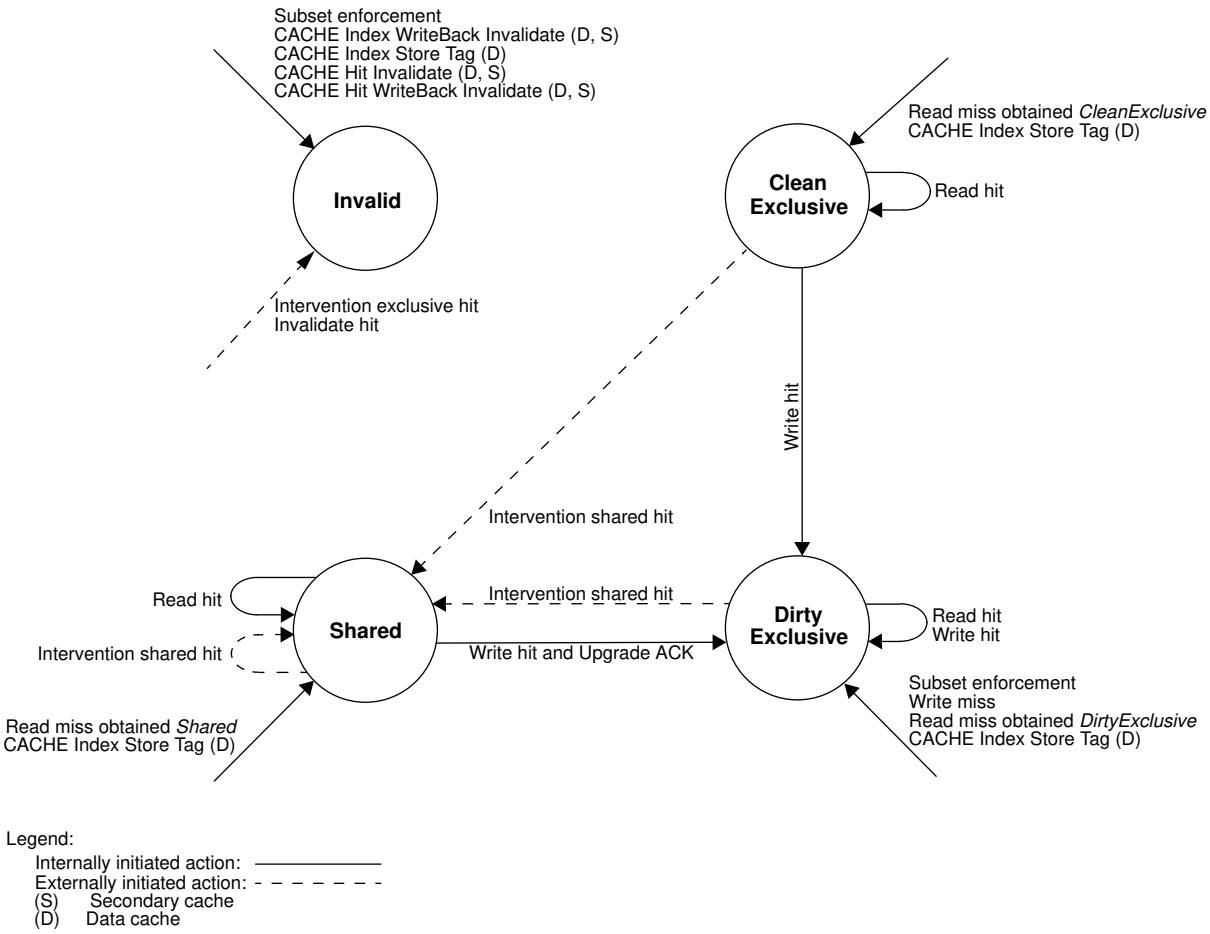

Figure 4-5 Primary Data Cache State Diagram

4.3 Secondary Cache

The R10000 processor must have an external secondary cache, ranging in size from 512 Kbytes to 16 Mbytes, in powers of 2, as set by the **SCSize** mode bit. The **SCBlkSize** mode bit selects a block size of either 16 or 32 words.

The secondary cache is two-way set associative (that is, two cache blocks are assigned to each set, as shown in Figure 4-6) with an LRU replacement algorithm.†

The secondary cache uses a write back protocol, which means a cache store writes data into the cache instead of writing it directly to memory. Some time later this data is independently written to memory.

The secondary cache is indexed with a physical address and tagged with a physical address.

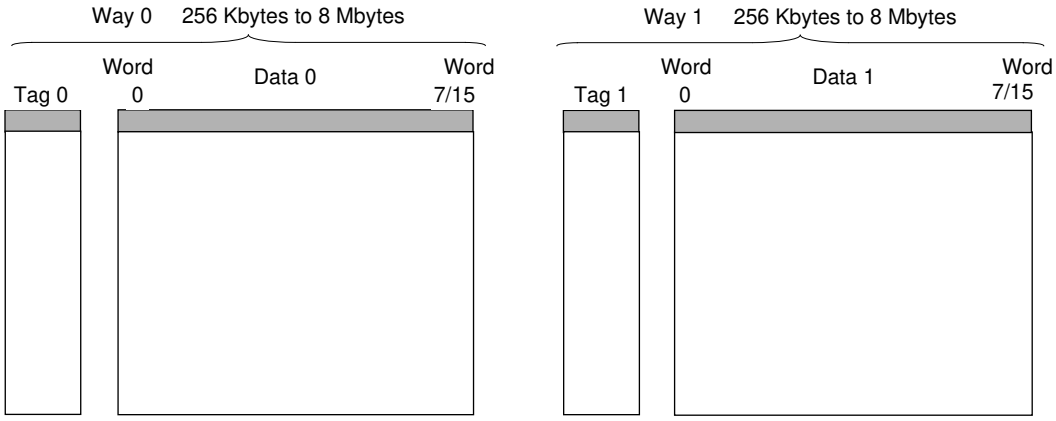

Figure 4-6 Organization of Secondary Cache

Each secondary cache block is in one of the following four states:

- *Invalid*
- *CleanExclusive*
- *DirtyExclusive*
- *Shared*

† The precise implementation of the LRU algorithm is affected by the speculative execution of instructions.

A secondary cache block can be changed from one state to another as a result of any of the following events:

- primary cache read/write miss
- primary cache write hit to a *Shared* or *CleanExclusive* block
- secondary cache read miss
- secondary cache write hit to a *Shared* or *CleanExclusive* block
- a CACHE instruction
- external intervention shared request
- intervention exclusive request
- invalidate request

These events are illustrated in Figure 4-7, which shows the secondary cache state diagram.

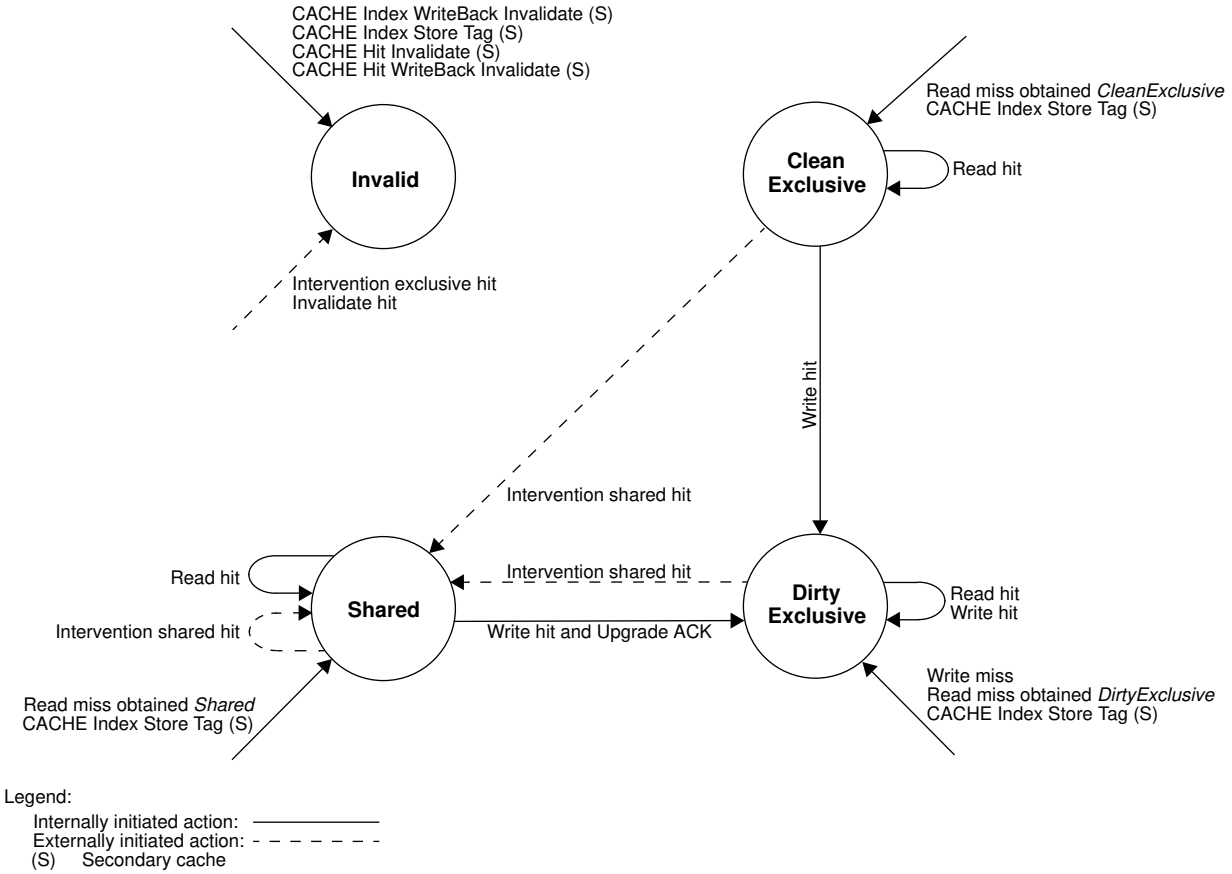

Figure 4-7 Secondary Cache State Diagram

4.4 Cache Algorithms

The behavior of the processor when executing load and store instructions is determined by the cache algorithm specified for the accessed address. The processor supports five different cache algorithms:

- uncached
- cacheable noncoherent
- cacheable coherent exclusive
- cacheable coherent exclusive on write
- uncached accelerated

Cache algorithms are specified in three separate places, depending upon the access:

- the cache algorithm for the mapped address space is specified on a per-page basis by the 3-bit cache algorithm field in the TLB
- the cache algorithm for the *kseg0* address space is specified by the 3-bit *K0* field of the CP0 *Config* register
- the cache algorithm for the *xkphys* address space is specified by **VA[61:59]**

Table 4-1 presents the encoding of the 3-bit cache algorithm field used in the TLB; *EntryLo0* and *EntryLo1* registers; CP0 *Config* register *K0* field for the *kseg0* address space; and **VA[61:59]** for the *xkphys* address space.

Table 4-1 Cache Algorithm Field Encodings

Value	Cache Algorithm
0	Reserved
1	Reserved
2	Uncached
3	Cacheable noncoherent
4	Cacheable coherent exclusive
5	Cacheable coherent exclusive on write
6	Reserved
7	Uncached accelerated

Descriptions of the Cache Algorithms

This section describes the cache algorithms listed in Table 4-1.

Uncached

Loads and stores under the *Uncached* cache algorithm bypass the primary and secondary caches. They are issued directly to the System interface using processor double/single/partial-word read or write requests.

Cacheable Noncoherent

Under the *Cacheable noncoherent* cache algorithm, load and store secondary cache misses result in processor noncoherent block read requests. External agents containing caches need not perform a coherency check for such processor requests.

Cacheable Coherent Exclusive

Under the *Cacheable coherent exclusive* cache algorithm, load and store secondary cache misses result in processor coherent block read exclusive requests. Such processor requests indicate to external agents containing caches that a coherency check must be performed and that the cache block must be returned in an *Exclusive* state.

Cacheable Coherent Exclusive on Write

The *Cacheable coherent exclusive on write* cache algorithm is similar to the *Cacheable coherent exclusive* cache algorithm except that load secondary cache misses result in processor coherent block read shared requests. Such processor requests indicate to external agents containing caches that a coherency check must be performed and that the cache block may be returned in either a *Shared* or *Exclusive* state.

Store hits to a *Shared* block result in a processor upgrade request. This indicates to external agents containing caches that the block must be invalidated.

Uncached Accelerated

The R10000 processor implements a new cache algorithm, *Uncached accelerated*. This allows the kernel to mark the TLB entries for certain regions of the physical address space, or certain blocks of data, as uncached while signalling to the hardware that data movement optimizations are permissible. This permits the hardware implementation to gather a number of uncached writes together, either a series of writes to the same address or sequential writes to all addresses in the block, into an uncached accelerated buffer and then issue them to the system interface as processor block write requests. The *uncached accelerated* algorithm differs from the *uncached* algorithm in that block write gathering is not performed.

There is no difference between an uncached accelerated load and an uncached load. Only word or doubleword stores can take advantage of this mode.

Stores under the *Uncached accelerated* cache algorithm bypass the primary and secondary caches. Stores to identical or sequential addresses are gathered in the uncached buffer, described in Chapter 6, Uncached Buffer.

Completely gathered uncached accelerated blocks are issued to the System interface as processor block write requests. Incompletely gathered uncached accelerated blocks are issued to the System interface using processor double/single-word write requests; this is also described in Chapter 6, Uncached Buffer.

4.5 Relationship Between Cached and Uncached Operations

Uncached and uncached accelerated load and store instructions are executed in order, and non-speculatively. Such accesses are buffered in the uncached buffer by the processor until they can be issued to the System interface.

All uncached and uncached accelerated accesses retain program order within the uncached buffer. The processor continues issuing cached accesses while uncached accesses are queued in the uncached buffer.

NOTE: Cached accesses do not probe the uncached buffer for conflicts.

Buffered uncached stores prevent a SYNC instruction from graduating. However buffered uncached accelerated stores do not prevent a SYNC instruction from graduating. The processor continues issuing cached accesses speculatively and out of order beyond a SYNC instruction that is waiting to graduate.

An uncached load may be used to guarantee that the uncached buffer is flushed of all uncached and uncached accelerated accesses.

A SYNC instruction and the **SysGlbPerf*** signal may be used to guarantee that all cache accesses and uncached stores have been globally performed as described in Chapter 6, SysGblPerf* Signal.

An uncached load followed by a SYNC instruction may be used to guarantee that all cache accesses, uncached accesses, and uncached accelerated accesses have been globally performed.

4.6 Cache Algorithms and Processor Requests

The cache algorithm determines the type of processor request generated for secondary cache load misses, secondary cache store misses, and store hits. Table 4-2 presents the relationship between the cache algorithm and processor requests.

Table 4-2 Cache Algorithms and Processor Requests

Cache Algorithm	Load Miss	Store Miss	Store Hit
Uncached	Double/single/partial-word read	Double/single/partial-word write	NA
Cacheable noncoherent	Noncoherent block read	Noncoherent block read	Upgrade if *Shared*[*]
Cacheable coherent exclusive	Coherent block read exclusive	Coherent block read exclusive	Upgrade if *Shared*[*]
Cacheable coherent exclusive on write	Coherent block read shared	Coherent block read exclusive	Upgrade if *Shared*
Uncached accelerated	Double/single/partial-word read	Gather identical or sequential double/single-word stores in the uncached buffer. Block write for completely gathered blocks. Double/single-word write for incompletely gathered blocks. Partial-word write for partial-word stores.	NA

* Should not occur under normal circumstances. Most systems return the *Exclusive* state for a cacheable noncoherent line; therefore, the *Shared* state is not normal.

4.7 Cache Block Ownership

The processor requires cache blocks to have a single owner at all times. The owner is responsible for providing the current contents of the cache block to any requestor.

The processor uses the following ownership rules:

- The processor assumes ownership of a cache block if the state of the cache block becomes *DirtyExclusive*. For a processor block read request, the processor assumes ownership of the block after receiving the last doubleword of a *DirtyExclusive* external block data response and an external ACK completion response. For a processor upgrade request, the processor assumes ownership of the block after receiving an external ACK completion response.
- The processor gives up ownership of a cache block if the state of the cache block changes to *Invalid*, *CleanExclusive*, or *Shared*.
- *CleanExclusive* and *Shared* cache blocks are always considered to be owned by memory.

Secondary Cache Interface

5

The processor supports a mandatory secondary cache by providing an internal secondary cache controller with a dedicated secondary cache port.

The cache's tag and data arrays each consist of an external bank of industry-standard synchronous SRAM (SSRAM). This SSRAM must have registered inputs and outputs, asynchronous output enables, and use the late write protocol (data is expected one cycle after the address).

5.1 Tag and Data Arrays

The secondary cache consists of a 138-bit wide data array (128 data bits + 9 ECC bits + 1 parity bit) and a 33-bit wide tag array (26 tag bits + 7 ECC bits), as shown in Figure 5-1. ECC is supported for both the data and tag arrays to improve data integrity.

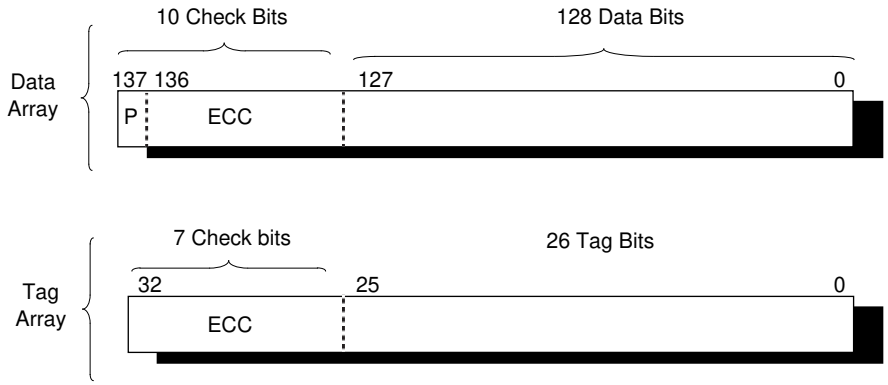

Figure 5-1 Secondary Cache Data and Tag Array

The secondary cache is implemented as a two-way set associative, combined instruction/data cache, which is physically addressed and physically tagged, as described in Chapter 4.

The **SCSize** mode bits specify the secondary cache size; minimum secondary cache size is 512 Kbytes and the maximum secondary cache size is 16 Mbytes, in increments (512 Kbytes, 1 Mbyte, 2 Mbytes, etc.).

The **SCBlkSize** mode bit specifies the secondary cache block size. When negated, the block size is 16 words, and when asserted, the block size is 32 words.

5.2 Secondary Cache Interface Frequencies

The secondary cache interface operates at the frequency of **SCClk**, which is derived from **PClk**. The **SCClkDiv** mode bits select a **PClk** to **SCClk** divisor of 1, 1.5, 2, 2.5, or 3, using the formula described in Chapter 7.

Synchronization between the **PClk** and **SCClk** is performed internally and is invisible to the system. The processor supplies six complementary copies of the secondary cache clock on **SCClk(5:0)** and **SCClk(5:0)***.

5.3 Secondary Cache Indexing

The secondary cache data array width is one quadword, and therefore **PA(3:0)**, which specify a byte within a quadword, are unused by the Secondary Cache interface.

Indexing the Data Array

Since the maximum secondary cache size is 16 Mbytes (8 Mbytes per way), each way requires a maximum of 23 bits to index a byte within a selected way, or 19 bits to index a quadword within a way. Consequently, the processor supplies **PA(22:4)** on **SC(A,B)Addr(18:0)** to index a quadword within a way. The processor selects a secondary cache data way with the **SC(A,B)DWay** signal.

Table 5-1 presents the secondary cache data array index for each secondary cache size; for instance, a 4 Mbyte cache uses the 17 address bits, **PA(20:4)** on **SC(A,B)Addr(16:0)**, concatenated with the way bit, **SC(A,B)DWay**, to index a quadword within a 2 Mbyte way.

Table 5-1 Secondary Cache Data Array Index

SCSize Mode Bits	Secondary Cache Size	Secondary Cache Data Array Index	Physical Address Bits Used
0	512 Kbyte	**SC(A,B)DWay || SC(A,B)Addr(13:0)**	**PA(17:4)**
1	1 Mbyte	**SC(A,B)DWay || SC(A,B)Addr(14:0)**	**PA(18:4)**
2	2 Mbyte	**SC(A,B)DWay || SC(A,B)Addr(15:0)**	**PA(19:4)**
3	4 Mbyte	**SC(A,B)DWay || SC(A,B)Addr(16:0)**	**PA(20:4)**
4	8 Mbyte	**SC(A,B)DWay || SC(A,B)Addr(17:0)**	**PA(21:4)**
5	16 Mbyte	**SC(A,B)DWay || SC(A,B)Addr(18:0)**	**PA(22:4)**

Indexing the Tag Array

The processor supplies the secondary cache tag array's least significant index bit on **SCTagLSBAddr** to support two block sizes without system hardware changes. This signal functions normally as a least significant index bit when the secondary cache block size is 16 words. However, when the secondary cache block size is 32 words, this signal is always negated, since only half as many tags are required. The processor supplies the secondary cache tag way on **SCTWay**.

Table 5-2 presents the secondary cache tag array index for each secondary cache size; it shows each index is composed of a physical address loaded onto **SC(A,B)Addr()**, concatenated with **SCTWay** and **SCTagLSBAddr**.

Table 5-2 Secondary Cache Tag Array Index

SCSize Mode Bits	Secondary Cache Size	Secondary Cache Tag Array Index
0	512 Kbyte	**SCTWay || SC(A,B)Addr(13:3) || SCTagLSBAddr**
1	1 Mbyte	**SCTWay || SC(A,B)Addr(14:3) || SCTagLSBAddr**
2	2 Mbyte	**SCTWay || SC(A,B)Addr(15:3) || SCTagLSBAddr**
3	4 Mbyte	**SCTWay || SC(A,B)Addr(16:3) || SCTagLSBAddr**
4	8 Mbyte	**SCTWay || SC(A,B)Addr(17:3) || SCTagLSBAddr**
5	16 Mbyte	**SCTWay || SC(A,B)Addr(18:3) || SCTagLSBAddr**

For a system design that only supports a secondary cache block size of 32 words, the secondary cache tag array need not use **SCTagLSBAddr** as an index bit.

5.4 Secondary Cache Way Prediction Table

The primary and secondary caches are two-way set associative. However, the implementation of the secondary cache is different than the primary caches.

The primary caches read simultaneously from two separate tag arrays, corresponding to each way in the cache, and then select the data based on the result of two parallel tag compares.

The secondary cache does not use this implementation because it would either require too many pins to read in two full copies of the data and tags, or add latency to externally multiplex two banks of memory. Instead, a way prediction table is used to determine which way to read from first.

The way prediction table is internal to the processor and has 8K one-bit entries, each entry corresponding to a pair of secondary cache blocks. The bit entry indicates which way of the addressed set has been most-recently used (MRU). When the secondary cache is accessed, this prediction bit is used as an address bit; thus the two ways in the secondary cache are shared in the same SSRAM bank.

The secondary cache way prediction table is indexed with a subset of 11 to 13 bits of the physical address, based on both the secondary cache block size, and the secondary cache size, as shown in Table 5-3. "0 || " indicates a zero bit concatenated to the address to pad the index out to a full 13-bits.

Table 5-3 Secondary Cache Way Prediction Table Index

SCSize Mode Bits	Secondary Cache Size	SCBlkSize Mode Bit	Secondary Cache Block Size	Secondary Cache Way Prediction Table Index
0	512 Kbyte	0	16-word	0 || PA(17:6)
		1	32-word	0 || 0 || PA(17:7)
1	1 Mbyte	0	16-word	PA(18:6)
		1	32-word	0 || PA(18:7)
2 to 5	2M to 16 Mbyte	0	16-word	PA(18:6)
		1	32-word	PA(19:7)

Three states are possible in the way prediction table:

- the desired data is in the predicted way
- the desired data is in the non-predicted way
- the desired data is not in the secondary cache

The tags for both ways are read "underneath" the data access cycles in order to discern as rapidly as possible which of these states are valid. This reading is possible because it takes two accesses to read a primary data block (8 words) and 4 cycles to read a primary instruction block (16 words); thus the bandwidth needed to read the tag array twice exists in all cases. Only an extra address pin to the tag array is needed to make this operation parallel and this is implemented by the **SCTWay** pin.

The three possible states are handled in the following manner:

- If, after reading the tags for both ways, it is discovered that the data exists in the predicted way, the processor continues normally.
- If the data exists in the non-predicted way, the processor accesses this non-predicted way in the secondary cache and updates the way prediction table to point to this way.
- If the access misses in both ways of the secondary cache, the data is fetched from the System interface into the non-predicted way, and the way prediction table is updated to point to this way since it is now the most-recently-used.

The way prediction table can cover up to a 2 Mbyte secondary cache when the secondary cache block size is 32 words. If the secondary cache exceeds this size, the accuracy of the way prediction table diminishes slightly. However, the extremely large performance gain made by making the secondary cache larger far outstrips any performance loss in the way prediction table.

5.5 Secondary Cache Tag

The secondary cache tag, transferred on the **SCTag(25:0)** bus, is divided into three fields, as shown in Figure 5-2 below.

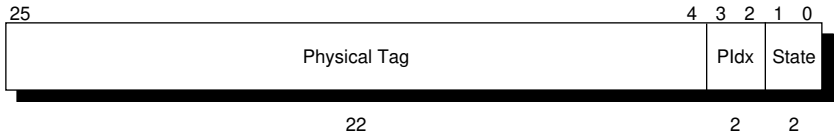

Figure 5-2 Secondary Cache Tag Fields

SCTag(25:4), Physical Tag

The minimum secondary cache size is 512 Kbytes (256 Kbytes per way), so a minimum of 18 bits are required to index a data byte within a selected way. Since the processor supports 40 physical bits, a maximum of 22 bits are required for the physical tag:

```
40 physical address bits - 18 minimum required = 22
```

Consequently, the processor supplies the 22 physical address bits, **PA(39:18)**, on **SCTag(25:4)** for the physical tag.

When the secondary cache is larger than the minimum size, the secondary cache tag array must still maintain the full physical tag supplied by the processor, even though some bits are redundant.

SCTag(3:2), PIdx

Bits **SCTag(3:2)** of the secondary cache tag contain the primary cache index, *PIdx*.

The *PIdx* field contains **VA(13:12)**, which are the two lowest virtual address bits above the minimum 4 Kbyte page size. This field is written into the secondary cache tag during a secondary cache refill. For each processor-initiated secondary cache access, the virtual address bits are compared with the *PIdx* field of the secondary cache tag. If a mismatch occurs, a virtual coherency condition exists and the value of the *PIdx* field is used by internal control logic to purge primary cache locations, so that all primary cache blocks holding valid data have indices known to the secondary cache. This mechanism, unlike that of the R4400 processor, is implemented in hardware. It helps preserve the integrity of cached accesses to a physical address using different virtual addresses, an occurrence called **virtual aliasing**. For each external coherency request, the *PIdx* field of the secondary cache tag provides a mechanism to locate subset lines in the primary caches.

SCTag(1:0), Cache Block State

The lower two bits of the secondary cache tag, **SCTag(1:0)**, contain the cache block state, which can be *Invalid*, *Shared*, *CleanExclusive*, or *DirtyExclusive* as shown in Table 5-4.

Table 5-4 Secondary Cache Tag State Field Encoding

SCTag(1:0)	State
0	*Invalid*
1	*Shared*
2	*CleanExclusive*
3	*DirtyExclusive*

Since the secondary cache tags are updated immediately for stores to the primary data cache, and all caches use a write back protocol, the data in the secondary cache may not always be consistent with data in the primary cache even though the tags always reflect the correct state of a secondary cache block.

5.6 Read Sequences

There are five basic read sequences:

- a 4-word read
- an 8-word read
- a 16-word read
- a 32-word read
- a tag read

4-Word Read Sequence

A 4-word read sequence is performed by a CACHE Index Load Data (S) instruction to read a doubleword of data and 10 check bits from the secondary cache data array.

Figure 5-3 depicts a secondary cache 4-word read sequence. A quadword is read from the index specified by **PA(23:6)**, and the way specified by **VA(0)** of the CACHE instruction.

The doubleword specified by **VA(3)** is then stored into the CP0 *TagHi* and *TagLo* registers, and the corresponding check bits are stored into the CP0 *ECC(9:0)* register. The data may be examined by copying the CP0 *TagHi*, *TagLo*, and *ECC* registers to the general registers with the MTC0 instruction.

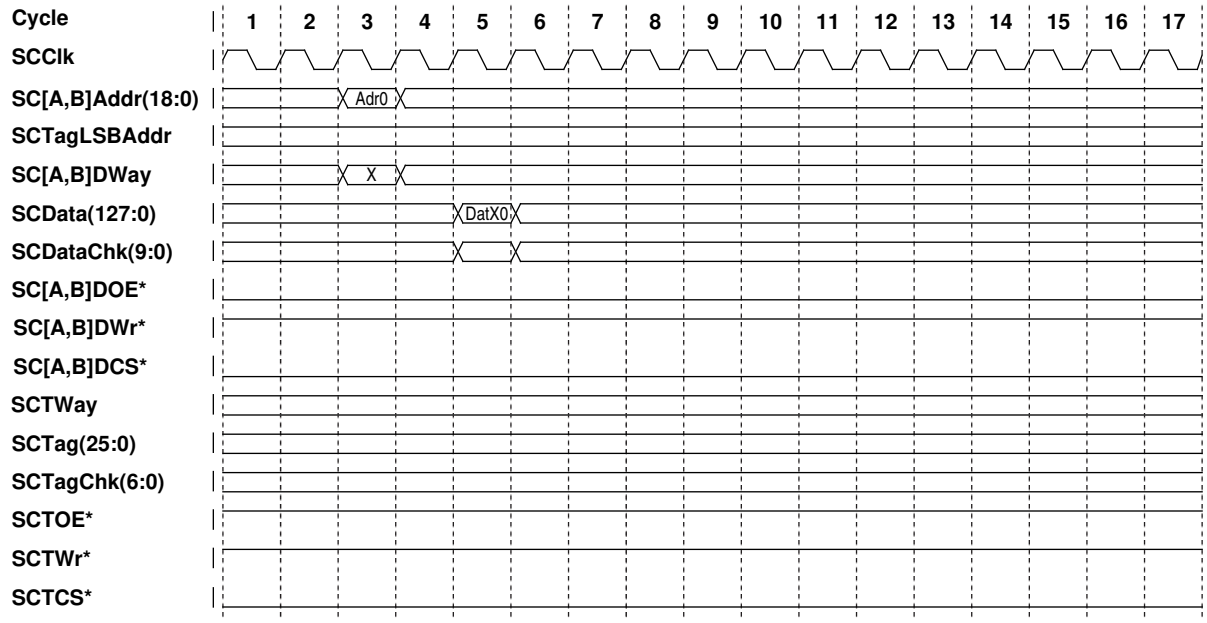

Figure 5-3 4-Word Read Sequence

8-Word Read Sequence

An 8-word read sequence refills the primary data cache from the secondary cache after a primary data cache miss.

Figure 5-4 depicts a secondary cache 8-word read sequence. In it, **SC(A,B)DWay** and **SCTWay** are driven with value X on the first address cycle, which is obtained from the way prediction table.

On the next address cycle, **SCTWay** is complemented in order to read the tag from the non-predicted way of the addressed set. **SC(A,B)DWay** is not changed since it is assumed that the way prediction table is correct and the read is likely to hit in the predicted way.

The tag for the non-predicted way is returned to the processor in the same cycle as the second quadword of data. Reads that miss in the predicted way, but hit in the non-predicted way, are noted by the internal control logic and reissued to the secondary cache as soon as possible.

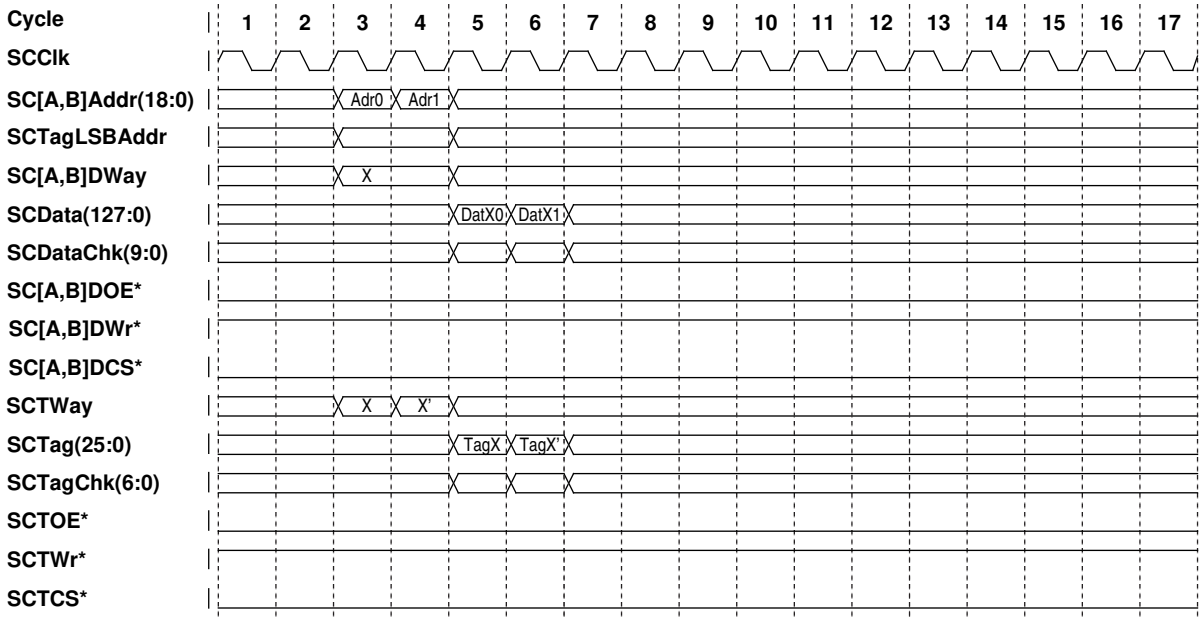

Figure 5-4 8-Word Read Sequence

16 or 32-Word Read Sequence

A 16-word read sequence refills the primary instruction cache from the secondary cache after a primary instruction cache miss. A 16-word read sequence is also performed when the secondary cache block size is 16 words, and a *DirtyExclusive* secondary cache block must be written back to the System interface.

A 32-word read sequence is performed when the secondary cache block size is 32 words, and a *DirtyExclusive* secondary cache block must be written back to the System interface.

Figure 5-5 depicts a secondary cache 16 or 32-word read sequence. This is similar to an 8-word read sequence except that more addresses must be issued, in order to read the appropriate number of quadwords.

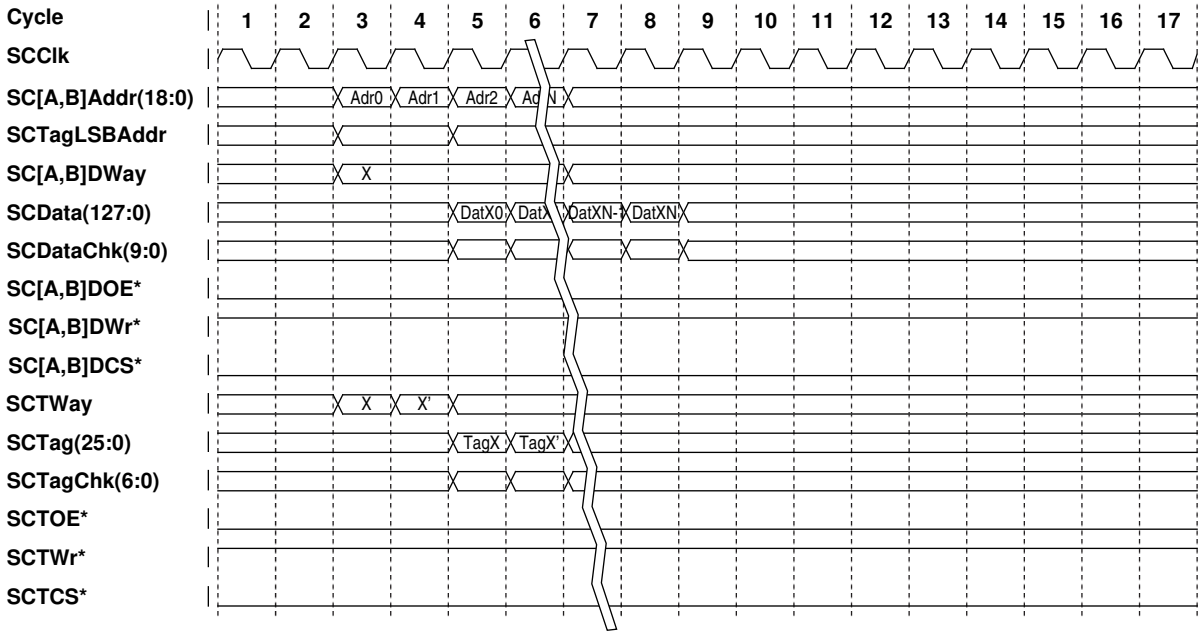

Figure 5-5 16 or 32-Word Read Sequence

Tag Read Sequence

A tag read sequence is performed when the state of a secondary cache block is required, but it is not necessary to access the data array. This sequence is used for the CACHE Index Load Tag (S) instruction.

Figure 5-6 depicts a secondary cache tag read sequence.

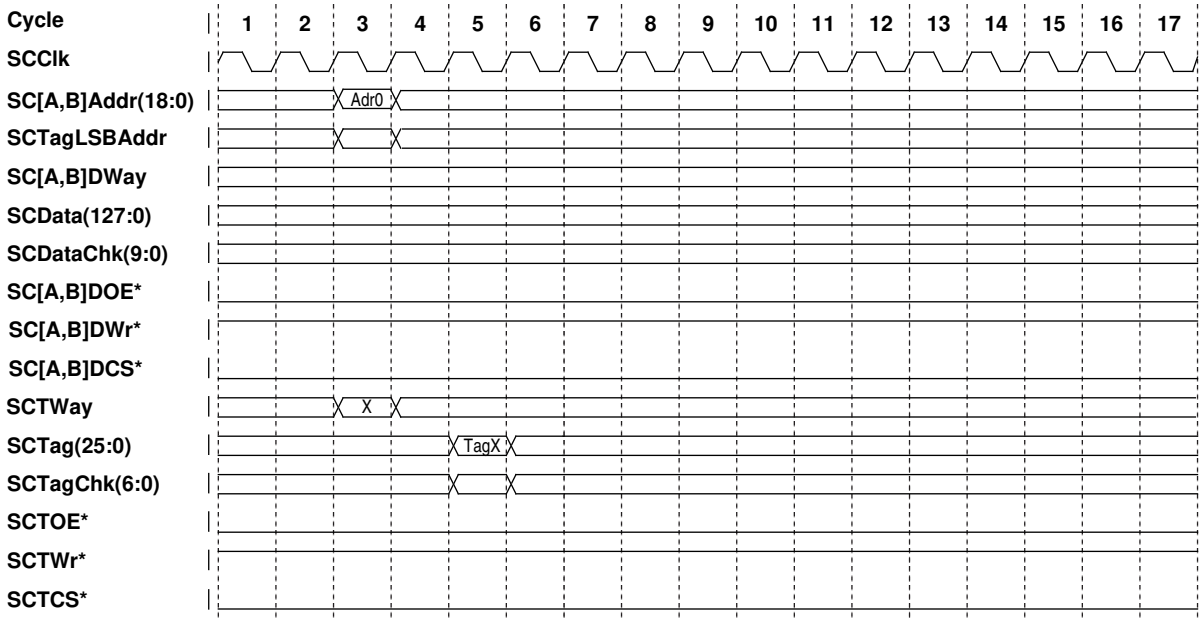

Figure 5-6 Tag Read Sequence

5.7 Write Sequences

There are five basic write sequences:

- a 4-word write.
- an 8-word write
- a 16-word write
- a 32-word write
- a tag write

4-Word Write Sequence

A 4-word write sequence is performed by a CACHE Index Store Data (S) instruction to store a quadword of data and 10 check bits into the secondary cache data array.

Figure 5-7 depicts a secondary cache 4-word write sequence. A quadword is written to the index specified by **PA(23:6)**, and the way specified by **VA(0)** of the CACHE instruction.

A doubleword specified by **VA(3)** is obtained from the CP0 *TagHi* and *TagLo* registers, and the other half of the doubleword is padded to zeros. Normal ECC and parity generation is bypassed and the check field of the data array is written with the contents of the CP0 *ECC(9:0)* register.

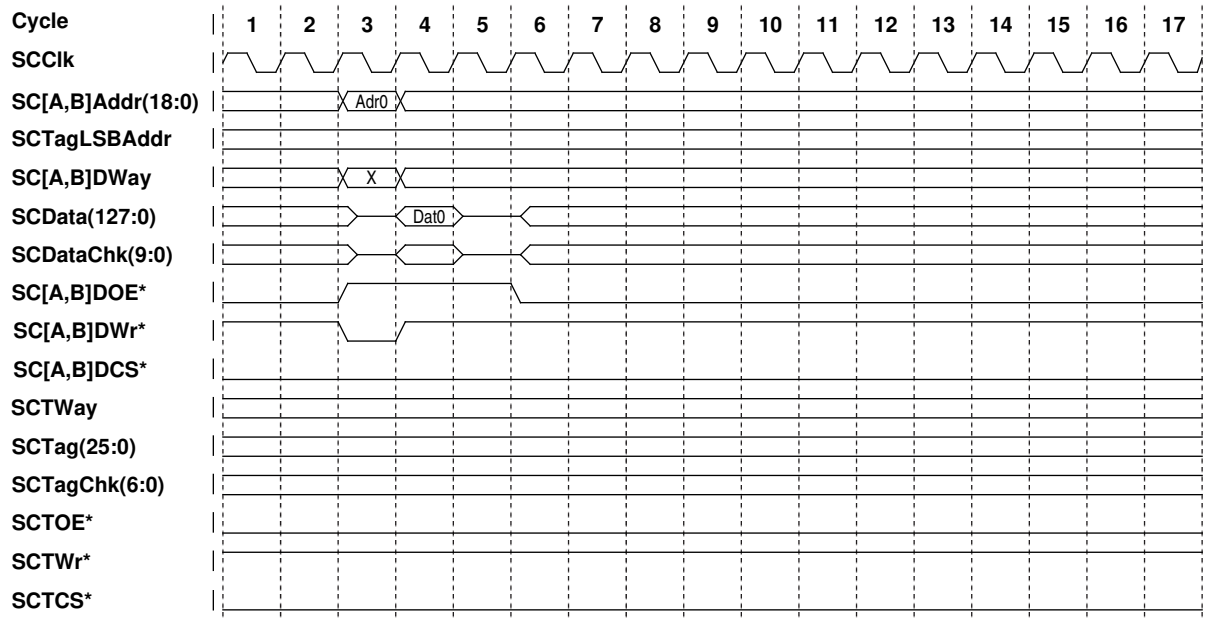

Figure 5-7 4-Word Write Sequence

8-Word Write Sequence

An 8-word write sequence writes back a dirty block from the primary data cache to the secondary cache.

Figure 5-8 depicts a secondary cache 8-word write sequence. **SC(A,B)DWay** are driven with the way bit obtained from the primary data cache tag. The secondary cache tag is not written since it was previously updated when the primary data cache block was modified.

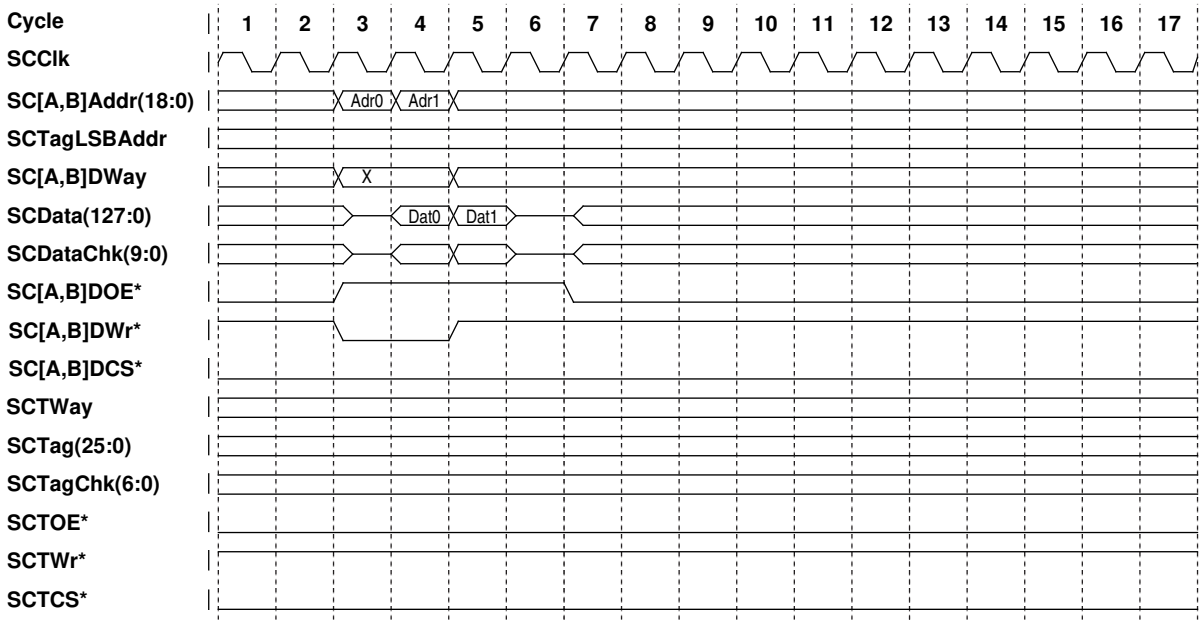

Figure 5-8 8-Word Write Sequence

16 or 32-Word Write Sequence

A 16- or 32-word write sequence refills a secondary cache block from the System interface after a secondary cache miss. A 16-word write sequence is performed when the secondary cache block size is 16 words, and a 32-word write sequence is performed when the secondary cache block size is 32 words.

Figure 5-9 depicts a secondary cache 16 or 32-word write sequence.

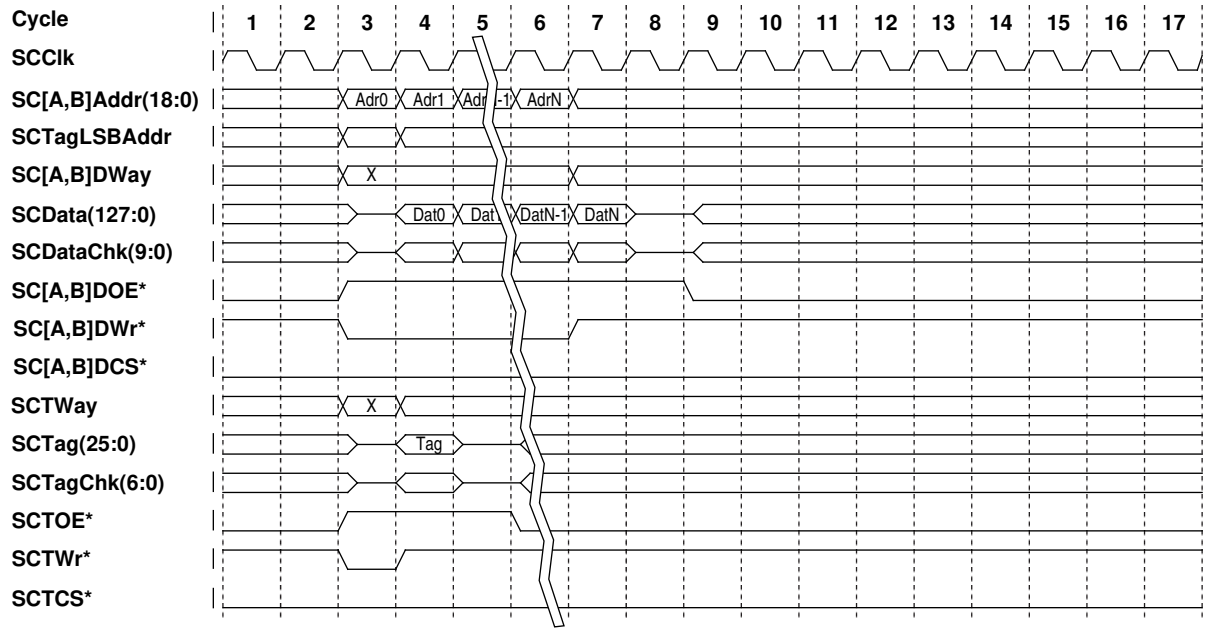

Figure 5-9 16/ 32-Word Write Sequence

Tag Write Sequence

A tag write sequence updates the secondary cache tag array without affecting the data array. This sequence is used for the following:

- to reflect primary cache state changes in the secondary cache
- for external coherency requests
- for the CACHE Index Store Tag (S) instruction

Figure 5-10 depicts the secondary cache tag write protocol.

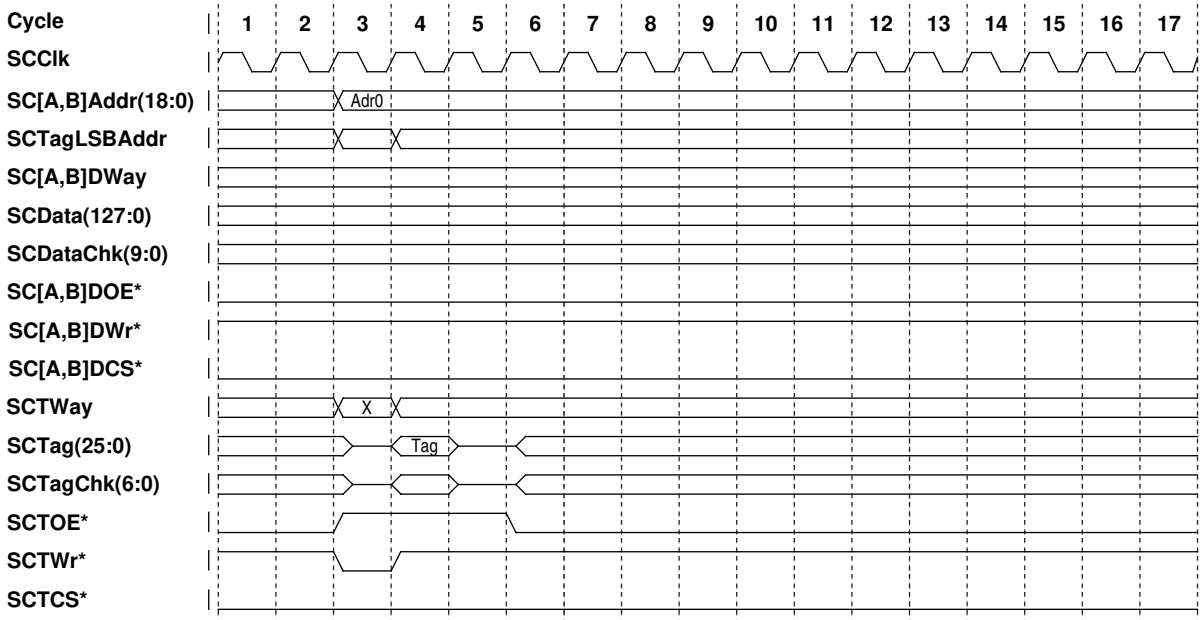

Figure 5-10 Tag Write Sequence

System Interface Operations

6

The R10000 System interface provides a gateway between processor, with its associated secondary cache, and the remainder of the computer system.

For convenience, any device communicating with the processor through the System interface is referred to as the **external agent**.

6.1 Request and Response Cycles

The System interface supports the following request and response cycles:

- **Processor requests** are generated by the processor, when it requires a system resource.
- **External responses** are supplied by an external agent in response to a processor request.
- **External requests** are generated by an external agent when it requires a resource within the processor.
- **Processor responses** are supplied by the processor in response to an external request.

6.2 System Interface Frequencies

The System interface operates at **SysClk** frequency, supplied by the external agent. The internal processor clock, **PClk**, is derived from this same **SysClk**.

The **SysClkDiv** mode bits select a **PClk** to **SysClk** divisor of 1, 1.5, 2, 2.5, 3, 3.5, or 4, using the formula described in Chapter 7.

6.3 Register-to-Register Operation

The System interface is designed to operate in the following register-to-register fashion with the external agent:

- all System interface outputs are sourced directly from registers clocked on the rising edge of **SysClk**
- all System interface inputs directly feed registers that are clocked on the rising edge of **SysClk**

This allows the System interface to run at the highest possible clock frequency.

6.4 System Interface Signals

The R10000 System interface is composed of:

- 3 arbitration signals
- 2 flow-control input signals
- a bidirectional 12-bit command bus
- a bidirectional 64-bit multiplexed address/data bus
- a 3-bit state output bus
- a 5-bit response input bus

6.5 Master and Slave States

At any time, the System interface is either in *master* or *slave* state.

In **master** state, the processor drives the bidirectional System interface signals and is permitted to issue processor requests to the external agent.

In **slave** state, the processor tristates the bidirectional System interface signals and accepts external requests from the external agent.

6.6 Connecting to an External Agent

In a uni- or multiprocessor system using dedicated external agents, the System interface connects to a single external agent.

In a multiprocessor system using the cluster bus (see below), the system can connect up to four R10000 processors to an external agent. This external agent is referred to as the **cluster coordinator**.

6.7 Cluster Bus

In a multiprocessor system using the cluster bus, the cluster coordinator performs the cluster bus arbitration and data flow management. The arbitration scheme assures that either one of the processors or the cluster coordinator is master at any given time, while the remaining devices are slave.

A processor request issued by the master processor is observed as an external request by all slave R10000 processors, as shown in Figure 6-1. Similarly, a processor coherency data response issued by a master processor is observed as an external data response by the slave processors.

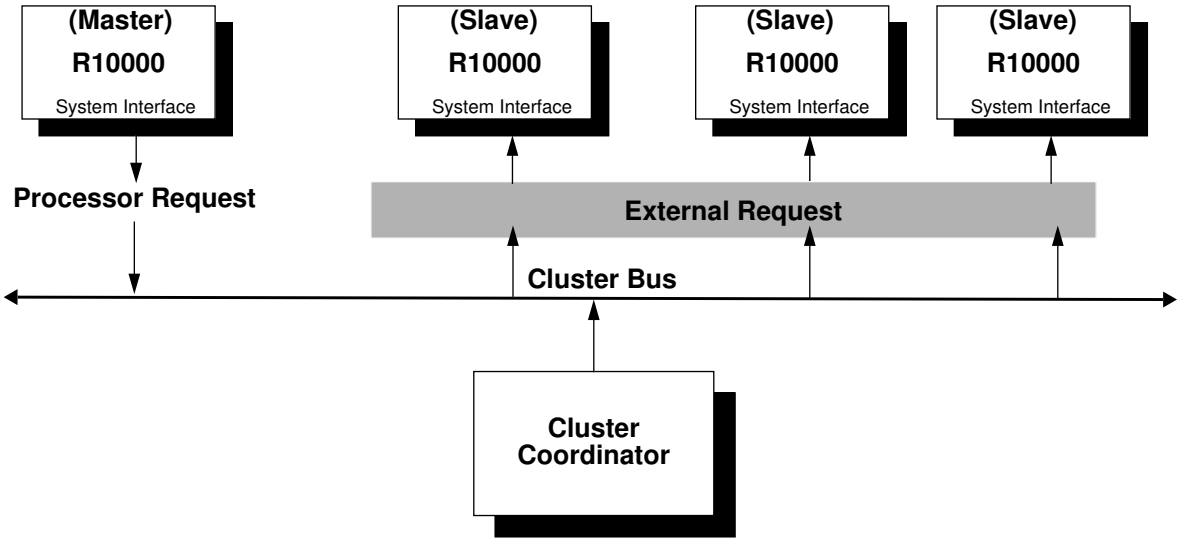

Figure 6-1 Processor Request Master/Slave Status

In a multiprocessor system using the cluster bus, a mode bit specifies whether processor coherent requests are to target the external agent only, or all processors and the external agent. This allows systems with efficient snoopy, duplicate tag, or directory-based coherency protocols to be created.

6.8 System Interface Connections

The major System interface connections required for various system configurations are presented in this section.

Uniprocessor System

Figure 6-2 shows the major System interface connections required for a typical uniprocessor system.

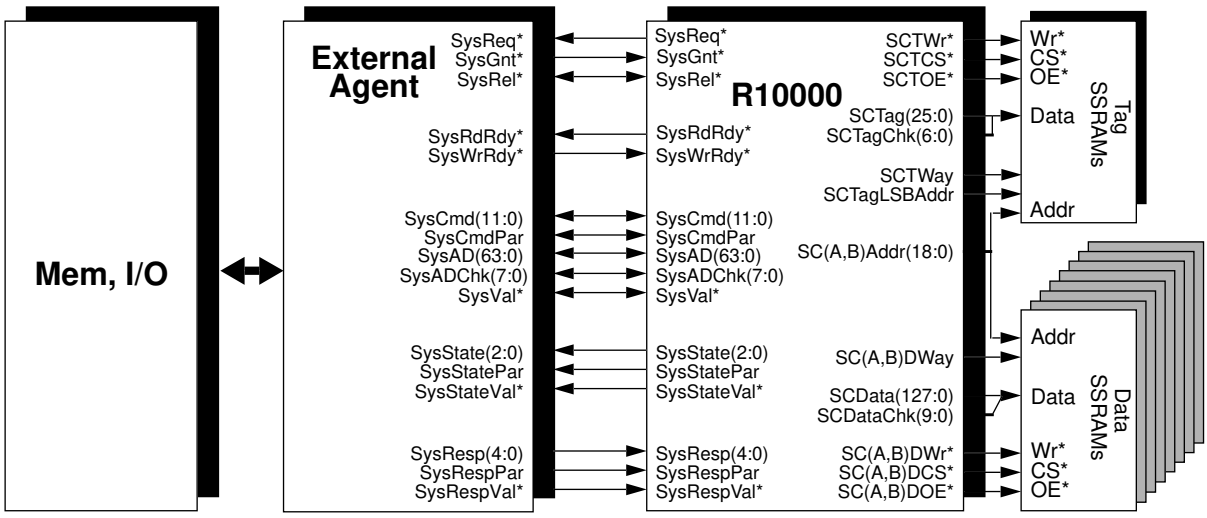

Figure 6-2 System Interface Connections for Uniprocessor System

Multiprocessor System Using Dedicated External Agents

Figure 6-3 shows the major System interface connections required for a typical multiprocessor system using dedicated external agents.

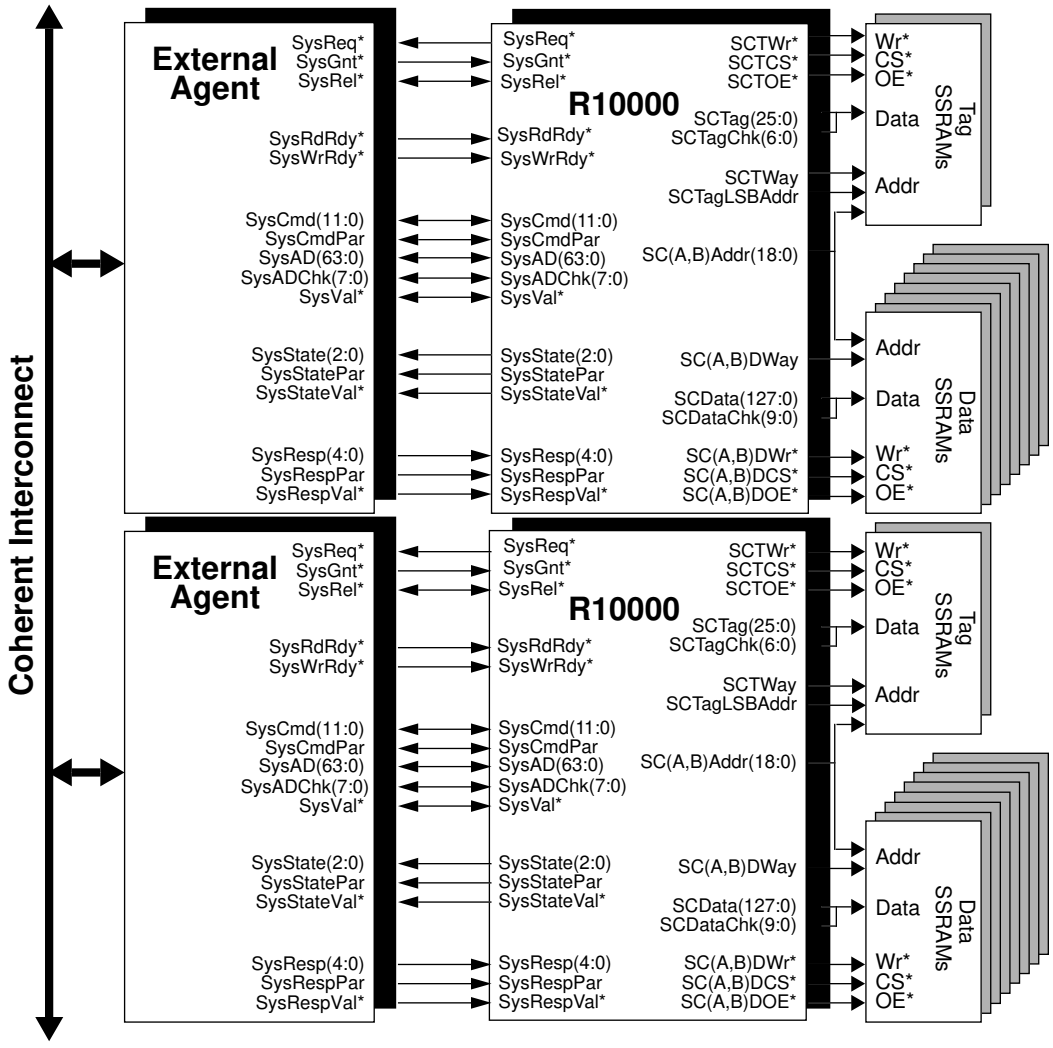

Figure 6-3 System Interface Connections for Multiprocessor using Dedicated External Agents

Multiprocessor System Using the Cluster Bus

Figure 6-4 presents the major System interface connections required for a typical multiprocessor system using the cluster bus.

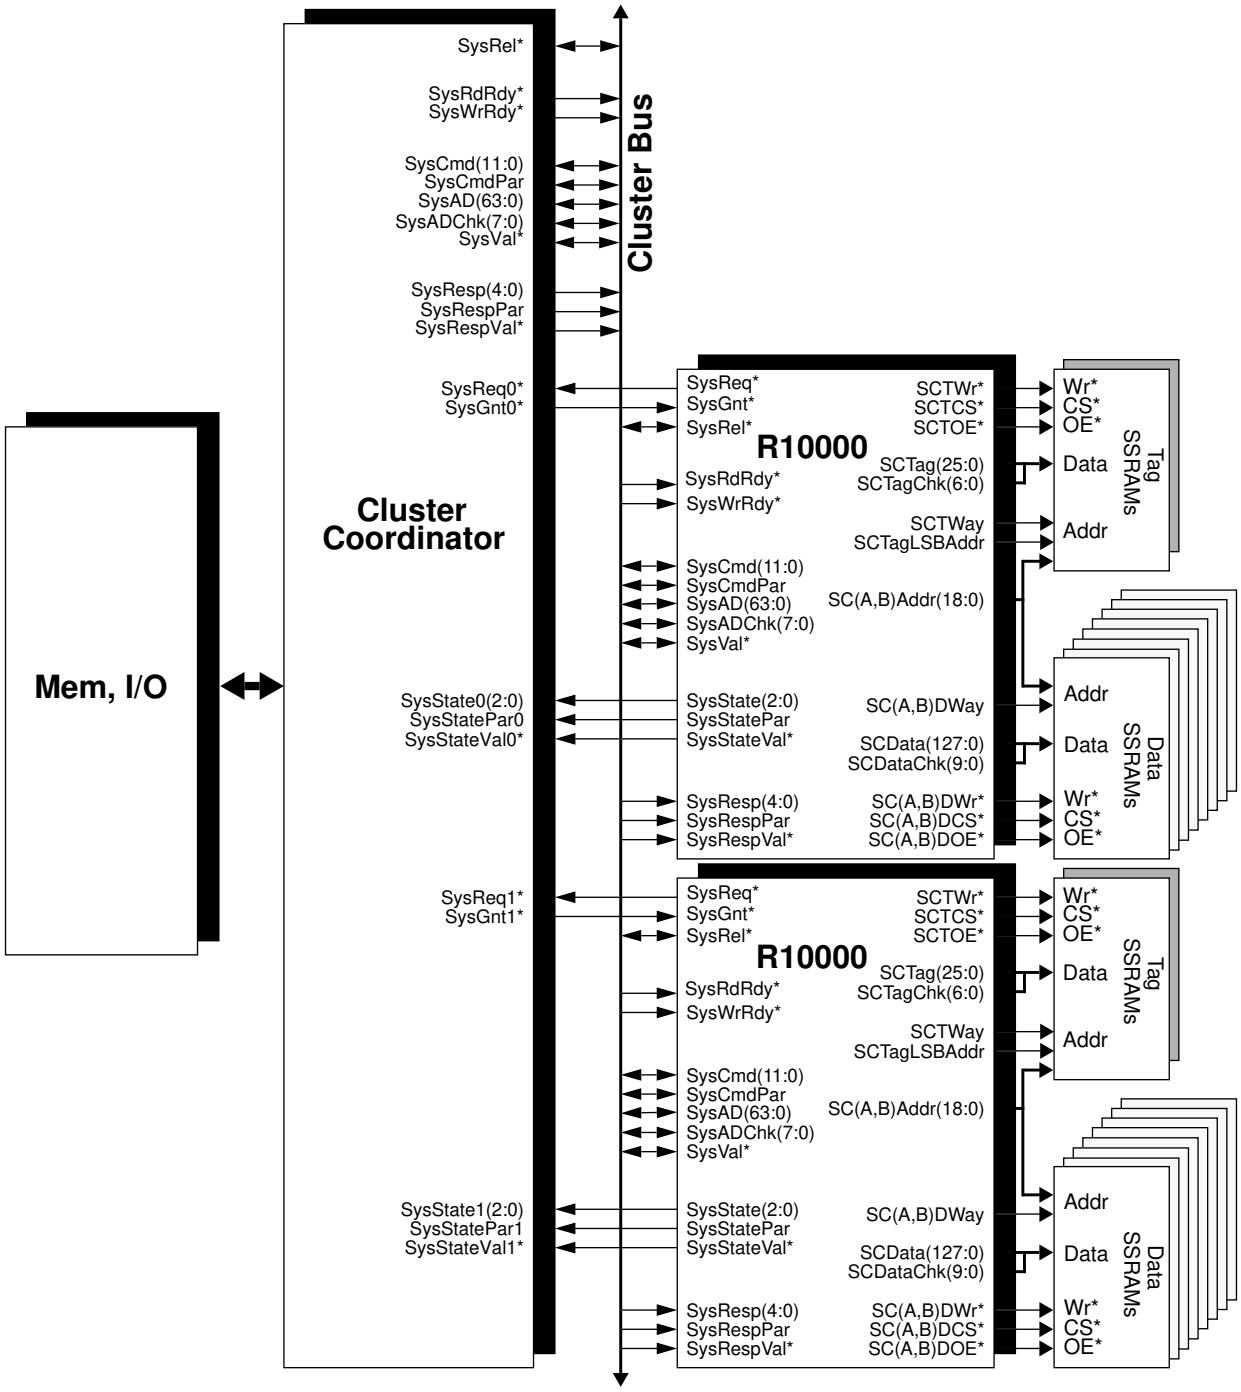

Figure 6-4 System Interface Connections for Multiprocessor Using the Cluster Bus

6.9 System Interface Requests and Responses

The System interface supports the following:

- processor request
- external response
- external request
- processor response

The following sections describe these request and response types, and their operations.

Processor Requests

Processor requests are generated by the processor when it requires a system resource. The following processor requests are supported:

- coherent block read shared request
- coherent block read exclusive request
- noncoherent block read request
- double/single/partial-word read request
- block write request
- double/single/partial-word write request
- upgrade request
- eliminate request

Processor write and eliminate requests do not require or expect a response by the external agent. However, if an external agent detects an error in a processor write or eliminate request, it may use an interrupt to signal the processor. It is not possible to generate precise exceptions for processor write and eliminate requests for which an external agent detects an error.

Processor read and upgrade requests require some type of response by the external agent.

External Responses

External responses are supplied by an external agent or another processor in response to a processor request. The following external responses are supported:

- block data response
- double/single/partial-word data response
- completion response

External Requests

External requests are generated by an external agent when it requires a resource within the processor. The following external requests are supported:

- intervention shared request
- intervention exclusive request
- allocate request number request
- invalidate request
- interrupt request

External intervention and invalidate requests require some type of response by the processor.

Processor Responses

Processor responses are supplied by the processor in response to an external request. The following processor responses are supported:

- coherency state response
- coherency data response

Outstanding Requests and Request Numbers

The processor allows requests and corresponding responses to be split transactions, which enables additional processor and external requests to be issued while waiting for a prior response. The System interface supports a request number field to link requests with their corresponding responses, so responses can be returned out of order.

The processor allows a maximum of eight outstanding requests on the System interface through a 3-bit request number. These outstanding requests may be composed of any mix of processor and external requests.

An individual processor (as opposed to the System interface, above) supports a maximum of four outstanding processor requests at any given time.

Request and Response Relationship

The relationship between processor and external requests, and their acceptable responses, is presented in Table 6-1. The data in this table is given with respect to a single processor, in either a uni- or multiprocessor system (independent of cluster/non-cluster configuration).

Table 6-1 Request and Response Relationship

Request	Acceptable Response Sequences
Processor block read request	External NACK or ERR completion response
	0 or more external block data responses followed by a final external block data response with a coincidental or subsequent external ACK, NACK, or ERR completion response
Processor double/single/partial-word read request	External NACK or ERR completion response
	0 or more external double/single/partial-word data responses followed by a final external double/single/partial-word data response with a coincidental or subsequent external ACK, NACK, or ERR completion response
Processor block write request	None
Processor double/single/partial-word write request	None
Processor upgrade request	External ACK, NACK, or ERR completion response
	0 or more external block data responses followed by a final external block data response with a coincidental or subsequent external ACK, NACK, or ERR completion response
Processor eliminate request	None
External intervention request	Processor coherency state response followed by processor coherency data response (if *DirtyExclusive)* with a coincidental or subsequent external ACK, NACK, or ERR completion response*
External allocate request number request	External ACK, NACK, or ERR completion response*
External invalidate request	Processor coherency state response followed by external ACK, NACK, or ERR completion response*
External interrupt request	None

* External completion response is required to free the request number.

6.10 System Interface Buffers

The processor contains the following five buffers to enhance the performance of the System interface and to simplify the system design:

- incoming buffer
- cached request buffer
- outgoing buffer
- uncached buffer
- cluster request buffer

These buffers are described in the following sections.

Incoming Buffer

The System interface contains an incoming buffer for external block and double/single/partial-word data responses. The four 32-word entries of the incoming buffer correspond to the four possible outstanding processor requests. Block data in each entry of the incoming buffer is stored in subblock order, beginning with a quadword-aligned address.

The incoming buffer eliminates the need for the processor to flow-control the external agent that is providing the external data responses. Regardless of the cache bandwidth or internal resource availability, the external agent may supply external data response data for all outstanding read and upgrade requests at the maximum System interface data rate.

The external agent may issue any number of external data responses for a particular request number before issuing a corresponding external completion response. An external data response remains in the incoming buffer until a corresponding external completion response is received. A former buffered external data response for a particular request number is over-written by a subsequent external data response for the same request number.

An external ACK completion response frees buffered data to be forwarded to the caches and other internal resources while an external NACK or ERR completion response purges any corresponding buffered data. For minimum latency, the external agent should issue an external ACK completion response coincident with the first doubleword of an external data response.

External coherency requests that target blocks residing in the incoming buffer are stalled until the incoming buffer data is forwarded to the secondary cache, and the instruction that caused the secondary miss is satisfied.

Each doubleword of the incoming buffer has an Uncorrectable Error flag. When an external data response provides a doubleword, the processor asserts the corresponding incoming buffer Uncorrectable Error flag if the data quality indicator, **SysCmd[5]**, is asserted, or if an uncorrectable ECC error is encountered on the system address/data bus and the ECC check indication on **SysCmd[0]** is asserted.

When the processor forwards block data from an incoming buffer entry after receiving an external ACK completion response, the associated incoming buffer Uncorrectable Error flags are checked, and if any are asserted, a single Cache Error exception is posted. When the processor forwards double/single/partial-word data from an incoming buffer entry after receiving an external ACK completion response, the associated incoming buffer Uncorrectable Error flag is checked, and if asserted, a Bus Error exception is posted.

Cached Request Buffer

The System interface contains a four-entry cached request buffer. This buffer holds the status of the four possible outstanding processor cached requests, including processor block read and upgrade requests. The relative order of the requests is maintained in the cached request buffer.

External coherency requests probe the cached request buffer to detect conflict conditions.

Outgoing Buffer

The System interface contains a five-entry outgoing buffer to provide buffering for the following:

- *DirtyExclusive* blocks that are cast out of the secondary cache because of a block replacement
- various CACHE instructions
- an external intervention request.

Four 32-word *typical* entries are associated with the four possible outstanding processor cached requests allowed by the processor. One 32-word *special* entry is reserved for external intervention requests only. The data is stored in each entry of the outgoing buffer in sequential order, beginning with a secondary cache block-aligned address.

An instruction or data access that misses in the secondary cache but targets an entry in the outgoing buffer is stalled until the outgoing buffer entry is issued as a processor block write request or coherency data response to the System interface bus.

External coherency requests probe the four typical outgoing buffer entries, with the following results:

- If an external intervention request hits a typical entry, that entry is converted from a processor block write request to a processor coherency data response.
- If an external invalidate request hits a typical outgoing buffer entry, that entry is deleted.
- If an external intervention request does not hit a typical outgoing buffer entry, but hits a *DirtyExclusive* block in the secondary cache, the special outgoing buffer entry is used to buffer the processor coherency data response.

A typical outgoing buffer entry containing a block write is ready for issue to the System interface bus when the first quadword is received from the secondary cache. The processor allows data to stream from the secondary cache to the System interface bus through the outgoing buffer.

An outgoing buffer entry containing a coherency data response for a secondary cache block-aligned external intervention request is ready for issue to the System interface bus when the first quadword is received from the secondary cache. In this case, the processor allows data to stream from the secondary cache to the System interface bus through the outgoing buffer. An outgoing buffer entry containing a coherency data response for a non-secondary cache block-aligned external intervention request is ready for issue to the System interface bus when all quadwords are received from the secondary cache.

Each quadword of the outgoing buffer maintains an Uncorrectable Error flag. If an uncorrectable error is encountered while a block is being cast out of the secondary cache, the associated outgoing buffer quadword Uncorrectable Error flag is asserted. When the processor empties an outgoing buffer entry by issuing a processor block write or coherency data response, the outgoing buffer Uncorrectable Error flags are reflected by the data quality indication on **SysCmd[5]**.

Uncached Buffer

The System interface contains an uncached buffer to provide buffering for uncached and uncached accelerated load and store operations. All operations retain program order within the uncached buffer.

The uncached buffer is organized as a 4-entry FIFO followed by a 2-entry gatherer. Each gathered entry has a capacity of 16 or 32 words, as specified by the **SCBlkSize** mode bit.

The uncached buffer begins gathering when an uncached accelerated double or singleword block-aligned store is executed. Gathering continues if the subsequent uncached operation executed is an uncached accelerated double or singleword store to a sequential or identical address. Once a second uncached accelerated store is gathered, the gathering mode is determined to be sequential or identical. Gathering continues until one of the following conditions occurs:

- a complete block is gathered
- an uncached or uncached accelerated load is executed
- an uncached or uncached accelerated partial-word store is executed
- an uncached store is executed
- a change in the current gathering mode is observed
- a change in the uncached attribute is observed

When gathering terminates, the data is ready for issue to the System interface bus. A processor uncached accelerated block write request is used to issue a completely gathered uncached accelerated block. One or more disjoint processor uncached accelerated double or singleword write requests are used to issue an incompletely gathered uncached accelerated block.

When gathering in an identical mode, uncached accelerated double or singleword stores may be freely mixed. The uncached buffer packs the associated data into the gatherer. When gathering in sequential mode, uncached accelerated singleword stores must occur in pairs, to prevent an address error exception. For instance, SW, SW, SD, SW, SW is legal. SD, SW, SD, is not.

External coherency requests have no effect on the uncached buffer.

CACHE instructions have no effect on the uncached buffer. SYNC instructions are prevented from graduating if an uncached store resides in the uncached buffer.

Cluster Request Buffer

The System interface contains an 8-entry cluster request buffer. This buffer maintains the status of the eight possible outstanding requests on the System interface. When the System interface is in master state, and it issues the address cycle of processor read or upgrade request, the processor places an entry into the cluster request buffer. When the System interface is in slave state, and an external agent issues an external coherency or allocate request number request, it places an entry into the cluster request buffer.

Once an entry is placed into the cluster request buffer, the associated request number transitions from *free* to *busy*. An entry remains busy until the processor receives an external completion response. Processor requests that are ready to be issued to the System interface bus probe the cluster request buffer to detect conflict conditions.

6.11 System Interface Flow Control

The System interface supports a maximum *request rate* of one request per **SysClk** cycle, and a maximum *data rate* of one doubleword per **SysClk** cycle.

Various flow control mechanisms are provided to limit these rates, as described below.

Processor Write and Eliminate Request Flow Control

The processor can only issue a processor write or eliminate request if:

- the System interface is in master state
- **SysWrRdy*** was asserted two **SysClk** cycles previously

Processor Read and Upgrade Request Flow Control

The processor can only issue a processor read or upgrade request if:

- the System interface is in master state
- **SysRdRdy*** was asserted two **SysClk** cycles previously
- the maximum number of outstanding processor requests specified by the **PrcReqMax** mode bits is not exceeded
- there is a free request number

Processor Coherency Data Response Flow Control

The processor can only issue a processor coherency data response if:

- the System interface is in master state
- **SysWrRdy*** was asserted two **SysClk** cycles previously

External Request Flow Control

When the System interface is in *Slave* state, it is capable of accepting external requests. An external agent may issue external requests in adjacent **SysClk** cycles.

External Data Response Flow Control

Since the processor has an incoming buffer, an external agent may supply external data response data in adjacent **SysClk** cycles, without regard to cache bandwidth or internal resource availability.

6.12 System Interface Block Data Ordering

During block data transfers on the System interface **SysAD[63:0]** bus, even doublewords (Dat0, Dat2,...) always correspond to **SCData[127:64]**, and odd doublewords (Dat1, Dat3,...) always correspond to **SCData[63:0]**.

External Block Data Responses

During the address cycle of processor block read and upgrade requests, the processor specifies a quadword-aligned address. The processor expects the external block data response to be supplied in a subblock order sequence, beginning at the specified quadword-aligned address.

Processor Coherency Data Responses

The address of external intervention requests are internally aligned by the processor to a quadword address. If the processor determines that it must issue a processor coherency data response, it supplies the data in a subblock order sequence beginning at the quadword-aligned address specified by the corresponding external coherency request.

Processor Block Write Requests

During the address cycle of processor block write requests, the processor specifies a cache block-aligned address. During the subsequent data cycles for typical processor block write requests, the processor supplies the data in sequence, beginning with the secondary cache block-aligned address.

6.13 System Interface Bus Encoding

This section presents the encoding of the following four System interface buses:

- **SysCmd[11:0]**
- **SysAD[63:0]**
- **SysState[2:0]**
- **SysResp[4:0]**

SysCmd[11:0] Encoding

This section describes address and data cycle encodings for the system command bus, **SysCmd[11:0]**.

SysCmd[11] Encoding

When **SysVal*** is asserted, **SysCmd[11]** indicates whether the **SysAD[63:0]** bus represents an address or a data cycle, as shown in Table 6-2.

Table 6-2 Encoding of **SysCmd[11]**

SysCmd[11]	Data/Address Cycle Indication
0	**SysAD[63:0]** address cycle
1	**SysAD[63:0]** data cycle

SysCmd[10:0] Address Cycle Encoding

During the address cycle of processor read and upgrade requests, **SysCmd[10:8]** contain the request number, as shown in Table 6-3. The request number provides a mechanism to associate an external response with the corresponding processor request.

Table 6-3 Encoding of **SysCmd[10:8]** *for Processor Read and Upgrade Requests*

SysCmd[10:8]	Request Number

During the address cycle of processor requests, **SysCmd[7:5]** contain the command, as shown in Table 6-4.

Table 6-4 Encoding of **SysCmd[7:5]** *for Processor Requests*

SysCmd[7:5]	Command
0	Coherent block read shared
1	Coherent block read exclusive
2	Noncoherent block read
3	Double/single/partial-word read
4	Block write
5	Double/single/partial-word write
6	Upgrade
7	Special

During the address cycle of processor read requests, **SysCmd[4:3]** contain the read cause indication, as shown in Table 6-5. This information is useful in handling the associated external response.

Table 6-5 Encoding of **SysCmd[4:3]** *for Processor Read Requests*

SysCmd[4:3]	Read Cause Indication
0	Instruction access
1	Data typical access
2	Data LL/LLD access
3	Data prefetch access

During the address cycle of processor write requests, **SysCmd[4:3]** contain the write cause indication, as shown in Table 6-6. This information is useful in handling the associated write data.

Table 6-6 Encoding of **SysCmd[4:3]** *for Processor Write Requests*

SysCmd[4:3]	Write Cause Indication
0	Reserved
1	Data typical access
2	Data uncached accelerated sequential access
3	Data uncached accelerated identical access

During the address cycle of processor upgrade requests, **SysCmd[4:3]** contain the upgrade cause indication, as shown in Table 6-7. This information useful in handling the associated external response.

Table 6-7 Encoding of **SysCmd[4:3]** *for Processor Upgrade Requests*

SysCmd[4:3]	Upgrade Cause Indication
0	Reserved
1	Data typical access
2	Data SC/SCD access
3	Data prefetch access

During the address cycle of processor special requests, **SysCmd[4:3]** contain the processor special cause indication, as shown in Table 6-8. This information differentiates between the various processor special requests.

Table 6-8 Encoding of **SysCmd[4:3]** *for Processor Special Requests*

SysCmd[4:3]	Special Cause Indication
0	Reserved
1	Eliminate
2	Reserved
3	Reserved

During the address cycle of processor block read, typical block write, upgrade, and eliminate requests, **SysCmd[2:1]** contain the secondary cache block former state, as shown in Table 6-9. This information may be useful for system designs implementing a duplicate tag or a directory-based coherency protocol.

Table 6-9 Encoding of **SysCmd[2:1]** *for Processor Block Read/Write, Upgrade, Eliminate Requests*

SysCmd[2:1]	Secondary Cache Block Former State
0	*Invalid*
1	*Shared*
2	*CleanExclusive*
3	*DirtyExclusive*

During the address cycle of processor double/single/partial-word read and write requests, **SysCmd[2:0]** contain the data size indication, as shown in Table 6-10.

Table 6-10 Encoding of **SysCmd[2:0]** *for Processor Double/Single/Partial-Word Read/Write Requests*

SysCmd[2:0]	Data Size Indication
0	One byte valid (Byte)
1	Two bytes valid (Halfword)
2	Three bytes valid (Tribyte)
3	Four bytes valid (Word)
4	Five bytes valid (Quintibyte)
5	Six bytes valid (Sextibyte)
6	Seven bytes valid (Septibyte)
7	Eight bytes valid (Doubleword)

During the address cycle of external intervention and invalidate requests, **SysCmd[10:8]** contain the request number, as shown in Table 6-11. The request number provides a mechanism to associate a potential processor coherency data response with the corresponding external coherency request.

Table 6-11 Encoding of **SysCmd[10:8]** *for External Intervention and Invalidate Requests*

SysCmd[10:8]	Request Number

During the address cycle of external requests, **SysCmd[7:5]** contain the command, as shown in Table 6-12.

Table 6-12 Encoding of **SysCmd[7:5]** *for External Requests*

SysCmd[7:5]	Command
0	Intervention shared
1	Intervention exclusive
2	Allocate request number
3	Allocate request number
4	NOP
5	NOP
6	Invalidate
7	Special

During the address cycle of external special requests, **SysCmd[4:3]** contain the external special cause indication, as shown in Table 6-13. This information is used to differentiate between the various external special requests.

Table 6-13 Encoding of **SysCmd[4:3]** *for External Special Requests*

SysCmd[4:3]	Special Cause Indication
0	Reserved
1	NOP
2	Interrupt
3	Reserved

During external address cycles, **SysCmd[0]** specifies whether ECC checking and correcting is to be performed for the **SysAD[63:0]** bus, as shown in Table 6-14. During the address cycle of processor block read, block write, upgrade, and eliminate requests, the processor asserts **SysCmd[0]**. Consequently, in a multiprocessor system using the cluster bus, ECC checking and correcting is enabled for external coherency requests resulting from processor coherent block read and upgrade requests.

Table 6-14 Encoding of **SysCmd[0]** *for External Address Cycles*

SysCmd[0]	ECC check indication
0	ECC checking and correcting disable
1	ECC checking and correcting enable

SysCmd[10:0] Data Cycle Encoding

During the data cycles of an external data response or a processor coherency data response, **SysCmd[10:8]** contain the request number associated with the original request, as shown in Table 6-15.

Table 6-15 Encoding of **SysCmd[10:8]** *for Data Responses*

SysCmd[10:8]	Request Number

During data cycles, **SysCmd[5]** indicates the data quality, as shown in Table 6-16.

Table 6-16 Encoding of **SysCmd[5]** *for Data Cycles*

SysCmd[5]	Data quality indication
0	Data is good
1	Data is bad

During data cycles, **SysCmd[4:3]** indicate the data type, as shown in Table 6-17. Processor block write and double/single/partial-word write requests use request data and request last data type indications. External data and processor coherency data responses use response data and response last data type indications.

Table 6-17 Encoding of **SysCmd[4:3]** *for Data Cycles*

SysCmd[4:3]	Data type Indication
0	Request data
1	Response data
2	Request last
3	Response last

During data cycles of an external block data response or processor coherency data response, **SysCmd[2:1]** contain the state of the cache block, as shown in Table 6-18.

Table 6-18 Encoding of **SysCmd[2:1]** *for Block Data Responses*

SysCmd[2:1]	Cache Block State
0	Reserved
1	*Shared*
2	*CleanExclusive*
3	*DirtyExclusive*

During data cycles, **SysCmd[0]** specifies whether ECC checking and correcting is to be performed for the **SysAD[63:0]** bus, as shown in Table 6-19. During processor data cycles, the processor asserts **SysCmd[0]**. Consequently, in a multiprocessor system using the cluster bus, ECC checking and correcting will be enabled for external block data responses resulting from processor coherency data responses.

Table 6-19 Encoding of **SysCmd[0]** *for External Data Cycles*

SysCmd[0]	ECC check indication
0	ECC checking and correcting disable
1	ECC checking and correcting enable

SysCmd[11:0] Map

Table 6-20 presents a map for the **SysCmd[11:0]** bus.

Table 6-20 SysCmd[11:0] Map

Cycle Type	Command		SysCmd[11:0] Bit											
			11	10	9	8	7	6	5	4	3	2	1	0
Processor address cycles	Coherent block read shared		0	Request number			0	0	0	Read cause		Block state		1
	Coherent block read exclusive						0	0	1					
	Noncoherent block read						0	1	0					
	Double/single/partial-word read						0	1	1			Data size		
	Block write			0			1	0	0	Write cause		Block state		1
	Double/single/partial-word write						1	0	1			Data size		
	Upgrade			Request number			1	1	0	Upgrade cause		Block state		1
	Special	Reserved		Reserved			1	1	1	0	0	Reserved		
		Eliminate		0						0	1	Block state		1
		Reserved		Reserved						1	0	Reserved		
										1	1			
Processor data cycles	Double/single/partial-word write		1	0			0		0	Data type		0		1
	Block write								Data quality			Block state		
	Coherency data response			Request number										
External address cycles	Intervention shared		0	Request number			0	0	0	X				ECC
	Intervention exclusive						0	0	1					
	Allocate request number						0	1	0					X
							0	1	1					
	NOP			X			1	0	0					
							1	0	1					
	Invalidate			Request number			1	1	0					ECC
	Special	NOP		X			1	1	1	0	0	X		X
										0	1			
		Interrupt								1	0			ECC
		NOP								1	1			X
External data cycles	Block data response		1	Request number			X		Data quality	Data type		Block state		ECC
	Double/single/partial-word data response											X		

SysAD[63:0] Encoding

This section describes the system address/data bus encoding.

SysAD[63:0] Address Cycle Encoding

Table 6-21 presents the encoding of the **SysAD[63:0]** bus for address cycles.

Table 6-21 Encoding of **SysAD[63:0]** *for Address Cycles*

SysAD[63:60]	Target Indication
SysAD[63]	Target processor with **DevNum** = 3
SysAD[62]	Target processor with **DevNum** = 2
SysAD[61]	Target processor with **DevNum** = 1
SysAD[60]	Target processor with **DevNum** = 0
SysAD[59:58]	Uncached attribute
SysAD[57]	Secondary cache block way indication
SysAD[56:40]	Reserved
SysAD[39:0]	Physical address

SysAD[63:60]

During the address cycle of processor noncoherent block read, double/single/partial-word read, block write, double/single/partial-word write, and eliminate requests, the processor always drives a target indication of 0 on **SysAD[63:60]**. This indicates that the request targets the external agent only. When the **CohPrcReqTar** mode bit is negated, during the address cycle of processor coherent block read and upgrade requests, the processor also drives a target indication of 0 on **SysAD[63:60]**. However, when the **CohPrcReqTar** mode bit is asserted, during the address cycle of processor coherent block read and upgrade requests, the processor drives a target indication of 0xF on **SysAD[63:60]**. This indicates that the request targets all processors, together with the external agent, on the cluster bus. In multiprocessor systems using the cluster bus, the **CohPrcReqTar** mode bit is asserted for a snoopy-based coherency protocol, and negated for a duplicate tag or directory-based coherency protocol.

When the processor is in slave state, an external agent uses the target indication field to specify which processors are targets of an external request.

SysAD[59:58] Uncached Attribute

During the address cycle of processor double/single/partial-word read and write requests and during the address cycle of processor *Uncached accelerated* block write requests, the processor drives the uncached attribute onto **SysAD[59:58]**. See the section titled, Support for Uncached Attribute, in this chapter for more information.

SysAD[57]

During the address cycle of processor block read, typical block write, upgrade, and eliminate requests, **SysAD[57]** contains the secondary cache block way indication. This information may be useful for system designs implementing a duplicate tag or a directory-based coherency protocol.

SysAD[56:40]

When processor is in master state, it drives **SysAD[56:40]** to zero during address cycles.

SysAD[39:0]

During the address cycle of processor and external requests, **SysAD[39:0]** contain the physical address.

Table 6-22 presents the processor request address cycle address alignment.

Table 6-22 Processor Request Address Cycle Alignment

Processor Request Type	Address Alignment	Address Bits Which Are Driven to 0
Block read	Quadword	3:0
Doubleword read/write	Doubleword	2:0
Singleword read/write	Singleword	1:0
Halfword read/write	Halfword	0
Byte, tribyte, quintibyte, sextibyte, septibyte read/write	Byte	-
Block write	Secondary cache block	5:0 (**SCBlkSize** = 0) 6:0 (**SCBlkSize** = 1)
Upgrade	Quadword	3:0
Eliminate	Secondary cache block	5:0 (**SCBlkSize** = 0) 6:0 (**SCBlkSize** = 1)

Table 6-23 presents the external coherency request address cycle address alignment.

Table 6-23 External Coherency Request Address Cycle Alignment

External Request Type	Address Alignment	Address Bits Which Are Ignored
Intervention	Quadword	3:0
Invalidate	Secondary cache block	5:0 (**SCBlkSize** = 0) 6:0 (**SCBlkSize** = 1)

SysAD[63:0] Data Cycle Encoding

During System interface data cycles, when less than a doubleword is transferred on the **SysAD[63:0]** bus, the valid byte lanes depend on the request address and the **MemEnd** mode bit.

For example, consider the data cycle for a byte request whose address modulo 8 is 1. When **MemEnd** is negated (little endian), the **SysAD[15:8]** byte lane is valid. When **MemEnd** is asserted (big endian), the **SysAD[55:48]** byte lane is valid.

SysState[2:0] Encoding

The processor provides a processor coherency state response by driving the targeted secondary cache block tag quality indication on **SysState[2]**, driving the targeted secondary cache block former state on **SysState[1:0]** and asserting **SysStateVal*** for one **SysClk** cycle. Table 6-24 presents the encoding of the **SysState[2:0]** bus when **SysStateVal*** is asserted.

Table 6-24 Encoding of **SysState[2:0]** *when* **SysStateVal*** *Asserted*

SysState[2]	**Secondary cache block tag quality indication**
0	Tag is good
1	Tag is bad
SysState[1:0]	**Secondary cache block former state**
0	*Invalid*
1	*Shared*
2	*CleanExclusive*
3	*DirtyExclusive*

When **SysStateVal*** is negated, **SysState[0]** indicates if a processor coherency data response is ready for issue. Table 6-25 presents the encoding of the **SysState[2:0]** bus when **SysStateVal*** is negated.

Table 6-25 Encoding of SysState **[2:0]** *When SysStateVal* Negated*

SysState[2:1]	**Reserved**
SysState[0]	**Processor coherency data response indication**
0	Not ready for issue
1	Ready for issue

SysResp[4:0] Encoding

An external agent issues an external completion response by driving the request number associated with the corresponding request on **SysResp[4:2]**, driving the completion indication on **SysResp[1:0]**, and asserting **SysRespVal*** for one **SysClk** cycle. Table 6-26 presents the encoding of the **SysResp[4:0]** bus.

Table 6-26 Encoding of **SysResp[4:0]**

SysResp[4:2]	Request number
SysResp[1:0]	**Completion indication**
0	Acknowledge (ACK)
1	Error (ERR)
2	Negative acknowledge (NACK)
3	Reserved

6.14 Interrupts

The processor supports five hardware, two software, one timer, and one nonmaskable interrupt. The Interrupt exception is described in Chapter 17.

Hardware Interrupts

Five hardware interrupts are accessible to an external agent via external interrupt requests.

An external interrupt request consists of a single address cycle on the System interface. During the address cycle, **SysAD[63:60]** specify the target indication, which allows an external agent to define the target processors of the external interrupt request. If a processor determines it is an external interrupt request target, **SysAD[20:16]** are the write enables for the five individual *Interrupt* register bits and **SysAD[4:0]** are the values to be written into these bits, as shown in Figure 6-5. This allows any subset of the *Interrupt* register bits to be set or cleared with a single external interrupt request.

The *Interrupt* register is an architecturally transparent, level-sensitive register that is directly readable as bits 14:10 of the *Cause* register. Since it is level-sensitive, an interrupt bit must remain asserted until the interrupt is taken, at which time the interrupt handler must cause a second external interrupt request to clear the bit.

The processor clears the *Interrupt* register during any of the reset sequences.

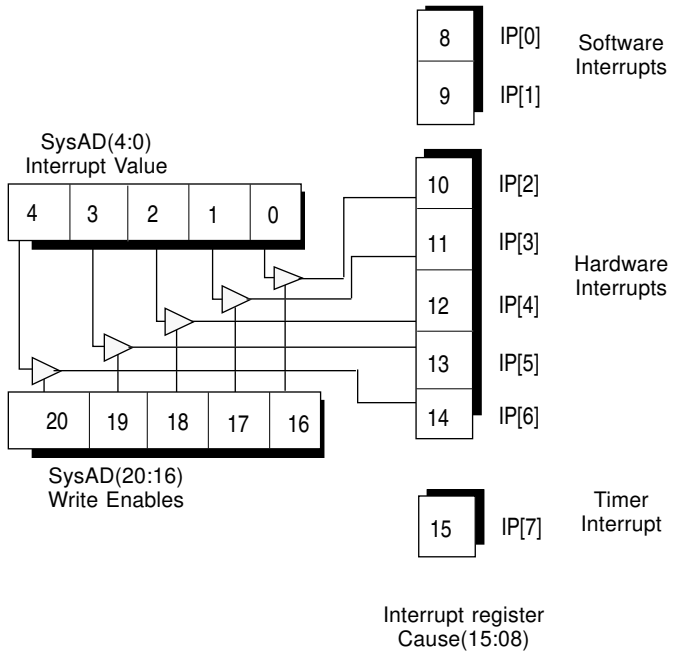

Figure 6-5 Hardware Interrupts

Software Interrupts

The two software interrupts are accessible as bits 9:8 of the *Cause* register, as shown in Figure 6-5. An MTC0 instruction is used to write these bits.

Timer Interrupt

The timer interrupt is accessible as bit 15 of the *Cause* register, **IP[7]**, as shown in Figure 6-5. This bit is set when one of the following occurs:

- the *Count* register is equal to the *Compare* register
- either one of the two performance counters overflows

Nonmaskable Interrupt

A nonmaskable interrupt is accessible to an external agent as the **SysNMI*** signal. To post a nonmaskable interrupt, an external agent asserts **SysNMI*** for at least one **SysClk** cycle.

The processor recognizes the nonmaskable interrupt on the first **SysClk** cycle that **SysNMI*** is asserted. After the nonmaskable interrupt is serviced, an external agent may post another nonmaskable interrupt by first negating **SysNMI*** for at least one **SysClk** cycle, and reasserting **SysNMI*** for at least one **SysClk** cycle.

6.15 Protocol Abbreviations

The following abbreviations are used in the System interface protocols:

SysCmd[11:0] Abbreviations

Cmd	Unspecified command
BlkRd	Block read request command
RdShd	Coherent block read shared request command
RdExc	Coherent block read exclusive request command
DSPRd	Double/single/partial-word read command
BlkWr	Block write request command
DSPWr	Double/single/partial-word write request command
Ugd	Upgrade request command
Elm	Eliminate request command
IvnShd	Intervention shared request command
IvnExc	Intervention exclusive request command
Alc	Allocate request number command
Ivd	Invalidate request command
Int	Interrupt request command
ExtCoh	External coherency request command
ReqDat	Request data
RspDat	Response data
ReqLst	Request last
RspLst	Response last
Empty	Empty; **SysCmd(11:0)** and **SysAD(63:0)** are undefined

SysAD[63:0] Abbreviations

Adr	Physical address
Dat	Unspecified data
Dat<n>	Doubleword n of a block

SysState[2:0] Abbreviations

State	Unspecified state
Ivd	*Invalid*
Shd	*Shared*
ClnExc	*CleanExclusive*
DrtExc	*DirtyExclusive*

SysResp[4:0] Abbreviations

Rsp	Unspecified completion response
ACK	Acknowledge completion response
ERR	Error completion response
NACK	Negative acknowledge completion response

Master Abbreviations

EA	External agent
P_n	R10000 processor whose device number is *n*
-	Dead cycle

6.16 System Interface Arbitration

The processor supports a simple System interface arbitration protocol, which relies on an external arbiter. This protocol is used in uniprocessor systems, multiprocessor systems using dedicated external agents, and multiprocessor systems using the cluster bus. System interface arbitration is handled by the **SysReq***, **SysGnt***, and **SysRel*** signals (request, grant, and release).

As described earlier in this chapter, the System interface resides in either master or slave state; the processor enters slave state during all of the reset sequences.

When mastership of the System interface changes, there is always one dead **SysClk** cycle during which the bidirectional signals are not driven; the processor ignores all bidirectional signals during this dead **SysClk** cycle.

The protocol supports overlapped arbitration which allows arbitration to occur in parallel with requests and responses. This results in fewer wasted cycles when mastership of the System interface changes.

Grant parking is also supported, allowing a device to retain mastership of the System interface as long as no other device requests the System interface.

In multiprocessor systems using the cluster bus, the external arbiter typically implements a round-robin priority scheme.

System Interface Arbitration Rules

The rules for the System interface arbitration are listed below:

- If the System interface is in slave state, and a processor request or coherency data response is ready for issue, and the required resources are available (e.g. a free request number, **SysRdRdy*** asserted, etc.), the processor asserts **SysReq***. The processor will not assert **SysReq*** unless all of the above conditions are met.
- The processor waits for the assertion of **SysGnt***.
- When the processor observes the assertion of **SysGnt*** it negates **SysReq*** two **SysClk** cycles later. Once the processor asserts **SysReq***, it does not negate **SysReq*** until the assertion of **SysGnt***, even if the need for the System interface bus is contravened by an external coherency request.
- When the processor observes the assertion of **SysRel***, it enters master state two **SysClk** cycles later, and begins to drive the System interface bus. **SysRel*** may be asserted coincidentally with or later than **SysGnt***.
- Once in master state, the processor does not relinquish mastership of the System interface until it observes the negation of **SysGnt***.
- The processor indicates it is relinquishing mastership of the System interface bus by asserting **SysRel*** for one **SysClk** cycle, two or more **SysClk** cycles after the negation of **SysGnt***. The processor may issue any type of processor request or coherency data response in the two **SysClk** cycles following the negation of **SysGnt***. This may delay the assertion of **SysRel***.

Uniprocessor System

Figure 6-6 shows how the System interface arbitration signals are used in a uniprocessor system. Note that this same configuration would be used in a multiprocessor system using dedicated external agents.

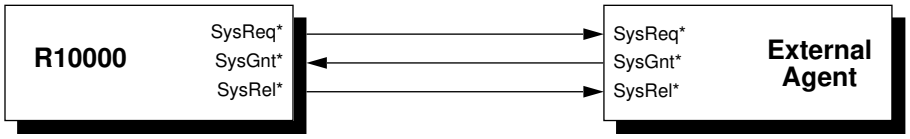

Figure 6-6 Arbitration Signals for Uniprocessor System

Figure 6-7 is an example of the operation of the System interface arbitration in a uniprocessor system. The *Master* row in the following figures indicates which device is driving the System interface bidirectional signals (P_0 and EA in Figure 6-7). When this row contains a dash (**-**), as shown in Cycle 12 of Figure 6-7, mastership of the System interface is changing and no device is driving the System interface bidirectional signals for this one dead **SysClk** cycle.

The external agent generally asserts the **SysGnt*** signal, which allows the processor to issue requests at any time.

When the external agent needs to return an external data response, it negates **SysGnt*** for a minimum of one cycle, waits for the processor to assert **SysRel***, and then begins driving the System interface bus after one dead **SysClk** cycle.

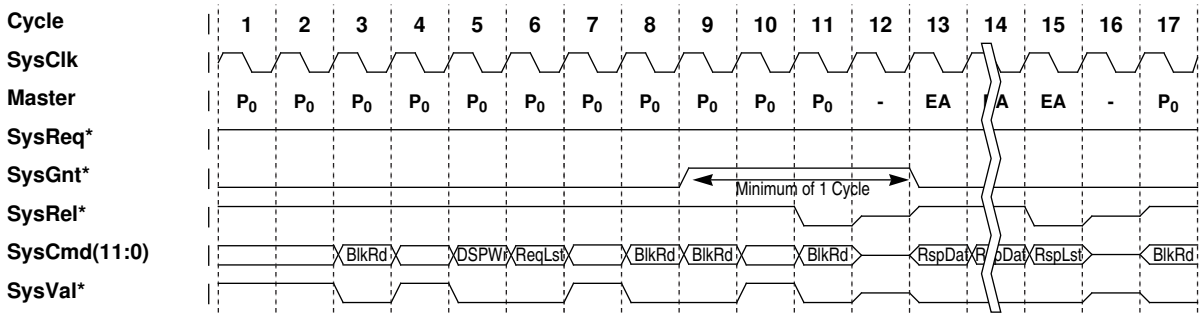

Figure 6-7 Arbitration Protocol for Uniprocessor System

Multiprocessor System Using Cluster Bus

Figure 6-8 shows how the System interface arbitration signals are used in a four-processor system using the cluster bus.

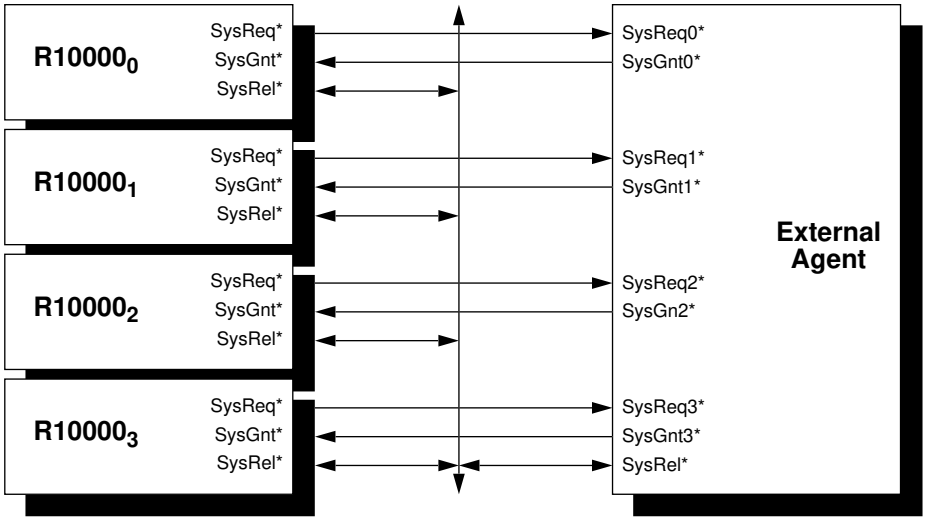

Figure 6-8 Arbitration Signals for Multiprocessor System Using the Cluster Bus

Figure 6-9 is an example of the System interface arbitration in a four-processor system using the cluster bus.

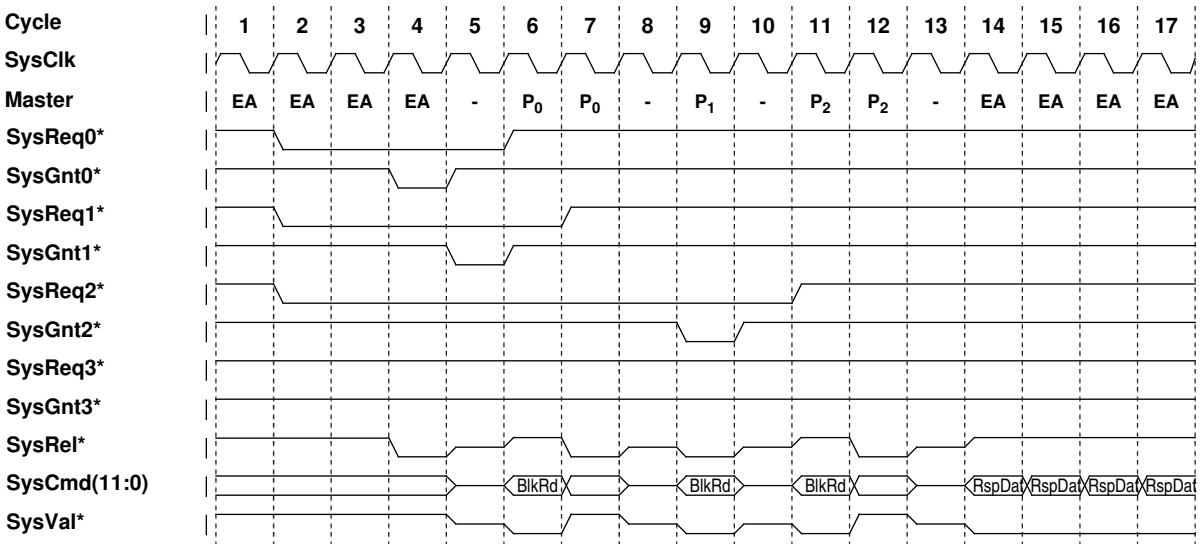

Figure 6-9 Arbitration Protocol for Multiprocessor System Using the Cluster Bus

6.17 System Interface Request and Response Protocol

The following sections detail the System interface request and response protocol. A 32-word secondary cache block size is assumed in the examples below.

Processor Request Protocol

A processor request is generated when the R10000 processor requires a system resource.

The processor may only issue a processor request when the System interface is in master state. If the System interface is in master state, the processor may issue a processor request immediately. Processor requests may occur in adjacent **SysClk** cycles. If the System interface is not in master state, the processor must first assert **SysReq***, and then wait for the external agent to relinquish mastership of the System interface bus by asserting **SysGnt*** and **SysRel***.

When multiple processor requests and/or coherency data responses are ready and meet all issue requirements, the processor uses the following priority:

- block read and upgrade requests have the highest priority, followed by
- processor coherency data responses,
- processor eliminate and typical block write requests,
- processor double/single/partial-word read/write and uncached accelerated block write requests, which have the lowest priority.

Processor Block Read Request Protocol

A processor block read request results from a cached instruction fetch, load, store, or prefetch that misses in the secondary cache.

The processor issues a processor block read request with a single address cycle. The address cycle consists of the following:

- negating **SysCmd[11]**
- driving a free request number on **SysCmd[10:8]**
- driving the block read command on **SysCmd[7:5]**
- driving the read cause indication on **SysCmd[4:3]**
- driving the secondary cache block former state on **SysCmd[2:1]**
- asserting **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the secondary cache block way on **SysAD[57]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

The processor may only issue a processor block read request address cycle when the following are true:

- the System interface is in master state
- **SysRdRdy*** was asserted two **SysClk** cycles earlier
- there is no conflicting entry in the outgoing buffer
- the maximum number of outstanding processor requests specified by the **PrcReqMax** mode bits is not exceeded
- there is a free request number
- the processor is not the target of a conflicting outstanding external coherency request

A single processor may have as many as four processor block read requests outstanding on the System interface at any given time.

Figure 6-10 depicts four processor block read requests. Since the System interface is initially in slave state, the processor must first assert **SysReq*** and then wait until the external agent relinquishes mastership of the System interface by asserting **SysGnt*** and **SysRel***.

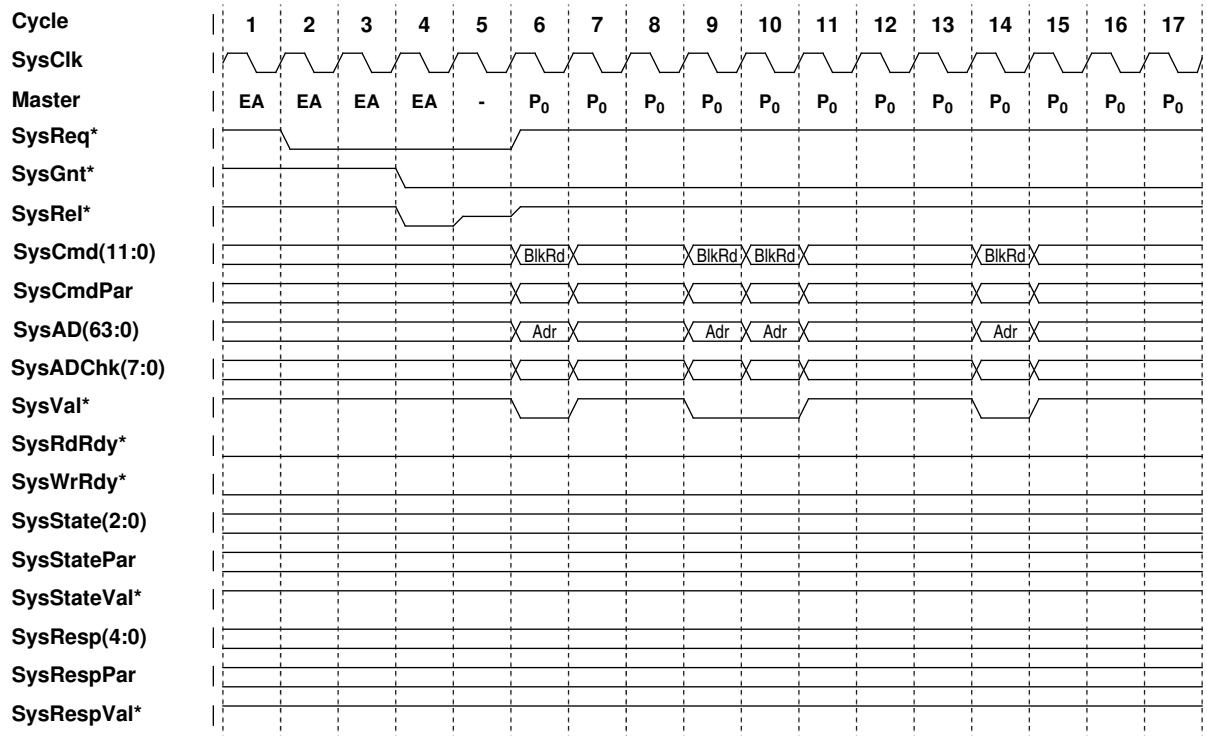

Figure 6-10 Processor Block Read Request Protocol

Processor Double/Single/Partial-Word Read Request Protocol

A processor double/single/partial-word read request results from an uncached instruction fetch or load.

The processor issues a processor double/single/partial-word read request with a single address cycle. The address cycle consists of:

- negating **SysCmd[11]**
- driving a free request number on **SysCmd[10:8]**
- driving the double/single/partial-word read command on **SysCmd[7:5]**
- driving the read cause indication on **SysCmd[4:3]**
- driving the data size indication on **SysCmd[2:0]**
- driving the target indication on **SysAD[63:60]**
- driving the uncached attribute on **SysAD[59:58]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

The processor may only issue a processor double/single/partial-word read request address cycle when:

- the System interface is in master state
- **SysRdRdy*** was asserted two **SysClk** cycles previously
- the maximum number of outstanding processor requests specified by the **PrcReqMax** mode bits is not exceeded
- there is a free request number

A single processor may have a maximum of one processor double/single/partial-word read request outstanding on the System interface at any given time.

Figure 6-11 depicts a processor double/single/partial-word read request. Since the System interface is initially in slave state, the processor must first assert **SysReq*** and then wait until the external agent gives up mastership of the System interface by asserting **SysGnt*** and **SysRel***.

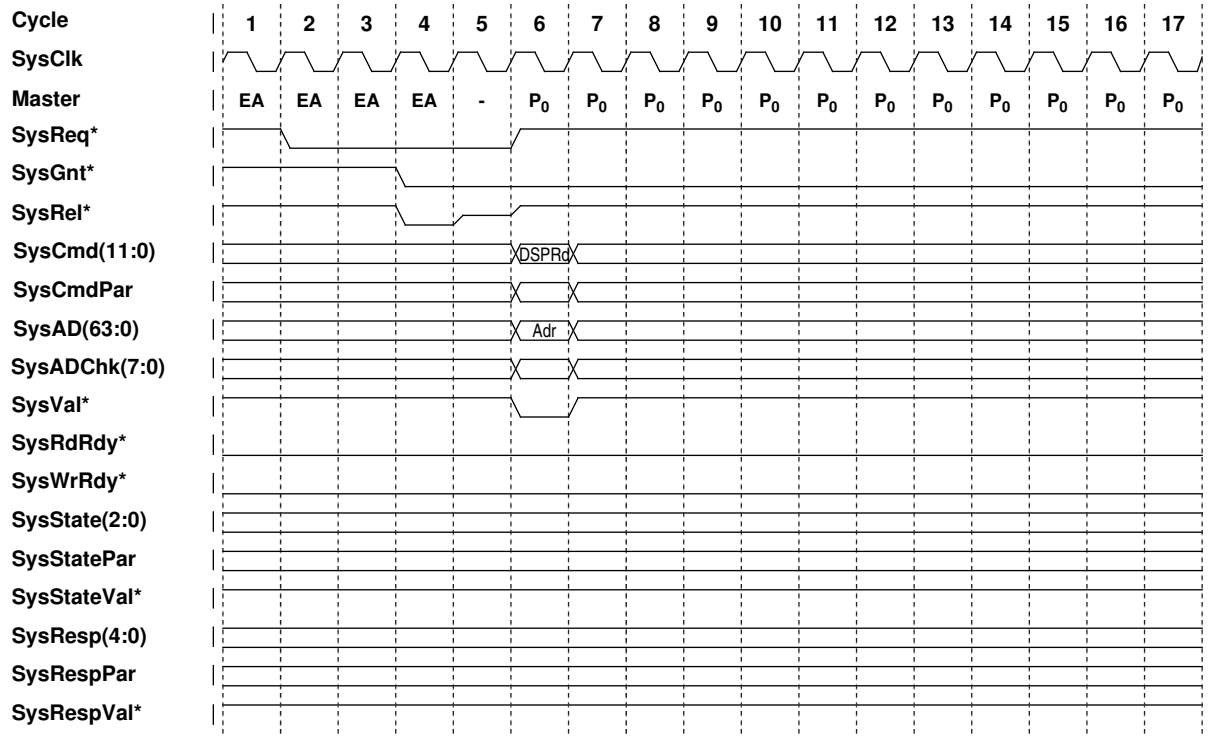

Figure 6-11 Processor Double/Single/Partial-Word Read Request Protocol

Processor Block Write Request Protocol

A processor block write request results from the following:

- replacement of a *DirtyExclusive* secondary cache block due to a load, store, or prefetch secondary cache miss
- a CACHE Index WriteBack Invalidate (S) or Hit WriteBack Invalidate (S) instruction
- a completely gathered uncached accelerated block

As shown in Figure 6-12, the processor issues a processor block write request with a single address cycle followed by 8 or 16 data cycles.

The address cycle consists of the following:

- negating **SysCmd[11]**
- driving the block write command on **SysCmd[7:5]**
- driving the write cause indication on **SysCmd[4:3]**
- asserting **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

If the processor block write request results from the writeback of a secondary cache block, the *DirtyExclusive* secondary cache block former state is driven on **SysCmd[2:1]**, and the secondary cache block way is driven on **SysAD[57]**.

If the processor block write request results from a completely gathered uncached accelerated block, the uncached attribute is driven on **SysAD[59:58]**.

Each data cycle consists of the following:

- asserting **SysCmd[11]**
- driving the data quality indication on **SysCmd[5]**
- driving the data type indication on **SysCmd[4:3]**
- driving the data on **SysAD[63:0]**
- asserting **SysVal***

The first 7 or 15 data cycles have a request data type indication, and the last data cycle has a request last data type indication.

The processor may negate **SysVal*** between data cycles of a processor block write request only if the **SCClk** frequency is less than half of the **SysClk** frequency.

The processor may only issue a processor block write request address cycle when the following are true:

- the System interface is in master state
- **SysWrRdy*** was asserted two **SysClk** cycles previously
- the processor is not the target of a conflicting outstanding external coherency request

Figure 6-12 depicts two adjacent processor block write requests. Since the System interface is initially in slave state, the processor must first assert **SysReq*** and then wait until the external agent relinquishes mastership of the System interface by asserting **SysGnt*** and **SysRel***.

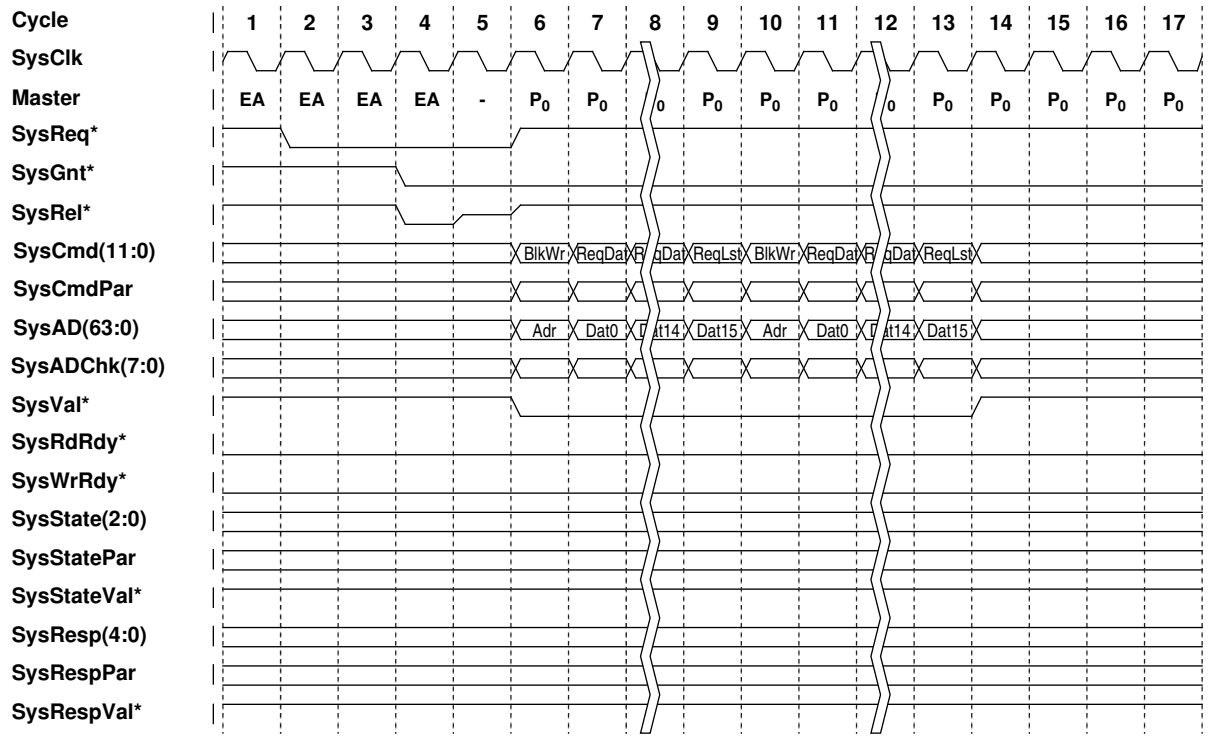

Figure 6-12 Processor Block Write Request Protocol

Processor Double/Single/Partial-Word Write Request Protocol

A processor double/single/partial-word write request results from an uncached store or incompletely gathered uncached accelerated block.

As shown in Figure 6-13, the processor issues a processor double/single/partial-word write request with a single address cycle immediately followed by a single data cycle.

The address cycle consists of the following:

- negating **SysCmd[11]**
- driving the double/single/partial-word write command on **SysCmd[7:5]**
- driving the write cause indication on **SysCmd[4:3]**
- driving the data size indication on **SysCmd[2:0]**
- driving the target indication on **SysAD[63:60]**
- driving the uncached attribute on **SysAD[59:58]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

The data cycle consists of the following:

- asserting **SysCmd[11]**
- driving the request last data type indication on **SysCmd[4:3]**
- driving the write data on **SysAD[63:0]**
- asserting **SysVal***

The processor may only issue a processor double/single/partial-word write request address cycle when the System interface is in master state and **SysWrRdy*** was asserted two **SysClk** cycles previously.

Figure 6-13 depicts three processor double/single/partial write requests. Since the System interface is initially in slave state, the processor must first assert **SysReq*** and then wait until the external agent relinquishes mastership of the System interface by asserting **SysGnt*** and **SysRel***.

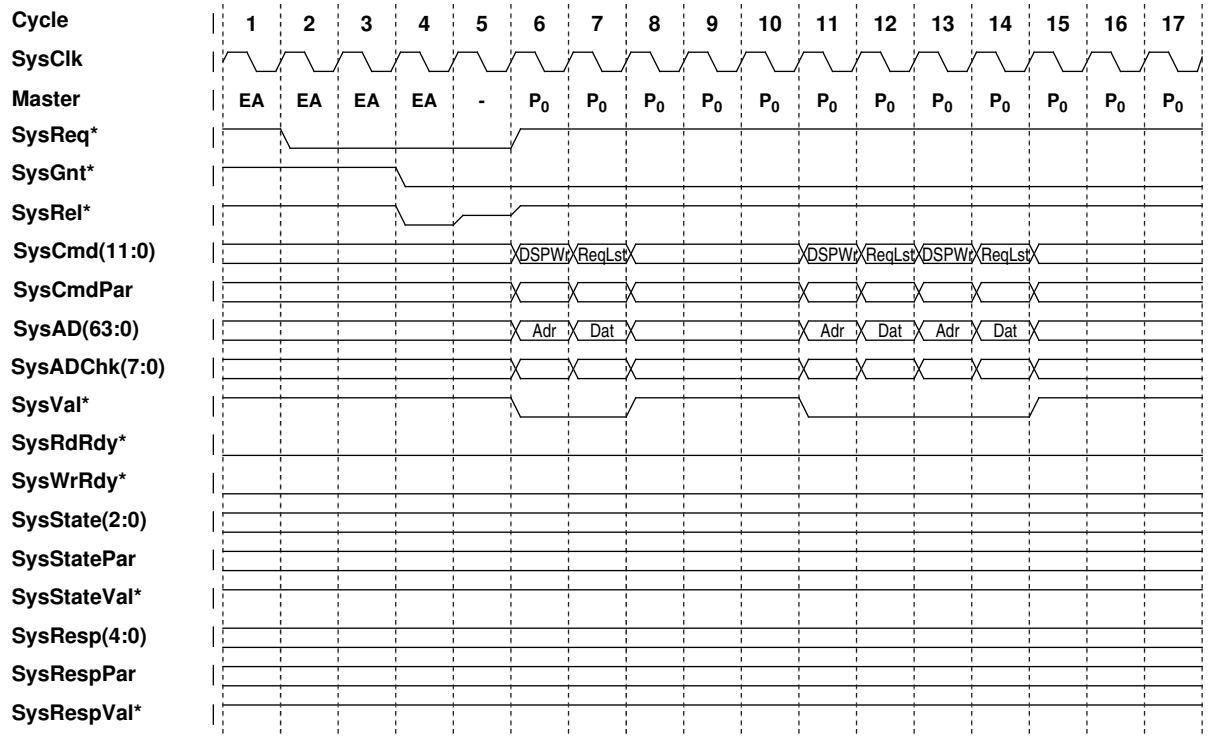

Figure 6-13 Processor Double/Single/Partial-Word Write Request Protocol

Processor Upgrade Request Protocol

A processor upgrade request results from a store or prefetch exclusive that hits a *Shared* block in the secondary cache.

As shown in Figure 6-14, the processor issues a processor upgrade request with a single address cycle. This address cycle consists of the following:

- negating **SysCmd[11]**
- driving a free request number on **SysCmd[10:8]**
- driving the upgrade command on **SysCmd[7:5]**
- driving the upgrade cause indication on **SysCmd[4:3]**
- driving the secondary cache block former state on **SysCmd[2:1]**
- asserting **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the secondary cache block way on **SysAD[57]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

The processor may only issue a processor upgrade request address cycle when the following are true:

- the System interface is in master state
- **SysRdRdy*** was asserted two **SysClk** cycles previously
- the maximum number of outstanding processor requests specified by the **PrcReqMax** mode bits is not exceeded
- there is a free request number
- the processor is not the target of a conflicting outstanding external coherency request

A single processor may have as many as four processor upgrade requests outstanding on the System interface at any given time.

Figure 6-14 depicts four processor upgrade requests. Since the System interface is initially in slave state, the processor must first assert **SysReq*** and then wait until the external agent relinquishes mastership of the System interface by asserting **SysGnt*** and **SysRel***.

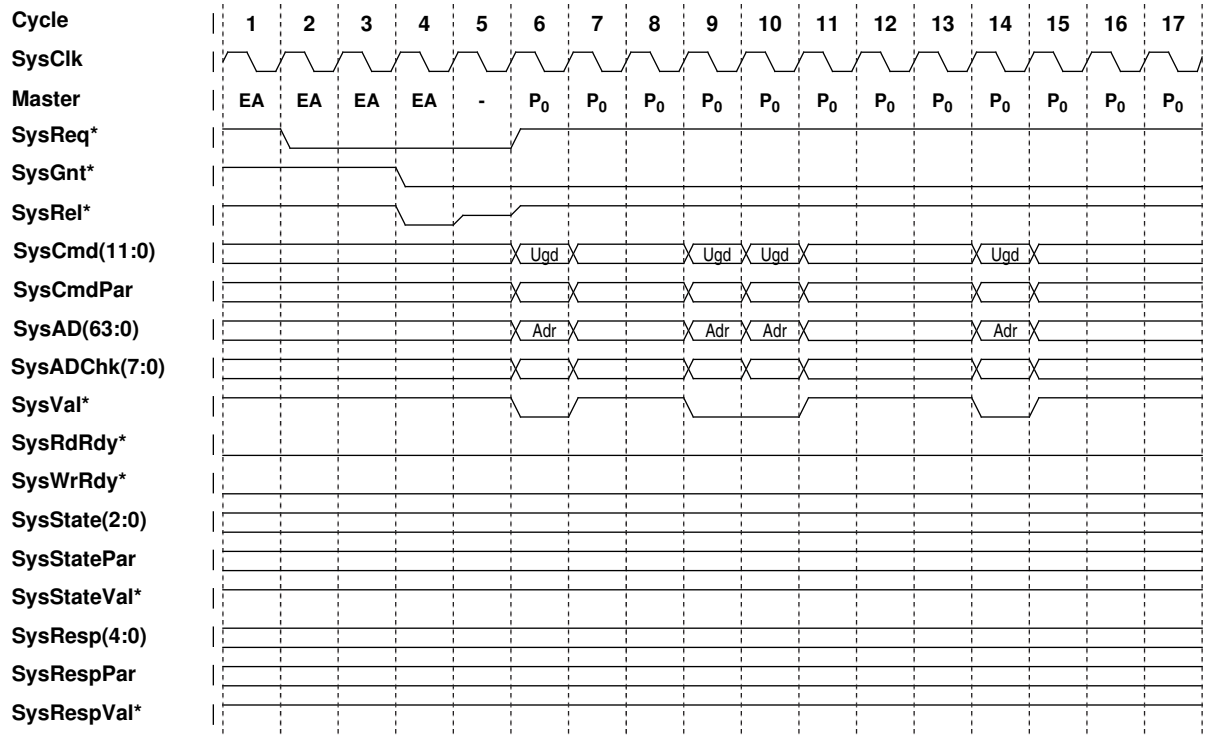

Figure 6-14 Processor Upgrade Request Protocol

Processor Eliminate Request Protocol

A processor eliminate request results from the following:

- a cached instruction fetch, load, store, or prefetch that misses in the secondary cache and forces the replacement of a *Shared* or *CleanExclusive* secondary cache block
- a CACHE Index WriteBack Invalidate (S), Hit Invalidate (S), or Hit WriteBack Invalidate (S) instruction that forces the invalidation of a *Shared* or *CleanExclusive* secondary cache block
- a CACHE Hit Invalidate (S) instruction that forces the invalidation of a *DirtyExclusive* secondary cache block.

A processor eliminate request notifies the external agent that a *Shared*, *CleanExclusive*, or *DirtyExclusive* block has been eliminated from the secondary cache. Such requests are useful for systems implementing a directory-based coherency protocol, and are enabled by asserting the **PrcElmReq** mode bit.

The processor issues a processor eliminate request with a single address cycle. This address cycle consists of the following:

- negating **SysCmd[11]**
- driving the special command on **SysCmd[7:5]**
- driving the eliminate special cause indication on **SysCmd[4:3]**
- driving the secondary cache block former state on **SysCmd[2:1]**
- asserting **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the secondary cache block way on **SysAD[57]**
- driving the physical address of the eliminated secondary cache block on **SysAD[39:0]**
- asserting **SysVal***

The processor may only issue a processor eliminate request address cycle when the following are true:

- the System interface is in master state
- **SysWrRdy*** was asserted two **SysClk** cycles previously
- the **PrcElmReq** mode bit is asserted
- the processor is not the target of a conflicting outstanding external coherency request

Figure 6-15 depicts three processor eliminate requests. Since the System interface is initially in slave state, the processor must first assert **SysReq*** and then wait until the external agent relinquishes mastership of the System interface by asserting **SysGnt*** and **SysRel***.

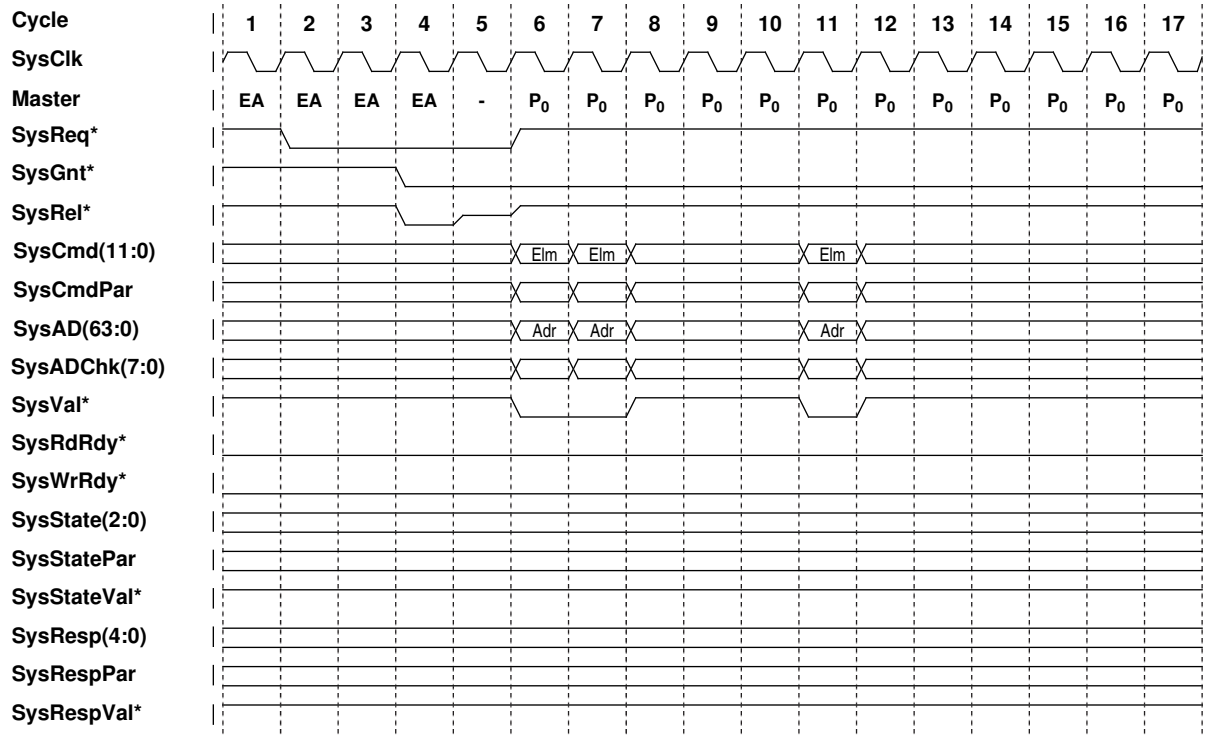

Figure 6-15 Processor Eliminate Request Protocol

Processor Request Flow Control Protocol

The processor provides the signals **SysRdRdy*** and **SysWrRdy*** to allow an external agent to control the flow of processor requests. **SysRdRdy*** controls the flow of processor read and upgrade requests whereas **SysWrRdy*** controls the flow of processor write and eliminate requests.

The processor can only issue a processor read or upgrade request address cycle to the System interface if **SysRdRdy*** was asserted two **SysClk** cycles previously. Similarly, the processor can only issue the address cycle of a processor write or eliminate request to the System interface if **SysWrRdy*** was asserted two **SysClk** cycles previously.

To determine the processor request buffering requirements for the external agent, note that the processor can issue any combination of processor requests in adjacent **SySClk** cycles. Also, since the System interface operates register-to-register with the external agent, a round trip delay of four **SysClk** cycles occurs between a processor request address cycle which prompts the external agent for flow control, and the flow control actually preventing any additional processor request address cycles from occurring. Consequently, if the maximum number of outstanding processor requests specified by the **PrcReqMax** mode bits is four, the external agent must be able to accept at least four processor read or upgrade requests. Also, the external agent must be able to accept at least four processor eliminate requests, two processor double/single/partial-word write requests, or one processor block write request.

Figure 6-16 depicts three processor double/single/partial-word write requests and four processor block read requests. After sensing the first processor double/single/partial-word write request, the external agent negates **SysWrRdy***. The external agent must have buffering sufficient for one additional processor write request before the flow control takes effect.

The external agent negates **SysRdRdy*** upon observing the first processor read request. The external agent must have buffering sufficient for three additional processor read requests before the flow control takes effect.

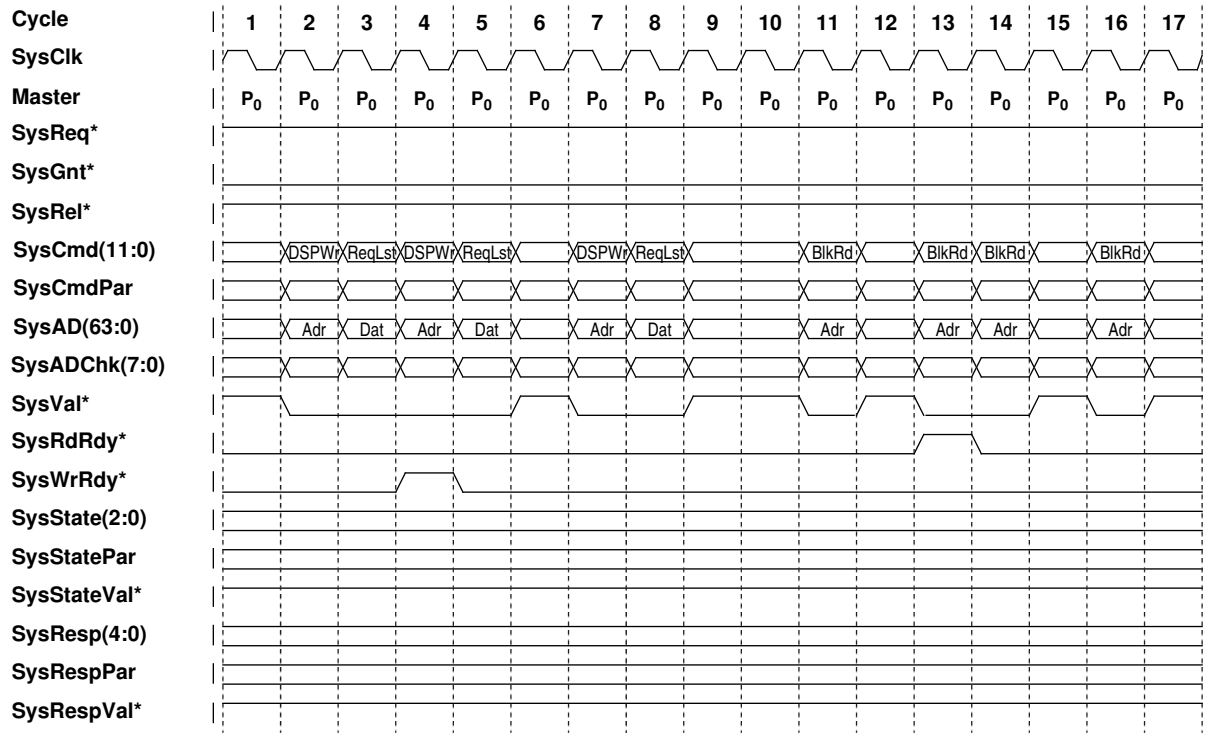

Figure 6-16 Processor Request Flow Control Protocol

External Response Protocol

The processor supports two classes of external responses:

- external data responses provide a double/single/partial-word of data or provide a block of data using the **SysAD[63:0]** bus
- external completion responses provide an acknowledge, error, or negative acknowledge indication using the **SysResp[4:0]** bus

An external agent may only issue an external data response to the processor when the System interface is in slave state. If the System interface is not already in slave state, the external agent must first negate **SysGnt*** and then wait for the processor to assert **SysRel***. If the System interface is already in slave state, the external agent may issue an external data response immediately.

External data responses may be accepted by the processor in adjacent **SysClk** cycles and in arbitrary order, relative to corresponding processor requests.

An external agent may issue an external completion response when the System interface is in either master or slave state. External completion responses may be accepted by the processor in adjacent **SysClk** cycles and in arbitrary order, relative to the corresponding processor requests.

External Block Data Response Protocol

An external agent may issue an external block data response in response to a processor block read or upgrade request.

An external agent issues an external block data response with 8 or 16 data cycles. Each data cycle consists of the following:

- asserting **SysCmd[11]**
- driving the request number associated with the corresponding processor request on **SysCmd[10:8]**
- driving the data quality indication on **SysCmd[5]**
- driving the data type indication on **SysCmd[4:3]**
- driving the cache block state on **SysCmd[2:1]**
- driving the ECC check indication on **SysCmd[0]**
- driving the data on **SysAD[63:0]**
- asserting **SysVal***

The first 7 or 15 data cycles have a response data type indication, and the last data cycle has a response last data type indication. The external agent may negate **SysVal*** between data cycles of an external block data response.

External block data response data must be supplied in subblock order, beginning with the quadword-aligned address specified by the corresponding processor request.

External block data responses for processor coherent block read shared or noncoherent block read requests may indicate a state of *Shared*, *CleanExclusive*, or *DirtyExclusive*. External block data responses for processor coherent block read exclusive or upgrade requests may indicate a state of *CleanExclusive* or *DirtyExclusive*.

Figure 6-17 depicts two processor block read requests and the corresponding external block data responses.

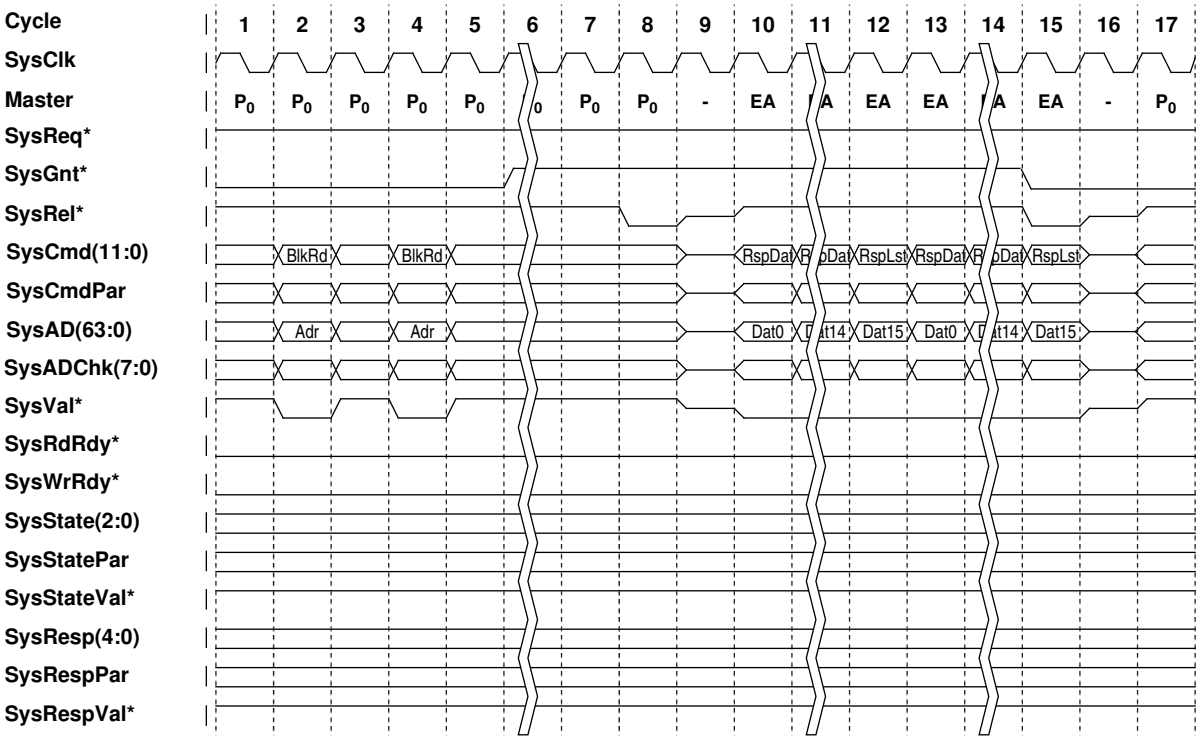

Figure 6-17 External Block Data Response Protocol

External Double/Single/Partial-Word Data Response Protocol

An external agent may issue an external double/single/partial-word data response in response to a processor double/single/partial-word read request.

An external agent issues an external double/single/partial-word data response with a single data cycle; the data cycle consists of:

- asserting **SysCmd[11]**
- driving the request number associated with the corresponding processor request on **SysCmd[10:8]**
- driving the data quality indication on **SysCmd[5]**
- driving the response last data type indication on **SysCmd[4:3]**
- driving the ECC check indication on **SysCmd[0]**
- driving the data on **SysAD[63:0]**
- asserting **SysVal***

Figure 6-18 depicts a processor double/single/partial-word read request and the corresponding external double/single/partial-word data response.

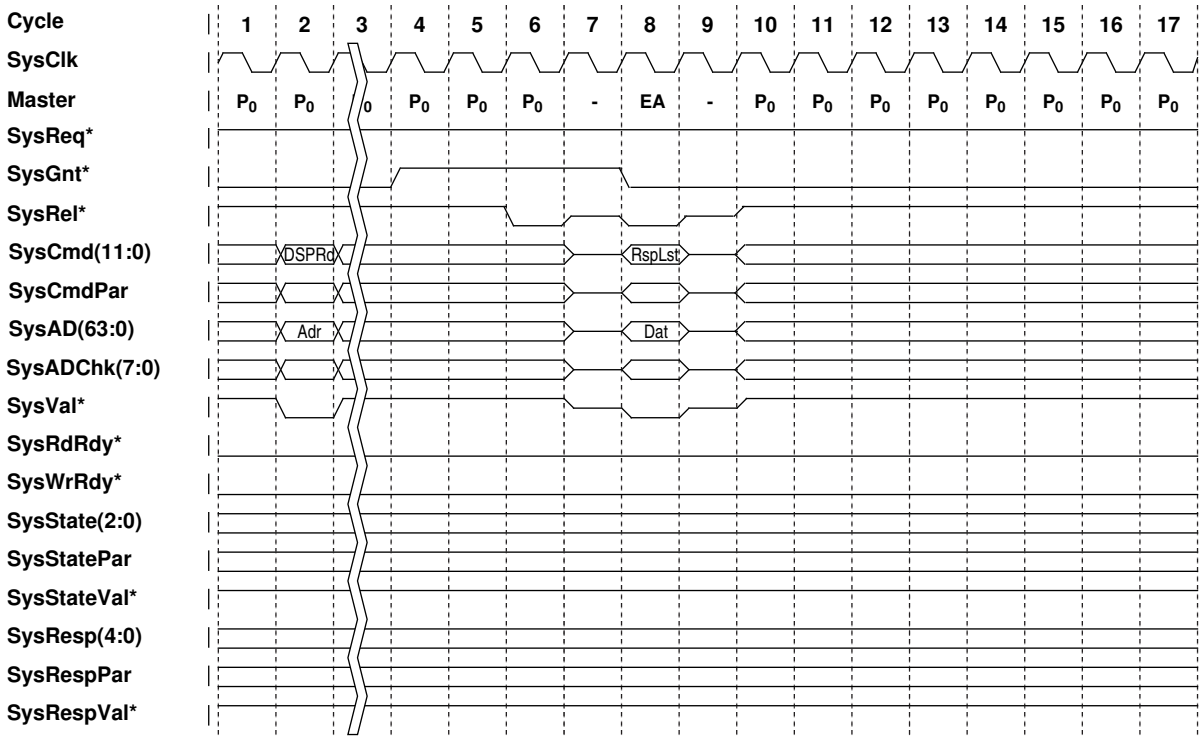

Figure 6-18 External Double/Single/Partial-Word Data Response Protocol

External Completion Response Protocol

An external agent issues an external completion response to provide an acknowledge, error, or negative acknowledge to an outstanding request, and to free the associated request number.

An external agent issues an external completion response by driving the response on **SysResp[4:0]** and asserting **SysRespVal*** for one **SysClk** cycle. **SysResp[4:2]** contains the request number associated with the corresponding outstanding request and **SysResp[1:0]** contains an acknowledge, error, or negative acknowledge indication, as described below:

- The external agent issues an external ACK completion response for a processor read or upgrade request to indicate that the request was successful. An external ACK completion response may only be issued for a processor read request if a corresponding external data response is coincidentally or previously issued.
- The external agent issues an external ERR completion response for a processor read or upgrade request to indicate that the request was unsuccessful. Upon receiving an external ERR completion response, the processor takes a Bus Error exception on the associated instruction. If the processor read or upgrade request was caused by a PREFETCH instruction, no exception is taken. Also, if the request was caused by a speculative instruction, no exception is taken.
- The external agent issues an external NACK completion response for a processor read or upgrade request to indicate that the request was not accepted. Upon receiving an external NACK completion response, the processor re-evaluates the associated instruction. Due to the speculative nature of the R10000 processor, the re-evaluation may or may not result in the reissue of a similar processor request.

An external ERR or NACK completion response issued in response to an external intervention, allocate request number, or invalidate has no affect on the processor except to free the request number.

Figure 6-19 depicts a processor upgrade request and a corresponding external completion response.

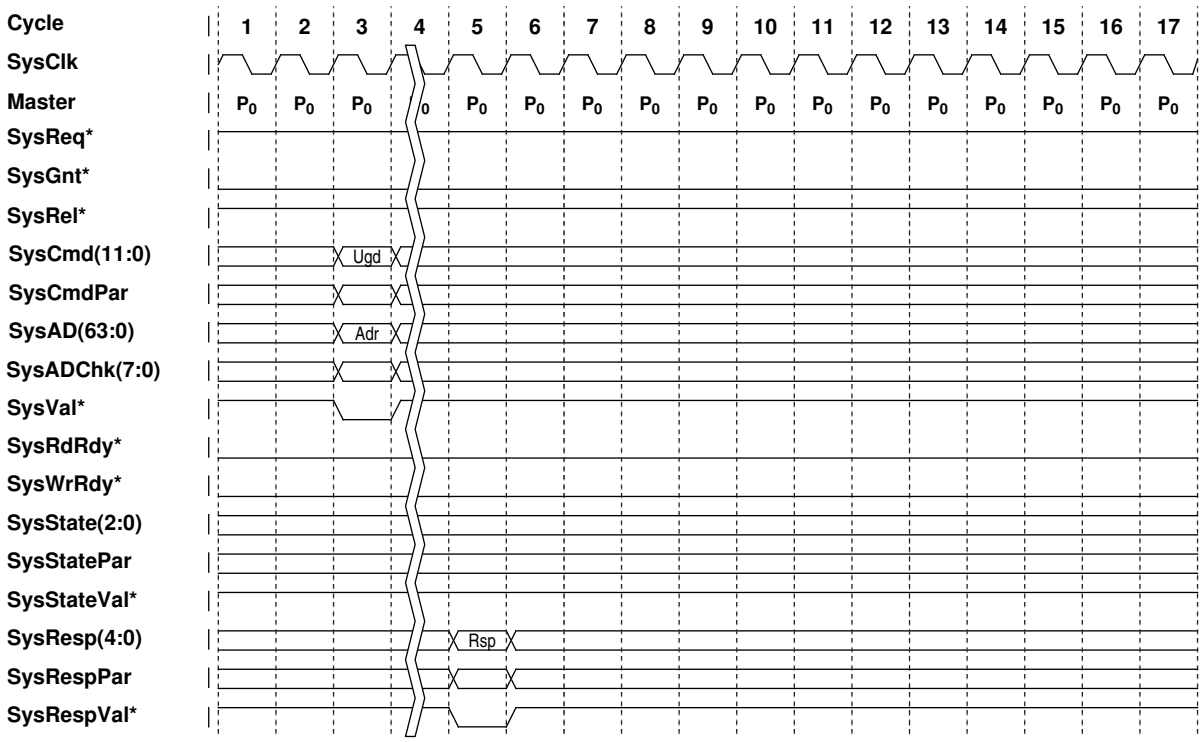

Figure 6-19 External Completion Response Protocol

External Request Protocol

An external agent issues an external request when it requires a resource within the processor.

An external agent may only issue an external request to the processor when the System interface is in slave state. If the System interface is not already in slave state, the external agent must first negate **SysGnt*** and then wait for the processor to assert **SysRel***. If the System interface is already in slave state, the external agent may issue an external request immediately. The total number of outstanding external requests, including interventions, allocate request numbers, and invalidates, cannot exceed eight.

External requests may be accepted by the processor in adjacent **SysClk** cycles. External intervention and invalidate requests are considered external coherency requests.

External Intervention Request Protocol

An external agent issues an external intervention request to obtain a *Shared* or *Exclusive* copy of a secondary cache block.

An external agent issues an external intervention request with a single address cycle; this address cycle consists of the following:

- negating **SysCmd[11]**
- driving a request number on **SysCmd[10:8]**
- driving the intervention command on **SysCmd[7:5]**
- driving the ECC check indication on **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

An external agent may only issue an external intervention request address cycle when the System interface is in slave state; typically a free request number is specified. An external agent may have as many as eight external intervention requests outstanding on the System interface at any given time.

Figure 6-20 depicts three external intervention requests. Since the System interface is initially in master state, the external agent must first negate **SysGnt*** and then wait until the processor relinquishes mastership of the System interface by asserting **SysRel***.

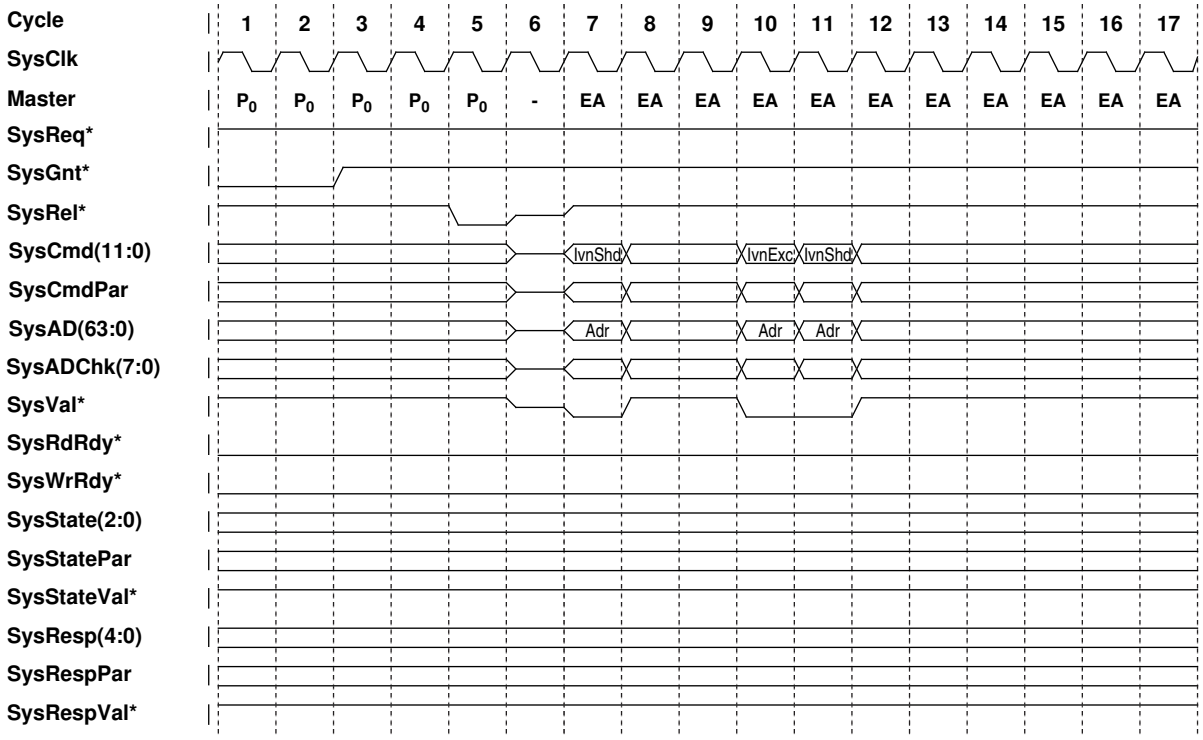

Figure 6-20 External Intervention Request Protocol

External Allocate Request Number Request Protocol

An external agent issues an external allocate request number request to reserve a request number for private use. Once allocated, the processor is prevented from using the request number until an external completion response for the request number is received.

An external agent issues an external allocate request number request with a single address cycle; this address cycle consists of the following:

- negating **SysCmd[11]**
- driving a free request number on **SysCmd[10:8]**
- driving the allocate request number command on **SysCmd[7:5]**
- asserting **SysVal***

An external agent may only issue an external allocate request number request address cycle when the System interface is in slave state and there is a free request number. The external agent may have as many as eight external allocate request number requests outstanding on the System interface at any given time.

Figure 6-21 depicts three external allocate request number requests. Since the System interface is initially in master state, the external agent must first negate **SysGnt*** and then wait until the processor relinquishes mastership of the System interface by asserting **SysRel***.

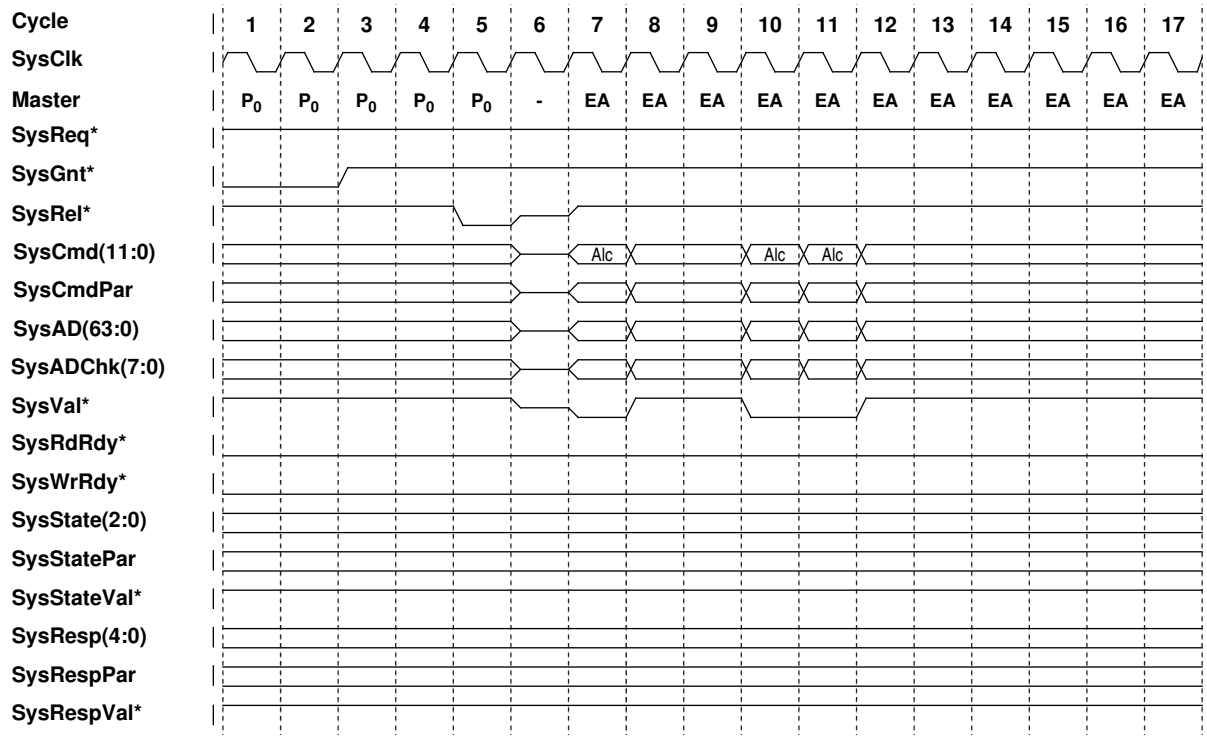

Figure 6-21 External Allocate Request Number Request Protocol

External Invalidate Request Protocol

An external agent issues an external invalidate request to invalidate a secondary cache block.

An external agent issues an external invalidate request with a single address cycle. This address cycle consists of the following:

- negating **SysCmd[11]**
- driving a request number on **SysCmd[10:8]**
- driving the invalidate command on **SysCmd[7:5]**
- driving the ECC check indication on **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the physical address on **SysAD[39:0]**
- asserting **SysVal***

An external agent may only issue an external invalidate request address cycle when the System interface is in slave state; typically a free request number is specified. An external agent may have as many as eight external invalidate requests outstanding on the System interface at any given time.

Figure 6-22 depicts three external invalidate requests. Since the System interface is initially in master state, the external agent must first negate **SysGnt*** and then wait until the processor relinquishes mastership of the System interface by asserting **SysRel***.

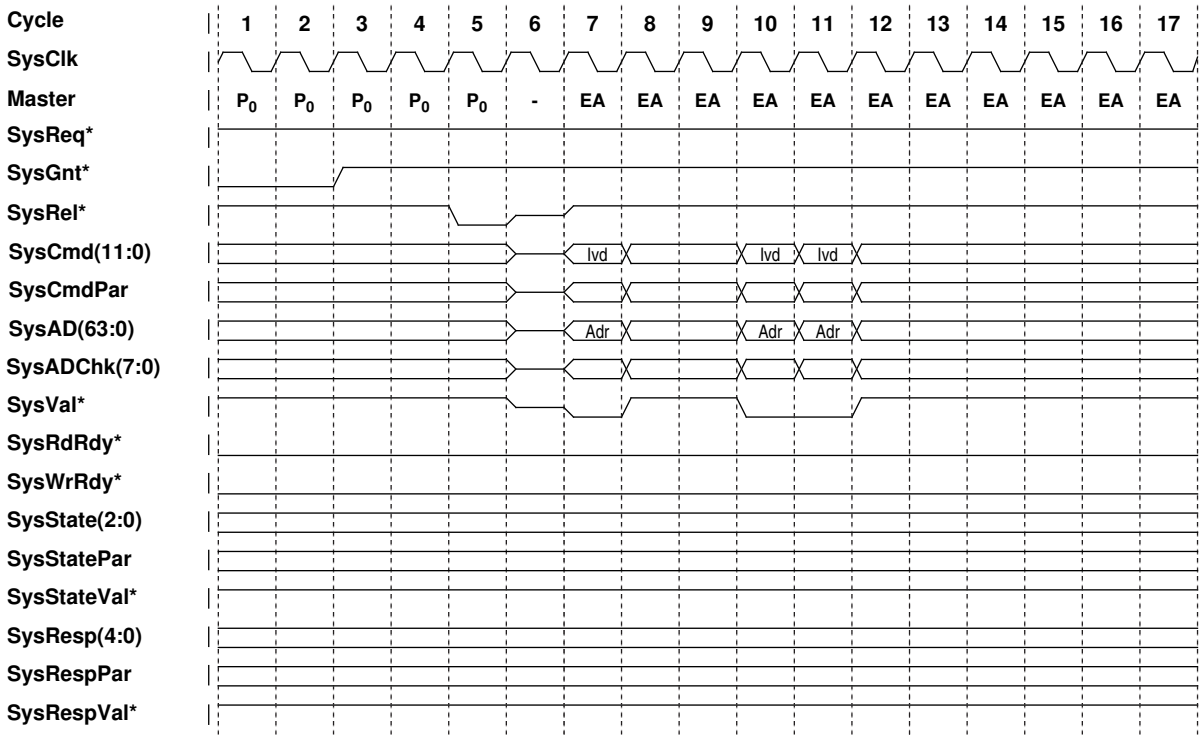

Figure 6-22 External Invalidate Request Protocol

External Interrupt Request Protocol

An external agent issues an external interrupt request to interrupt the normal instruction flow of the processor.

An external agent issues an external interrupt request with a single address cycle. This address cycle consists of the following:

- negating **SysCmd[11]**
- driving the special command on **SysCmd[7:5]**
- driving the interrupt special cause indication on **SysCmd[4:3]**
- driving the ECC check indication on **SysCmd[0]**
- driving the target indication on **SysAD[63:60]**
- driving the *Interrupt* register write enables on **SysAD[20:16]**
- driving the *Interrupt* register values on **SysAD[4:0]**
- asserting **SysVal***

An external agent may only issue an external interrupt request address cycle when the System interface is in slave state.

Figure 6-23 depicts three external interrupt requests. Since the System interface is initially in master state, the external agent must first negate **SysGnt*** and then wait until the processor relinquishes mastership of the System interface by asserting **SysRel***.

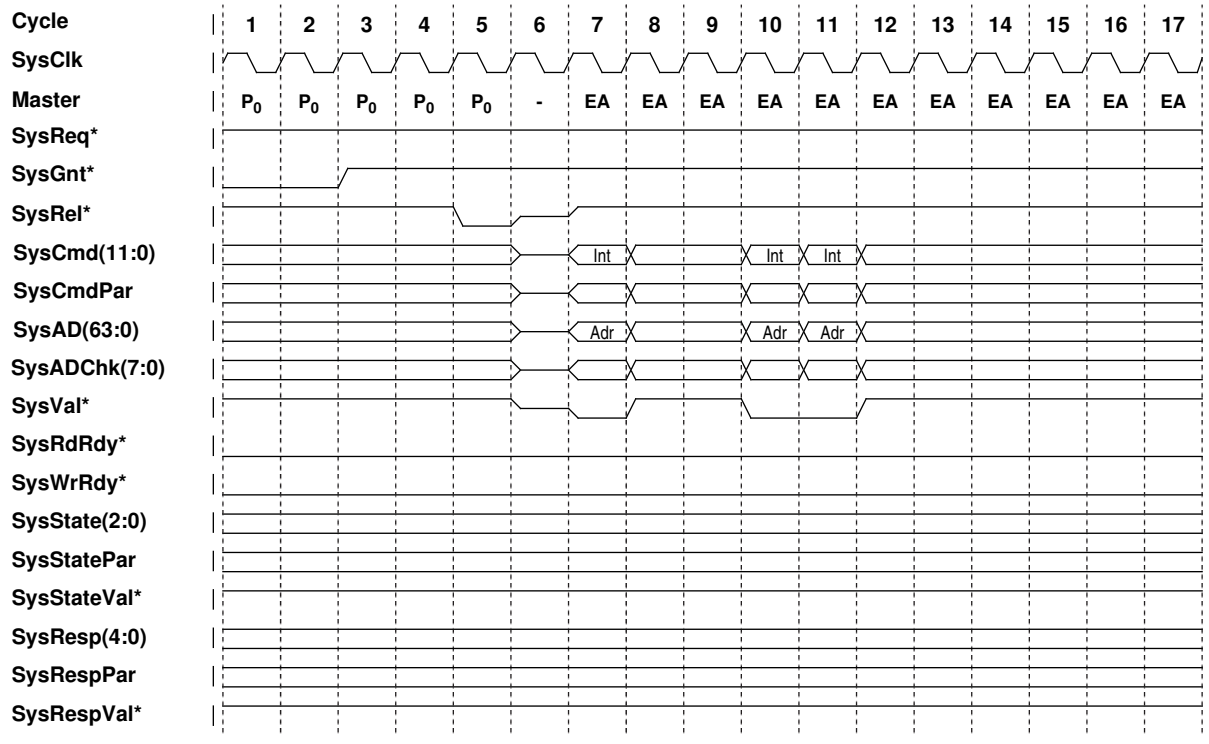

Figure 6-23 External Interrupt Request Protocol

Processor Response Protocol

Processor responses are supplied by the processor in response to external coherency requests that target the processor. The R10000 processor issues a processor coherency state response for each external coherency request that targets the processor. The processor issues a processor coherency data response for each external intervention request that targets the processor and hits a *DirtyExclusive* secondary cache block.

Processor coherency state responses are issued by the processor in the same order that the corresponding external coherency requests are received. Processor coherency state and data responses may occur in adjacent **SysClk** cycles.

Processor Coherency State Response Protocol

A processor coherency state response results from an external coherency request that targets the processor.

The processor issues a processor coherency state response by driving the secondary cache block tag quality indication on **SysState[2]**, driving the secondary cache block former state on **SysState[1:0]**, and asserting **SysStateVal*** for one **SysClk** cycle.

Figure 6-24 depicts two external coherency requests and the resulting processor coherency state responses.

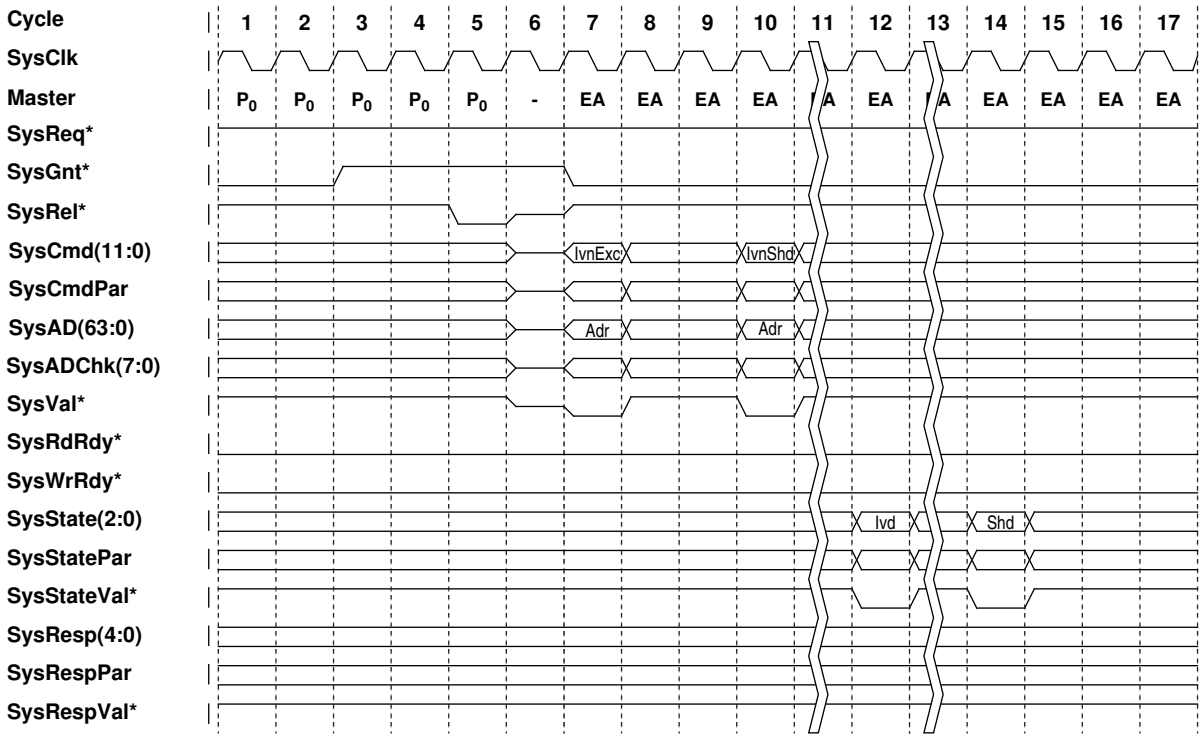

Figure 6-24 Processor Coherency State Response Protocol

Processor Coherency Data Response Protocol

A processor coherency data response results from an external intervention request that targets the processor and hits a *DirtyExclusive* secondary cache block.

The processor issues a processor coherency data response with a single empty cycle followed by either 8 or 16 data cycles. The empty cycle consists of negating **SysVal*** for a single **SysClk** cycle. The data cycles consist of the following:

- asserting **SysCmd[11]**
- driving the request number associated with the corresponding external coherency request on **SysCmd[10:8]**
- driving the data quality indication on **SysCmd[5]**
- driving the data type indication on **SysCmd[4:3]**
- driving the state of the cache block on **SysCmd[2:1]**
- asserting **SysCmd[0]**
- driving the data on **SysAD[63:0]**,
- asserting **SysVal***

The first 7 or 15 data cycles have a response data type indication, and the last data cycle has a response last data indication. The processor may negate **SysVal*** between data cycles of a processor coherency data response only if the **SCClk** frequency is less than half of the **SysClk** frequency.

The processor may only issue a processor coherency data response when the System interface is in master state and **SysWrRdy*** was asserted two **SysClk** cycles previously. Note that the empty cycle is considered the issue cycle for a processor coherency data response. If the System interface is not already in master state, the processor must first assert **SysReq***, and then wait for the external agent to relinquish mastership of the System interface bus by asserting **SysGnt*** and **SysRel***. If the System interface is already in master state, the processor may issue a processor coherency data response immediately.

When **SysStateVal*** is negated, **SysState[0]** indicates if a processor coherency data response is ready to be issued. Processor coherency data response data is supplied in subblock order, beginning with the quadword-aligned address specified by the corresponding external coherency request.

Figure 6-25 depicts one external coherency request and the resulting processor coherency state and data responses.

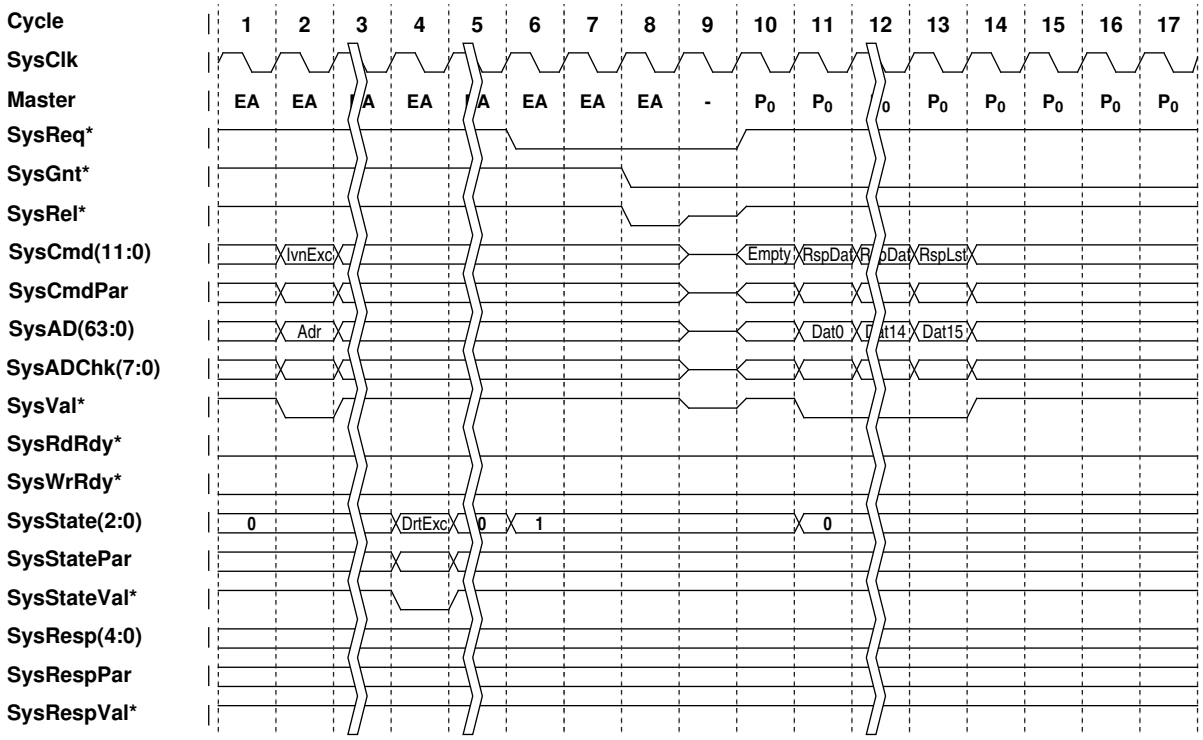

Figure 6-25 Processor Coherency Data Response Protocol

6.18 System Interface Coherency

The System interface supports external intervention shared, intervention exclusive, and invalidate coherency requests. These requests are used by an external agent or other R10000 processors on the cluster bus to maintain cache coherency.

Each external coherency request that targets an R10000 results in a processor coherency state response. Additionally, each external intervention request that targets the R10000 and hits a *DirtyExclusive* secondary cache block results in a processor coherency data response.

External coherency requests and the corresponding processor coherency state responses are handled in FIFO order.

External Intervention Shared Request

An external intervention shared request is used by an external agent to obtain a *Shared* copy of a cache block. If the desired block resides in the processor cache, it is marked *Shared*.

If the secondary cache block's former state was *DirtyExclusive*, the processor issues a processor coherency data response.

External Intervention Exclusive Request

An external intervention exclusive request is used by an external agent to obtain an *Exclusive* copy of a cache block. If the desired block resides in the processor cache, it is marked *Invalid*.

If the secondary cache block's former state was *DirtyExclusive*, the processor issues a processor coherency data response.

External Invalidate Request

An external invalidate request is used by an external agent to invalidate a cache block. If the desired block resides in the processor cache, it is marked *Invalid*.

Under normal circumstances, the secondary cache block former state should not be *CleanExclusive* or *DirtyExclusive*.

External Coherency Request Action

Table 6-27 indicates the action taken for external coherency requests that target the processor.

Table 6-27 Action Taken for External Coherency Requests that Target the R10000 Processor

Secondary Cache Block Former State	Type of External Request	Secondary Cache Block New State	Processor Coherency State Response SysState[1:0]	Processor Coherency Data Response Required?	Processor Coherency Data Response State SysCmd[2:1]
Invalid	Intervention shared	*Invalid*	0	No	N/A
	Intervention exclusive	*Invalid*	0	No	N/A
	Invalidate	*Invalid*	0	No	N/A
Shared	Intervention shared	*Shared*	1	No	N/A
	Intervention exclusive	*Invalid*	1	No	N/A
	Invalidate	*Invalid*	1	No	N/A
CleanExclusive	Intervention shared	*Shared*	2	No	N/A
	Intervention exclusive	*Invalid*	2	No	N/A
	Invalidate*	*Invalid*	2	No	N/A
DirtyExclusive	Intervention shared†	*Shared*	3	Yes	*Shared*
	Intervention exclusive	*Invalid*	3	Yes	*DirtyExclusive*
	Invalidate*	*Invalid*	3	No	N/A

* This should not occur under normal circumstances.

† The processor coherency data response must be written back to memory.

Coherency Conflicts

Coherency conflicts arise when a processor request and an external request target the same secondary cache block. Coherency conflicts may be categorized as either internal or external, and are described in this section.

Internal Coherency Conflicts

A processor request is considered to be **pending issue** when it is buffered in the processor and has not yet been issued to the System interface bus. Internal coherency conflicts occur when the processor has a processor request pending issue and a conflicting external coherency request is received. Internal coherency conflicts are unavoidable and cannot be anticipated by the external agent since it cannot anticipate when the processor will have processor requests pending issue.

Table 6-28 describes the manner in which the processor resolves internal coherency conflicts.

Table 6-28 Internal Coherency Conflict Resolution

Processor Request Pending Issue	Conflicting External Coherency Request	Resolution
Coherent block read	Intervention shared	The processor allows the conflicting external coherency request to proceed and provides an *Invalid* processor coherency state response. The processor stalls the processor coherent block read request until the conflicting external coherency request has received an external completion response.
	Intervention exclusive	
	Invalidate	
Upgrade	Intervention shared	The processor allows the conflicting external coherency request to proceed and provides a *Shared* processor coherency state response. Once the conflicting external coherency request has received an external completion response, the processor internally NACKs the processor upgrade request that is pending issue.
	Intervention exclusive	
	Invalidate	
Block write	Intervention shared	The processor provides a *DirtyExclusive* processor coherency state response and changes the processor block write request that is pending issue into a *DirtyExclusive* processor coherency data response.
	Intervention exclusive	
	Invalidate	The processor provides a *DirtyExclusive* processor coherency state response and deletes the processor block write request that is pending issue.
Eliminate	Intervention shared	The processor provides a *Shared* or *CleanExclusive* processor coherency state response and deletes the processor eliminate request that is pending issue.*
	Intervention exclusive	
	Invalidate	

* If the processor eliminate request that is pending issue has a *DirtyExclusive* state, a *CleanExclusive* processor coherency state response is provided.

External Coherency Conflicts

A processor request is considered to be **pending response** when it has been issued to the System interface bus but has not yet received an external data or completion response. External coherency conflicts occur when the processor has a processor request that is pending response and a conflicting external coherency request is received. The processor relies on the external agent to detect and resolve external coherency conflicts.

External coherency conflicts may be avoided if the point of coherence is the processor System interface bus and only one request is allowed to be outstanding for any given secondary cache block. However, in some system designs external coherency conflicts are unavoidable.

Processor block write and eliminate requests are never pending response, and therefore cannot cause external coherency conflicts.

Table 6-29 describes the manner in which the external agent resolves external coherency conflicts.

Table 6-29 External Coherency Conflict Resolution

Processor Requests that are Pending Response	Conflicting External Coherency Request	Resolution
Coherent block read	Intervention shared	The external agent responds to the external coherency requestor that the block is *Invalid*. At some later time, the external agent supplies an external response to the processor coherent block read request that is pending response.*
	Intervention exclusive	
	Invalidate	
Upgrade	Intervention shared	The external agent responds to the external coherency requestor that the block is *Shared*. At some later time, the external agent supplies an external response to the processor upgrade request that is pending response.*
	Intervention exclusive	The external agent issues the conflicting external coherency request to the processor. The processor allows the conflicting external coherency request to proceed and supplies a *Shared* processor coherency state response. After observing the processor coherency state response, the external agent provides an external ACK completion response for the conflicting external coherency request. At some later time, the external agent supplies an external response for the processor upgrade request that is pending response. This external response may not be an external ACK completion response unless it is associated with an external block data response.
	Invalidate	

* The external agent may choose to issue the conflicting external coherency request to R10000, although it is not required.

External Coherency Request Latency

This section describes the R10000 external coherency request latency. Figure 6-26 depicts the following:

- an external coherency request which targets the processor
- the resulting processor coherency state response
- the potential processor coherency data response

Two external coherency request latency parameters are also defined:

- the processor coherency state response latency, t_{pcsr}, specifies the time from external coherency request to processor coherency state response
- the processor coherency data response latency, t_{pcdr}, specifies the time from the external coherency request to the processor coherency data response if a master, or to the assertion of the processor coherency data response indication on **SysState[0]** if a slave.

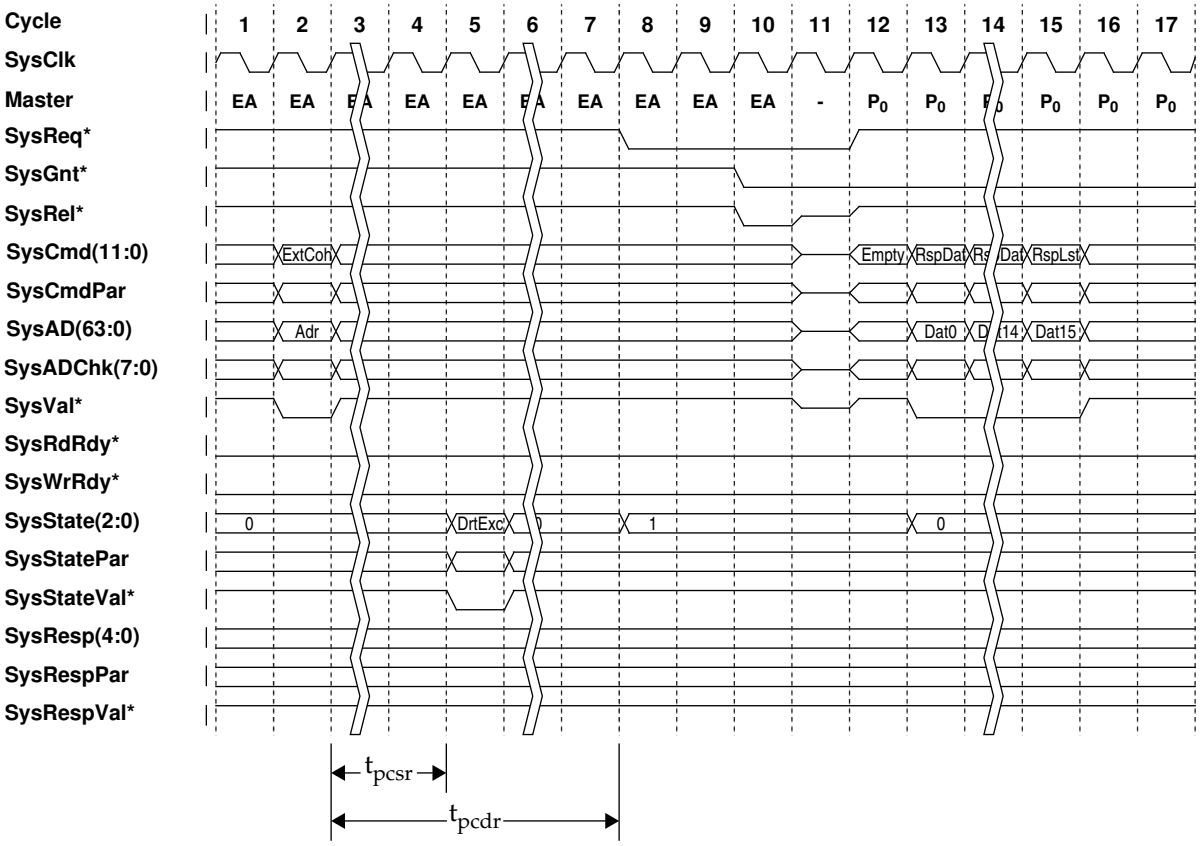

Figure 6-26 External Coherency Request Latency Parameters

The external coherency request latency is presented in Table 6-30.

Table 6-30 External Coherency Request Latency

	Latency* (PClk cycles)					
	Processor Coherency State Response (t_{pcsr})			Processor Coherency Data Response† (t_{pcdr})		
SCClkDiv	**Min‡**	**Typ^**	**Max@**	**Min§**	**Typƒ**	**Max¤**
1	5	10	39	8	28	70
1.5	5	13	48	8	33	88
2	5	14	59	8	38	105
2.5	5	16	71	8	43	128
3	5	17	79	8	43	141

* This latency assumes no other previously issued external coherency requests are outstanding. 1 to 3 additional **PClk** cycles may be required for synchronization with **SysClk** depending on the **SysClkDiv** mode bits.

† This value assumes a 32-word secondary cache block size.

‡ This value assumes the external coherency request hits a cached or outgoing buffer entry.

^ This value assumes the external coherency request does not hit a cached or outgoing buffer entry, the secondary cache is not busy, and the external coherency request hits in the MRU way of the secondary cache. If the external coherency request misses in the most-recently used (MRU) way of the secondary cache, 1 to 3 additional **PClk** cycles are required to query the LRU way of the secondary cache, depending on the **SCClkDiv** mode bits.

@ This value assumes the external coherency request does not hit a cached or outgoing buffer entry, the secondary cache just commenced an index-conflicting CACHE Hit WriteBack Invalidate (S), and the external coherency request misses in the secondary cache MRU way.

§ This value assumes the external coherency request hits an outgoing buffer entry.

ƒ This value assumes the external coherency request does not hit a cached or outgoing buffer entry, the secondary cache is not busy, the external coherency request hits in the MRU way of the secondary cache, no subset primary data cache blocks are inconsistent, and the external coherency request is secondary cache block-aligned. If the external coherency request misses in the MRU way of the secondary cache, 1 to 3 additional **PClk** cycles are required to query the LRU way of the secondary cache, depending on the **SCClkDiv** mode bits.

¤ This value assumes the external coherency request does not hit a cached or outgoing buffer entry, the secondary cache just commenced an index-conflicting CACHE Hit WriteBack Invalidate (S), the external coherency request hits in the LRU way of the secondary cache, all subset primary data cache blocks are inconsistent, and the external coherency request is not secondary cache block-aligned.

SysGblPerf* Signal

The **SysGblPerf*** signal is provided for systems implementing a relaxed consistency memory model. The external agent asserts this signal when all processor requests are globally performed, thereby allowing the processor to graduate SYNC instructions. The external agent negates this signal when some processor requests are not yet globally performed, thereby preventing the processor from graduating SYNC instructions.

To prevent a SYNC instruction from graduating, the external agent must negate the **SysGblPerf*** signal no later than the same **SysClk** cycle in which it issued the external completion response for a processor read or upgrade request which is not yet globally performed. Also, the external agent must negate the **SysGblPerf*** signal no later than two **SysClk** cycles after the address cycle of a processor double/single/partial-word write request which has not yet been globally performed.

The **SysGblPerf*** signal may be permanently asserted in systems implementing a sequential consistency memory model.

6.19 Cluster Bus Operation

A R10000 multiprocessor cluster may be created by directly attaching the System interfaces of 2 to 4 R10000 processors, and providing an external cluster coordinator to handle arbitration and coherency management.

The cluster coordinator arbitrates the multiprocessors using the **SysReq***, **SysGnt***, and **SysRel*** signals.

A processor request issued by an R10000 processor in master state is observed as an external request by any R10000 processors in the slave state on the cluster bus. This is described Table 6-31.

Table 6-31 Relationship Between Processor and External Requests for the Cluster Bus

Processor Request	External Request
Coherent block read shared	Intervention shared
Coherent block read exclusive	Intervention exclusive
Noncoherent block read	Allocate request number
Double/single/partial-word read	Allocate request number
Block write	NOP
Double/single/partial-word write	NOP
Upgrade	Invalidate
Eliminate	NOP

In the same manner, a processor coherency data response issued by a processor in the master state is observed as an external block data response by any processors in the slave state.

External coherency requests that target a processor are handled in FIFO order and result in processor coherency state responses. If an external coherency request that targets a processor hits a *DirtyExclusive* secondary cache block, the processor also provides a processor coherency data response.

Figure 6-27 presents an example of a processor read request with four R10000 processors residing on the cluster bus. The **CohPrcReqTar** mode bit is asserted for a snoopy-based coherency protocol. $R10000_0$ issues a processor coherent read exclusive request. This is observed as an external intervention exclusive request by $R10000_1$, $R10000_2$, and $R10000_3$. $R10000_1$ and $R10000_3$ respond with *Invalid* processor coherency state responses. $R10000_2$ responds with a *DirtyExclusive* processor coherency state response. Based on these processor coherency state responses, the cluster coordinator allows $R10000_2$ to become master of the System interface so that it may provide a processor coherency data response, which will be observed as an external block data response by $R10000_0$. Finally, the cluster coordinator issues an external ACK completion response to forward the external block data response and to free the request number.

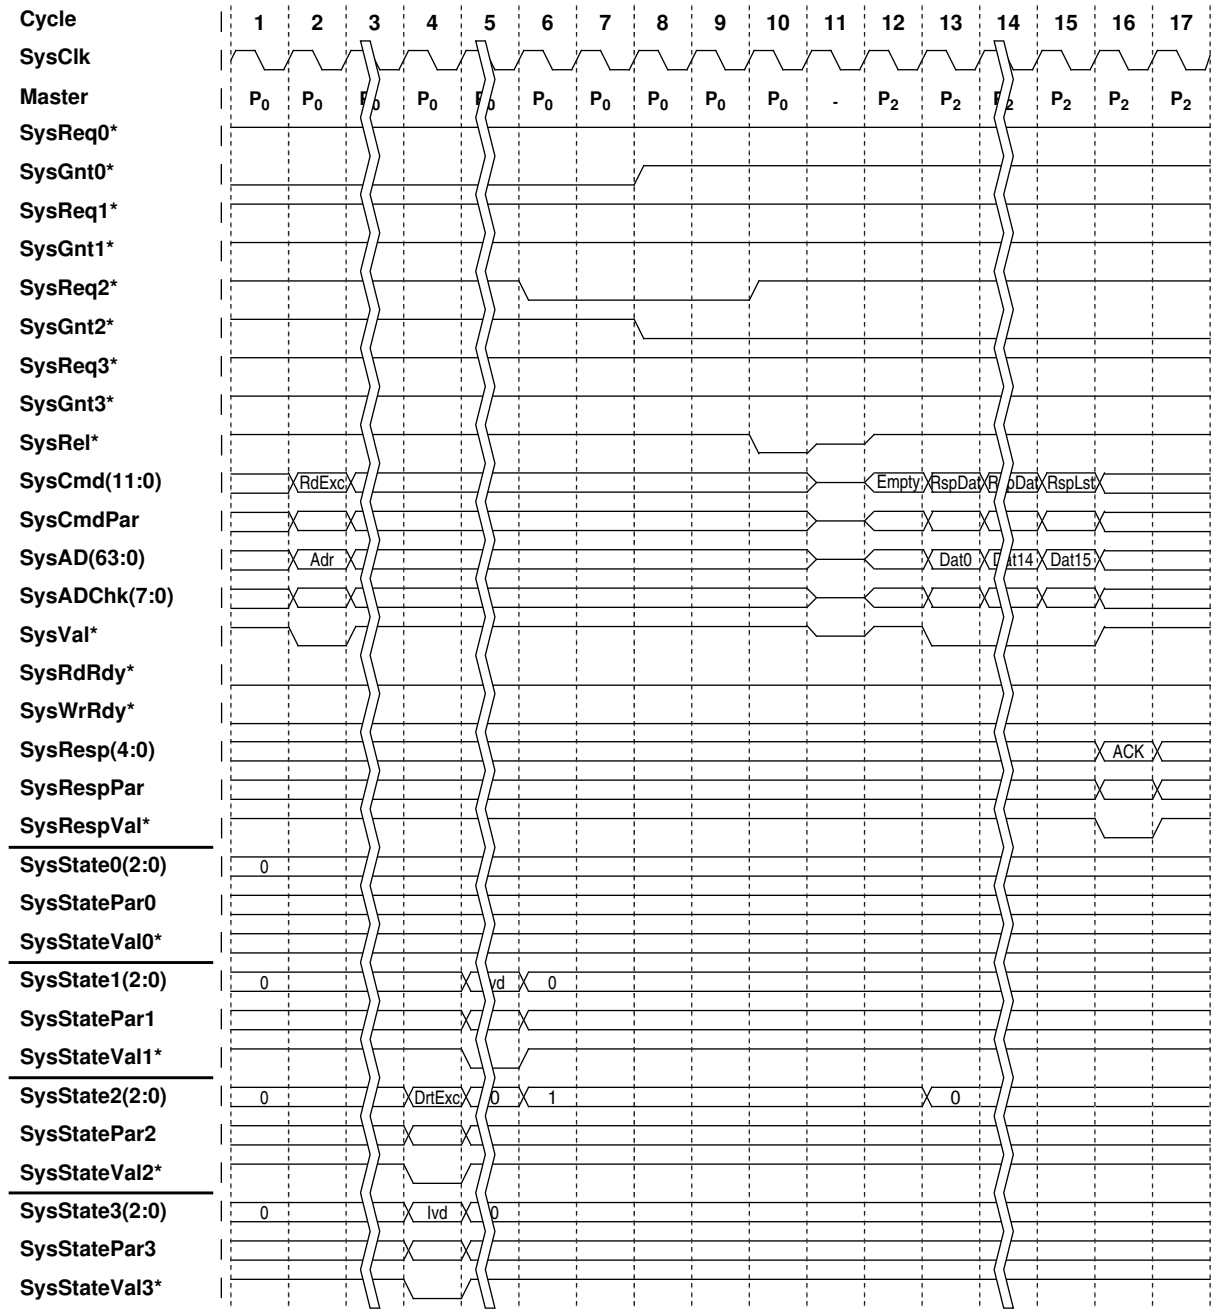

Figure 6-27 R10000 Multiprocessor Cluster Processor Read Request Example

Figure 6-28 presents an example of a processor upgrade request with four R10000 processors residing on the cluster bus. The **CohPrcReqTar** mode bit is asserted for a snoopy-based coherency protocol. $R10000_0$ issues a processor upgrade request, observed as an external invalidate request by $R10000_1$, $R10000_2$, and $R10000_3$. $R10000_2$ and $R10000_3$ provide *Shared* processor coherency state responses. $R10000_1$ provides an *Invalid* processor coherency state response. Based

on these processor coherency state responses, the cluster coordinator issues an external ACK completion response for the processor upgrade request to indicate that the request was successful and to free the request number.

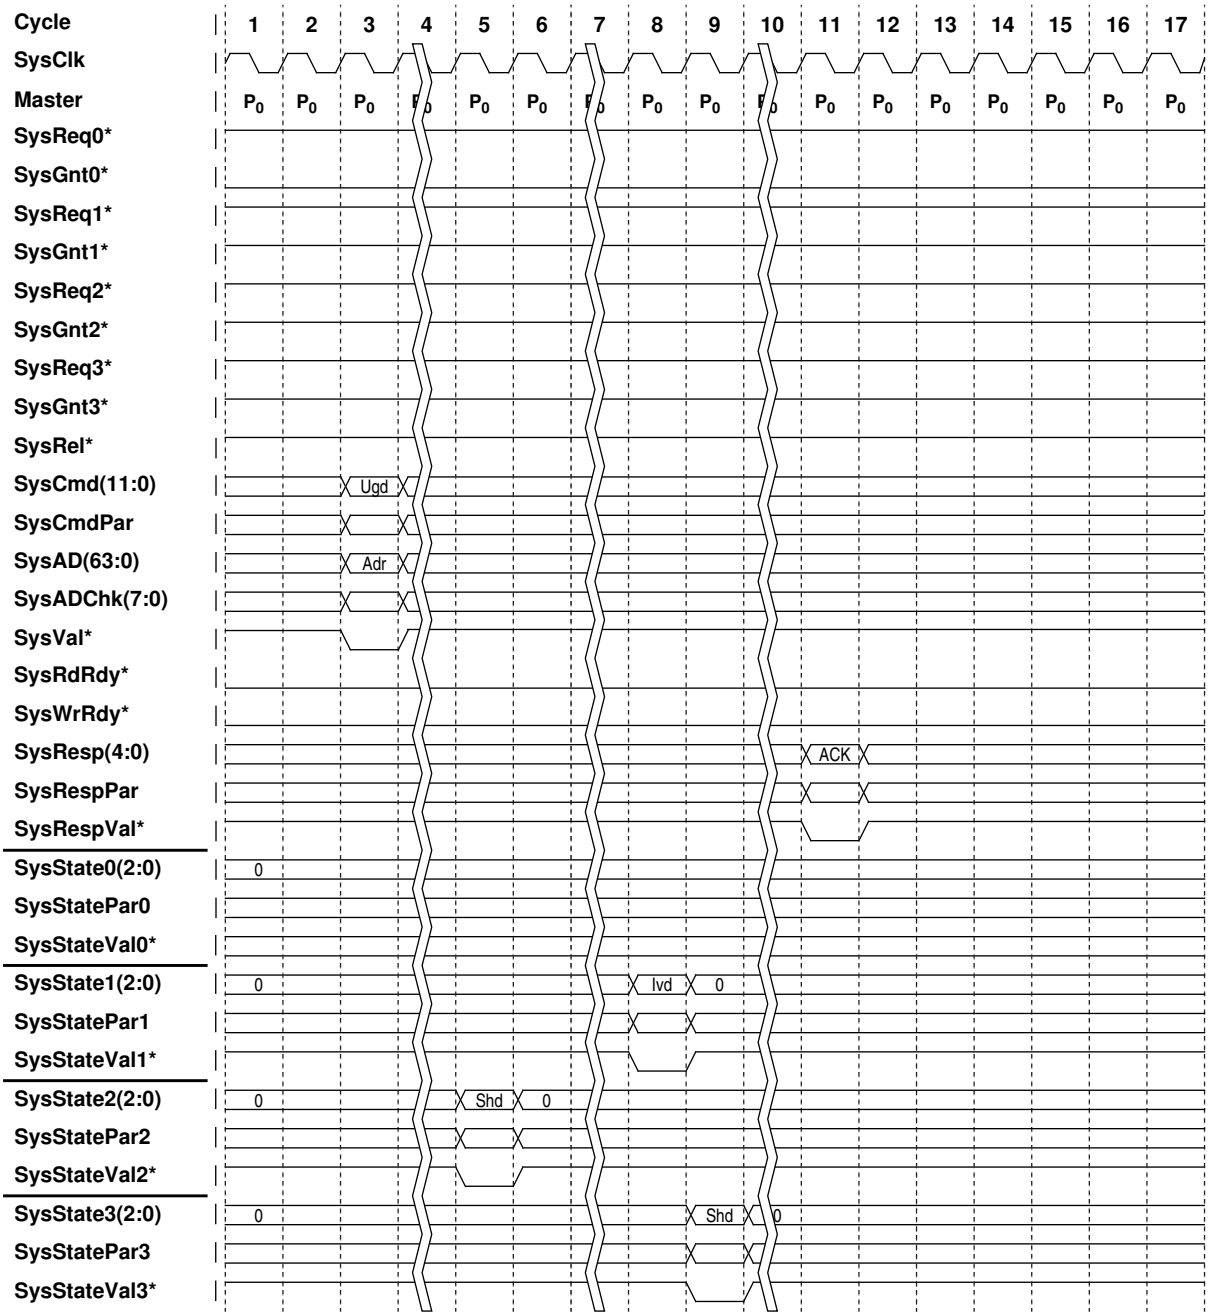

Figure 6-28 R10000 Multiprocessor Cluster Processor Upgrade Request Example

6.20 Support for I/O

The processor assumes a memory-mapped I/O model. Consequentially, no special System interface encodings are provided, or required to designate I/O accesses. It is left to the programmer to ensure that I/O addresses have the appropriate TLB mappings.

The processor supports system designs utilizing hardware or software for coherent I/O. The external coherency requests are useful for creating systems with hardware I/O coherency, and the CACHE instruction is sufficient for creating a system with software I/O coherency.

6.21 Support for External Duplicate Tags

Some system designs implement an external duplicate copy of the secondary cache tags to reduce the coherency request latency and also filter out unnecessary external coherency requests made to the R10000 processor.

For such systems, it must be remembered that blocks may reside in either the secondary cache or in the outgoing buffer. During the address cycle of processor block read requests, the secondary cache block former state is provided. The external agent may use this information to maintain the external duplicate tags.

Typically, in a multiprocessor system using the cluster bus, the cluster coordinator specifies a free request number for an external coherency request. However, in a system using a duplicate-tag or directory-based coherency protocol, where the **CohPrcReqTar** mode bit is negated, the cluster coordinator may specify a busy request number for an external coherency request, providing each targeted R10000 processor has the request number busy due to an outstanding processor coherency request from another processor.

For example, suppose the processor in master state issues a processor coherent block read or upgrade request. The processors in slave state observe the processor request as an external coherency request that targets the external agent only, causing the associated request number to become busy. The cluster coordinator checks the duplicate tag or directory structure to determine if the block resides in the cache of one of the processors that was in slave state. If necessary, the cluster coordinator issues an external coherency request targeted at one or more of the processors that were in slave state. By using the same request number as the original processor request, this external coherency request does not consume a free request number, and allows a potential processor coherency data response to be supplied as an external block data response to the original processor request.

6.22 Support for a Directory-Based Coherency Protocol

Some system designs implement a directory-based coherency protocol.

For such systems, the processor provides the processor eliminate request cycle. If the **PrcElmReq** mode bit is asserted, the processor issues a processor eliminate request whenever it intends to eliminate a *Shared*, *CleanExclusive*, or *DirtyExclusive* block from the secondary cache. During the address cycle of the processor eliminate request, the physical address and the secondary cache block former state are provided. The external agent may then use this information to maintain an external directory structure.

6.23 Support for Uncached Attribute

The processor supports a 2-bit user-defined *Uncached Attribute*, which is driven on **SysAD[59:58]** during the address cycle of the following:

- processor double/single/partial-word read requests
- double/single/partial-word write requests
- block write requests resulting from completely gathered uncached accelerated blocks

For unmapped accesses, the uncached attribute is sourced from **VA[58:57]**.

For mapped accesses, the uncached attribute is sourced from the TLB *Uncached Attribute* field. The TLB *Uncached Attribute* field may be initialized in 64-bit mode using bits 63:62 of the CP0 *EntryLo0* and *EntryLo1* registers.

6.24 Support for Hardware Emulation

When using the R10000 processor in hardware emulation, it is desirable to operate the System interface at a relative low frequency (typically 1 MHz or below). Since the R10000 processor contains dynamic circuitry, an external agent cannot simply provide low frequency **SysClk**, so a **SysCyc*** input to the processor allows an external agent to define a virtual system clock, and yet supply a **SysClk** within the acceptable operating range. The assertion of **SysCyc*** in a particular **SysClk** cycle creates a virtual system clock pulse four **SysClk** cycles later. **SysCyc*** may be asserted aperiodically.

In a normal system environment, the **SysCyc*** input should be permanently asserted.

Figure 6-29 depicts the use of **SysCyc*** to create a virtual **SysClk** of one-third the normal **SysClk** frequency.

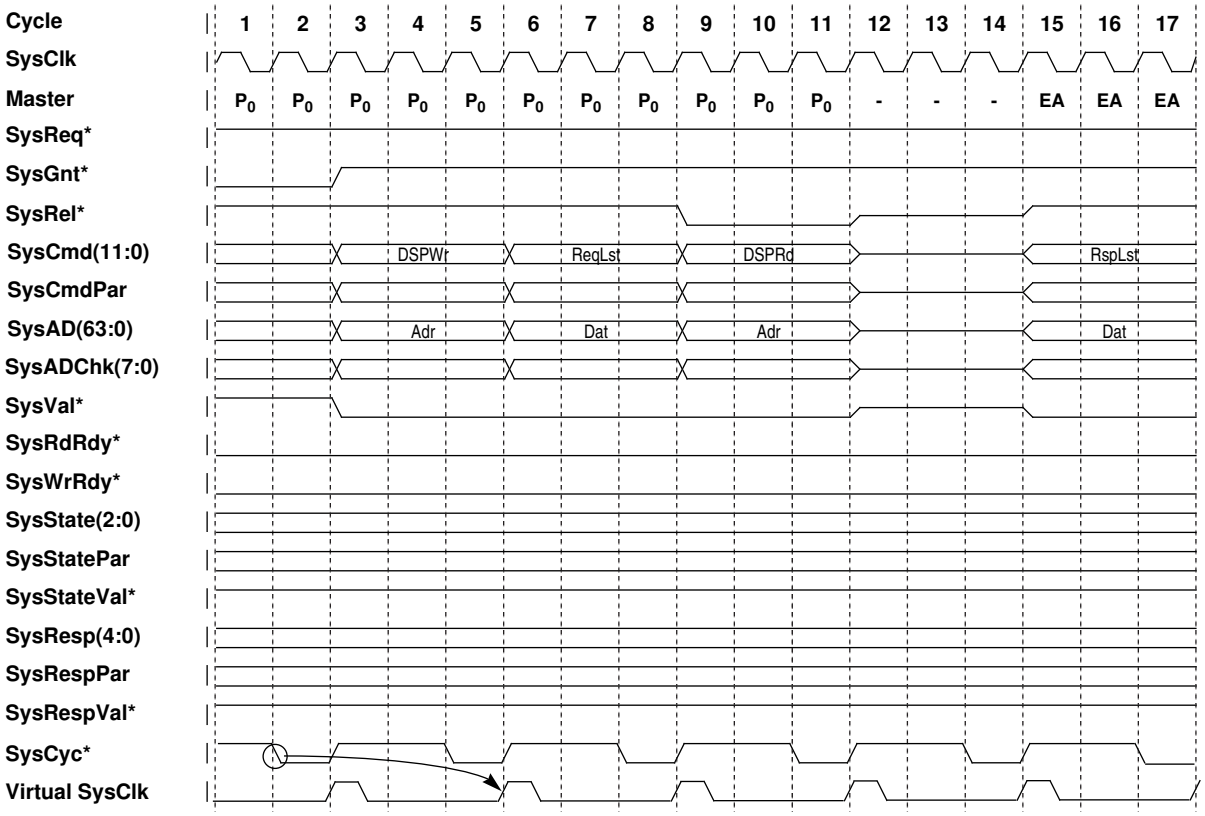

Figure 6-29 Hardware Emulation Protocol

Clock Signals

7

The R10000 processor has differential PECL clock inputs, **SysClk** and **SysClk***, from which all processor internal clock signals and secondary cache clock signals are derived.

Three major clock domains are in the processor:

- the **System interface clock domain**, which operates at the system clock frequency and controls the System interface signals
- the **internal processor clock domain**, which controls the processor core logic
- the **secondary cache clock domain**, which controls signals communicating with the external secondary cache synchronous SRAM

These domains are described in this chapter.

7.1 System Interface Clock and Internal Processor Clock Domains

In high performance systems, PECL-level differential clocks are routinely used to minimize system clock skews. The R10000 processor receives differential system clock signals at the **SysClk** and **SysClk*** pins; two additional pins, **SysClkRet** and **SysClkRet***, are the return paths for termination of these signals.

SysClk and **SysClk*** are used to drive an on-chip phase-locked loop (PLL), which multiplies the system clock to create an internal processor clock, **PClk**.

The R10000 processor always communicates with the system at the **SysClk** frequency, and **PClk** always runs at a frequency-multiple of **SysClk**, according to the following formula:

```
PClk = SysClk*(SysClkDiv+1)/2
```

For example, in a 50 MHz system with **SysClkDiv** = 7 and **SCClkDiv**=2, **PClk**= 50*8/2 = 200 MHz.

> **NOTE:** It is preferred that the R10000 processor uses a differential PECL clock input. However, in a less-aggressive system, a CMOS/TTL single-ended clock can be used to drive the processor, provided its complementary clock input, **SysClk***, is tied to an appropriate reference voltage (1.4V for TTL, Vcc/2 for CMOS). In any case, the reference voltage applied to SysClk* should not be less than 1.2V.

7.2 Secondary Cache Clock

The processor uses registered synchronous SRAMs for its secondary cache, to allow pipelined accesses.

The processor provides 6 pairs of differential **HSTL** clock outputs, **SCClk(5:0)** and **SCClk*(5:0)**, to be used by the secondary cache synchronous SRAMs. The **SCClkTap** mode bits (mode bits are described in Chapter 8) specify the alignment of **SCClk(5:0)** and **SCClk*(5:0)** relative to the internal secondary cache clock. Note that the output buffer delay is not included.

The secondary cache interface clock is generated by dividing down the internal processor clock, **PClk**.

SCClk is related to **SysClk** according to the following formula:

```
SCClk = SysClk*(SysClkDiv+1)/(SCClkDiv+1)
```

For example, in a 50 MHz system with **SysClkDiv**=7 and **SCClkDiv**=2, **SCClk** = 50*8/3 = 133 MHz.

7.3 Phase-Locked-Loop

The processor uses the internal PLL for clock generation and multiplication as shown in Figure 7-1.

Values of the termination resistors for the **SysClkRet/SysClkRet*** signals are system-dependent. The system designer must select a value based upon the characteristic impedance of the board, therefore it is beyond the scope of this manual to specify values for these termination resistors.

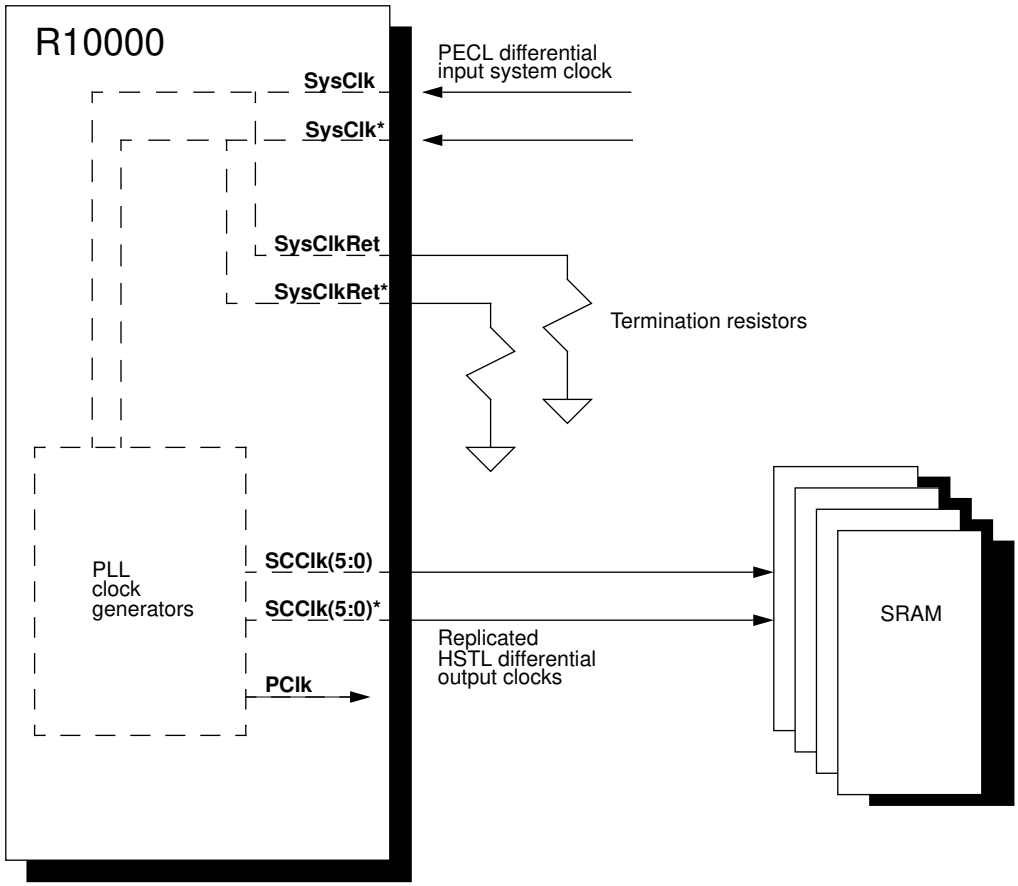

Figure 7-1 R10000 System and Secondary Cache Clock Interface

Initialization

8

This section describes initialization of the R10000 processor, including initialization of logical registers.

Initialization of the processor occurs during a reset sequence. The processor supports three separate reset sequences:

- Power-on reset
- Cold reset
- Soft reset

These sequences are described in this chapter.

Also described are the mode bits.

8.1 Initialization of Logical Registers

After a power-on or cold reset sequence, all logical registers (both in the integer and the floating-point register files) must be written before they can be read. Failure to write any of these registers before reading from them will have an unpredictable result.

8.2 Power-On Reset Sequence

The Power-on Reset sequence is used to reset the processor after the initial power-on, or whenever power or **SysClk** are interrupted.

The Power-on Reset sequence is as follows:

- The external agent negates **DCOk**.
- The external agent asserts **SysReset***.
- The external agent negates **SysGnt***.
- The external agent negates **SysRespVal***.
- Once **Vcc**, **VccQ[SC,Sys]**, **Vref[SC,Sys]**, **Vcc[Pa,Pd]**, and **SysClk** stabilize, the external agent waits at least 1ms and then asserts **DCOk**.
- At this time, the System interface resides in slave state and all internal state is initialized.
- The **SysClkDiv** mode bits default to divide-by-1.
- The **SCClkDiv** mode bits default to divide-by-3.
- After waiting at least 100 ms for the internal clocks to stabilize, the external agent loads the mode bits into the processor by driving the mode bits on **SysAD(63:0)**, waiting at least two **SysClk** cycles, and then asserting **SysGnt*** for at least one **SysClk** cycle.
- After waiting at least another 100 ms for the internal clocks to restabilize, the external agent synchronizes all clocks internal to the processor. This is performed by asserting **SysRespVal*** for one **SysClk** cycle.
- After waiting at least 100 ms for the internal clocks to again restabilize, (a third 100 ms restabilization period) the external agent negates **SysReset***.
- The external agent must retain mastership of the System interface, refrain from issuing external requests or nonmaskable interrupts, and ignore the system state bus until the processor asserts **SysReq***. The assertion of **SysReq*** indicates the processor is ready for operation. In a cluster arrangement, all processors must assert **SysReq***, indicating they are ready for operation.

During a Power-on Reset sequence, all internal state is initialized. A Power-on Reset sequence causes the processor to start with the Reset exception.

Figure 8-1 shows the Power-on Reset sequence.

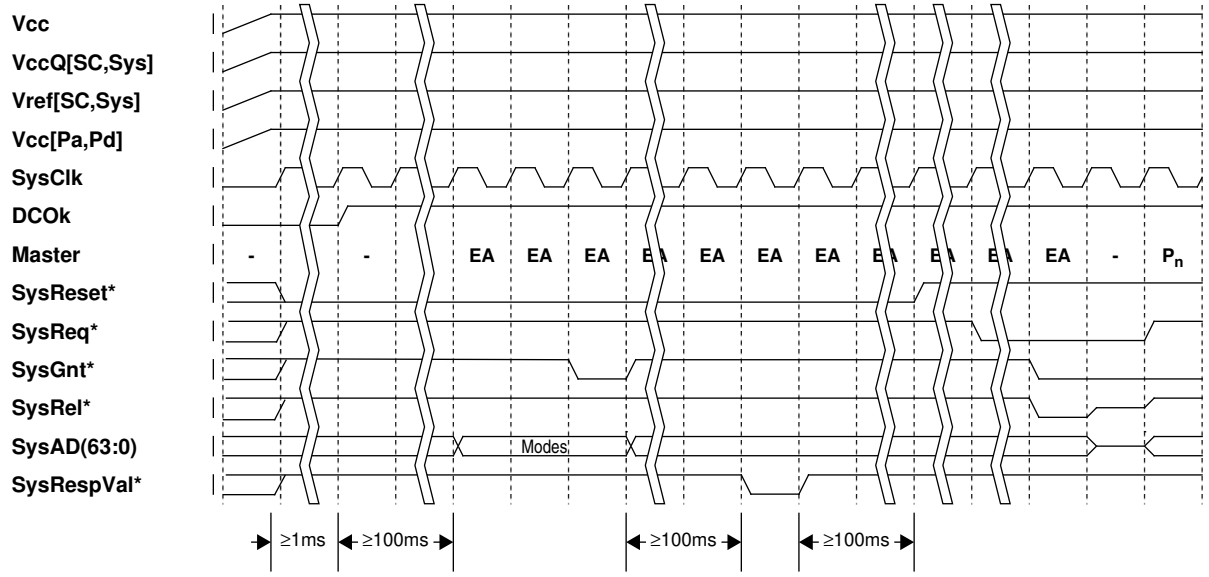

Figure 8-1 Power-On Reset Sequence

8.3 Cold Reset Sequence

The Cold Reset sequence is used to reset the entire processor, and possibly alter the mode bits while power and **SysClk** are stable.

The Cold Reset sequence is as follows:

- The external agent negates **SysGnt*** and **SysRespVal***.
- After waiting at least one **SysClk** cycle, the external agent asserts **SysReset***.
- After waiting at least 100 ms, the external agent loads the mode bits into R10000. This is performed by driving the mode bits on **SysAD(63:0)**, waiting at least two **SysClk** cycles, and then asserting **SysGnt*** for at least one **SysClk** cycle.
- After waiting at least another 100 ms for the internal clocks to restabilize, the external agent synchronizes all processor internal clocks by asserting **SysRespVal*** for one **SysClk** cycle.
- After waiting at least 100 ms for the internal clocks to again restabilize, (a third 100 ms restabilization period) the external agent negates **SysReset***.
- The external agent must retain mastership of the System interface, refrain from issuing external requests or nonmaskable interrupts, and ignore the system state bus until the processor asserts **SysReq***. The assertion of **SysReq*** indicates the processor is ready for operation. In a cluster arrangement, all processors must assert **SysReq***, indicating they are ready for operation.

During a Cold Reset sequence all processor internal state is initialized. A Cold Reset sequence causes the processor to start with a Reset exception.

Figure 8-2 shows the cold reset sequence.

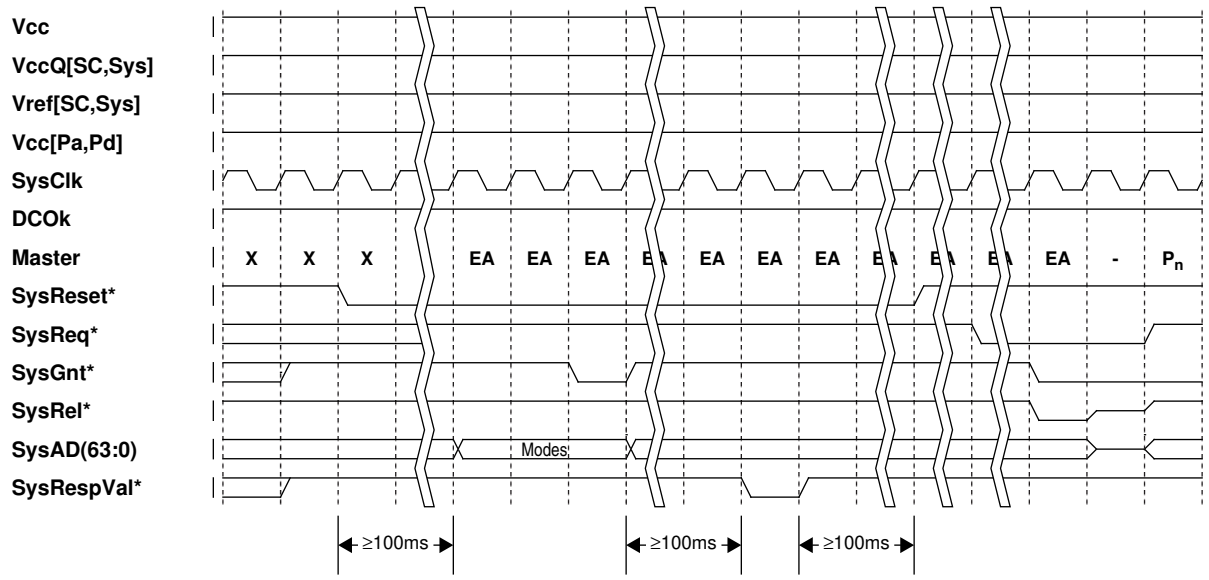

Figure 8-2 Cold Reset Sequence

8.4 Soft Reset Sequence

A Soft Reset sequence is used to reset the external interface of the processor without altering the mode bits while power and **SysClk** are stable.

The Soft Reset sequence is as follows:

- The external agent negates **SysGnt*** and **SysRespVal***.
- After waiting at least one **SysClk** cycle, the external agent asserts **SysReset*** for at least 16 **SysClk** cycles.
- The external agent must retain mastership of the System interface, refrain from issuing external requests or nonmaskable interrupts, and ignore system state bus until the processor asserts **SysReq***. The assertion of **SysReq*** indicates the processor is ready for operation. In a cluster arrangement, all processors must assert **SysReq***, indicating they are ready for operation.

During a Soft Reset sequence, all external interface state is initialized. The internal and secondary cache clocks are not affected by a Soft Reset sequence. The general purpose, CP0, and CP1 registers are preserved, as well as the primary and secondary caches.

A Soft Reset sequence causes a Soft Reset exception, in which the Soft Reset exception handler executes instructions from uncached space and uses CACHE instructions to analyze and dump the contents of the primary and secondary caches. To resume normal operation, a Cold Reset sequence must be initiated.

Figure 8-3 presents the Soft Reset sequence.

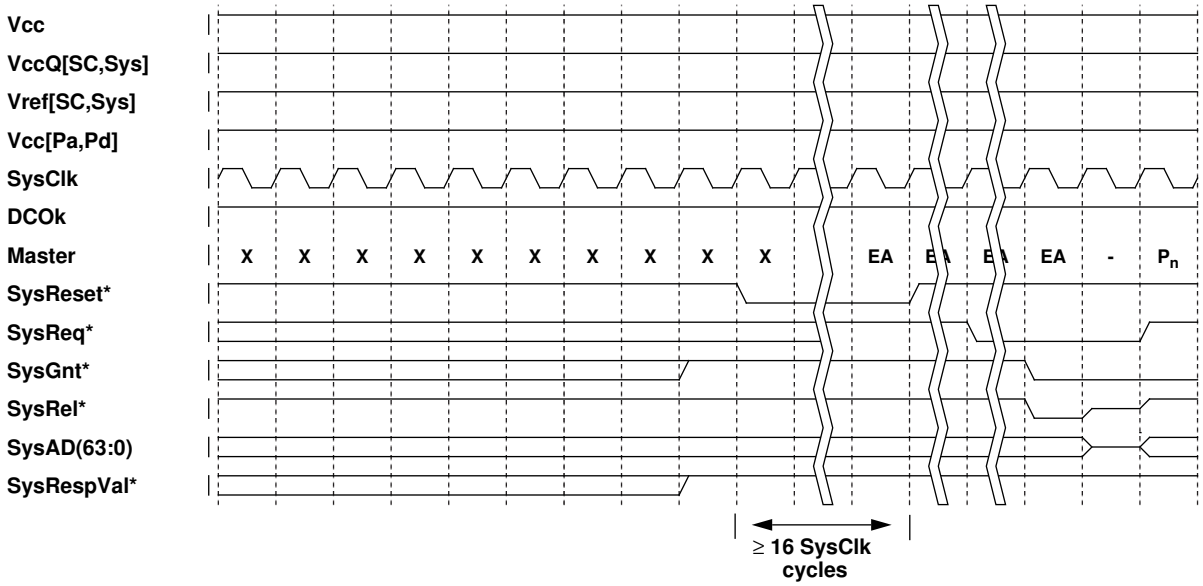

Figure 8-3 Soft Reset Sequence

8.5 Mode Bits

The R10000 processor uses mode bits to configure the operation of the microprocessor. These mode bits are loaded into the processor from the **SysAD(63:0)** bus during a power-on or cold reset sequence while **SysGnt*** is asserted. The **SysADChk(7:0)** bus does not have to contain correct ECC during mode bit initialization. During the reset sequence, the mode bits obtained from **SysAD(24:0)** are written into bits 24:0 of the CP0 *Config* register.

The mode bits are described in Table 8-1.

Table 8-1 Mode Bits

SysAD Bit	Name and Function	Value	Mode Setting
2:0	**Kseg0CA** Specifies the *kseg0* cache algorithm.	0 1 2 3 4 5 6 7	Reserved Reserved Uncached Cacheable noncoherent Cacheable coherent exclusive Cacheable coherent exclusive on write Reserved Uncached accelerated
4:3	**DevNum** Specifies the processor device number.	0-3	
5	**CohPrcReqTar** Specifies the target of processor coherent requests issued on the System interface by the processor.	0 1	External agent only Broadcast
6	**PrcElmReq** Specifies whether to enable processor eliminate requests onto the System interface by the processor.	0 1	Disable Enable
8:7	**PrcReqMax** Specifies the maximum number of outstanding processor requests allowed on the System interface by the processor.	0 1 2 3	1 outstanding processor request 2 outstanding processor requests 3 outstanding processor requests 4 outstanding processor requests

Table 8-1 (cont.) Mode Bits

SysAD Bit	Name and Function	Value	Mode Setting
12:9	**SysClkDiv** Sets **PClk** to **SysClk** ratio; determines the System interface clock frequency; see Chapter 7.	0 1 2 3 4 5 6 7 8 9 A B C D E F	Reserved Result of division by 1 Result of division by 1.5 Result of division by 2 Result of division by 2.5 Result of division by 3 Result of division by 3.5 Result of division by 4 Reserved Reserved Reserved Reserved Reserved Reserved Reserved Reserved
13	**SCBlkSize** Specifies the secondary cache block size.	0 1	16-word 32-word
14	**SCCorEn** Specifies the method of correcting secondary cache data array ECC errors.	0 1	Retry access through corrector Always access through corrector
15	**MemEnd** Specifies the memory system endianness.	0 1	Little endian Big endian
18:16	**SCSize** Specifies the size of the secondary cache.	0 1 2 3 4 5 6 7	512 Kbyte 1 Mbyte 2 Mbyte 4 Mbyte 8 Mbyte 16 Mbyte Reserved Reserved
21:19	**SCClkDiv** Sets **PClk** to **SCClk** ratio; determines the secondary cache clock frequency; see Chapter 7.	0 1 2 3 4 5 6 7	Reserved Result of division by 1 Result of division by 1.5 Result of division by 2 Result of division by 2.5 Result of division by 3 Reserved Reserved
24:22	Reserved	0	

Table 8-1 (cont.) Mode Bits

SysAD Bit	Name and Function	Value	Mode Setting
28:25	**SCClkTap** Specifies the alignment† of **SysClk(5:0)** and **SysClk*(5:0)** relative to the internal secondary cache clock.	0 1 2 3 4 5 6 7 8 9 A B C D E F	**SCClk** same phase as internal clock **SCClk** 1/12 **PClk** period earlier than internal clock **SCClk** 2/12 **PClk** period earlier than internal clock **SCClk** 3/12 **PClk** period earlier than internal clock **SCClk** 4/12 **PClk** period earlier than internal clock **SCClk** 5/12 **PClk** period earlier than internal clock undefined undefined **SCClk** 6/12 **PClk** period earlier than internal clock **SCClk** 7/12 **PClk** period earlier than internal clock **SCClk** 8/12 **PClk** period earlier than internal clock **SCClk** 9/12 **PClk** period earlier than internal clock **SCClk** 10/12 **PClk** period earlier than internal clock **SCClk** 11/12 **PClk** period earlier than internal clock undefined undefined
29	Reserved	0	
30	**ODrainSys** Specifies whether or not to configure select‡ System interface bidirectional and output signals as open drain.	0 1	Push-pull Open drain
31	**CTM** Specifies whether or not to enable cache test mode.	0 1	Disable Enable
63:32	Reserved	0	

†Does not include the output buffer delay.

‡ **SysReq***, **SysRel***, **SysCmd(11:0)**, **SysCmdPar**, **SysAD(63:0)**, **SysADChk(7:0)**, **SysVal***, **SysState(2:0)**, **SysStatePar**, **SysStateVal***, **SysCorErr***, **SysUncErr***

Error Protection and Handling

9

This chapter presents the error protection and handling features provided by the R10000 processor.

Two types of errors can occur in an R10000 system:

- correctable
- uncorrectable

The following two sections describe them.

9.1 Correctable Errors

Correctable errors consist of:

- secondary cache tag array correctable ECC errors
- secondary cache data array correctable ECC errors
- System interface address/data bus correctable ECC errors

When the processor detects a correctable error, the error is automatically corrected, and normal operation continues. Secondary cache array scrubbing is not performed.

The processor informs the external agent that a correctable error was detected and then corrected by asserting the **SysCorErr*** signal for one **SysClk** cycle.

9.2 Uncorrectable Errors

Uncorrectable errors consist of:

- Primary instruction cache array parity errors
- Primary data cache array parity errors
- Secondary cache tag array uncorrectable ECC errors
- Secondary cache data array uncorrectable ECC errors
- System interface command bus parity errors
- System interface address/data bus uncorrectable ECC errors
- System interface response bus parity errors

When the processor detects an uncorrectable error, a Cache Error exception is posted. In general, the protection of an uncorrectable error does not disrupt any ongoing operations. However, the instruction fetch and load/store units never use data which contains an uncorrectable error.

The processor informs the external agent that an uncorrectable tag error has been detected by asserting **SysUncErr*** for one **SysClk** cycle.

9.3 Propagation of Uncorrectable Errors

The processor assists the external agent in limiting the propagation of uncorrectable errors in the following manner:

- During external block data response cycles, if the data quality indication on **SysCmd(5)** is asserted, or if an uncorrectable ECC error is encountered on the system address/data bus while the ECC check indication on **SysCmd(0)** is asserted, the processor intentionally corrupts the ECC of the corresponding secondary cache quadword after receiving an external ACK completion response.
- During processor data cycles, the processor asserts the data quality indication on **SysCmd(5)** if the data is known to contain uncorrectable errors. The System interface ECC is never intentionally corrupted; the **SysCmd(5)** bit is used to indicate corrupted data.
- If an uncorrectable cache tag error is detected, the processor asserts **SysUncErr*** for one **SysClk** cycle.
- An external coherency request that detects a secondary cache tag array uncorrectable error asserts the secondary cache block tag quality indication on **SysState(2)** during the corresponding processor coherency state response.
- If an external coherency request requires a processor coherency data response, and a primary data cache tag parity error is encountered during the primary cache interrogation, or a secondary cache tag array uncorrectable error is encountered during the secondary cache interrogation, the processor asserts the data quality indication on **SysCmd(5)** for all doublewords of the corresponding processor coherency data response.

9.4 Cache Error Exception

The processor indicates an uncorrectable error has occurred by asserting a Cache Error exception.

The following four internal units detect and report uncorrectable errors:

- instruction cache
- data cache
- secondary cache
- System interface

Each of these four units maintains a unique local *CacheErr* register.

A Cache Error exception is imprecise; that is, it is not associated with a particular instruction. When any of the four units post a Cache Error exception, completed instructions are graduated before the Cache Error exception is taken. If there are Cache Error exceptions posted from more than one of the units, the exceptions are prioritized in the following order:

1. instruction cache
2. data cache
3. secondary cache
4. System interface.

The corresponding local *CacheErr* register is transferred to the CP0 *CacheErr* register and the CP0 *Status* register *ERL* bit is asserted. Instruction fetching begins from 0xa0000100 or 0xbfc00300, depending on the CP0 *Status* register *BEV* bit. The CP0 *ErrorEPC* register is loaded with the virtual address of the next instruction that has not been graduated, so that execution can resume after the Cache Error exception handler completes.

When ERL=1, the user address region becomes a 2-Gbyte uncached space mapped directly to the physical addresses. This allows the Cache Error handler to save registers directly to memory without having to use a register to construct the address.

The processor does not support nested Cache Error exception handling. While the CP0 *Status* register *ERL* bit is asserted, any subsequent Cache Error exceptions are ignored. However, the detection of additional uncorrectable errors is not inhibited, and additional Cache Error exceptions may be posted.†

† The hardware does not handle the case of multiple Cache Error exceptions in any special manner; caches are refilled as normal, and data forwarded to the appropriate functional units.

9.5 CP0 *CacheErr* Register *EW* Bit

When a unit detects an uncorrectable error, it records information about the error in its local *CacheErr* register and posts a Cache Error exception. If a subsequent uncorrectable error occurs while waiting for the Cache Error exception to be taken and transfer of the local *CacheErr* register to the CP0 *CacheErr* register to complete, the *EW* bit is set in its local *CacheErr* register. Once the Cache Error exception is taken, the *EW* bit in the CP0 *CacheErr* register is set and the Cache Error exception handler now determines that a second error has occurred.

Once the CP0 *CacheErr* register *EW* bit is set, it can only be cleared by a reset sequence.

9.6 CP0 *Status* Register *DE* Bit

Asserting the CP0 *Status* register *DE* bit suppresses the posting of future Cache Error exceptions. All local *CacheErr* registers are also prevented from being updated. Unlike the R4400 processor architecture, when the *DE* bit is asserted, cache hits are not inhibited when an uncorrectable error is detected. Correctable errors are handled normally when the *DE* bit is set.

NOTE: Be careful when setting this bit, since it may cause erroneous data and/or instructions to be propagated.

9.7 CACHE Instruction

Uncorrectable error protection is suppressed for the Index Load Tag, Index Store Tag, Index Load Data, and Index Store Data CACHE instruction variations. These four variations may be used within a Cache Error exception handler to examine the cache tags and data without the occurrence of further uncorrectable errors.

9.8 Error Protection Schemes Used by R10000

Error protection schemes used in the R10000 processor are:

- parity
- sparse encoding
- ECC

These schemes are described in this section, and listed in Table 9-1.

Table 9-1 Error Protection Schemes Used in the R10000 Processor

Error Detection Used	What is Protected
Parity	Primary caches Secondary cache data System interface buses
Sparse encoding	Primary data cache state mod array
ECC (SECDED)	Secondary cache tag Secondary cache data System interface address/data bus

Parity

Parity is used to protect the primary caches and various System interface buses. The processor uses both odd and even parity schemes:

- in an odd parity scheme, the total number of ones on the protected data and the corresponding parity bit should be odd
- in an even parity scheme, the total number of ones on the protected data and the corresponding parity bit should be even.

Sparse Encoding

A sparse encoding is used to protect the primary data cache state mod array. In such a scheme, valid encodings are chosen so that altering a single bit creates an invalid encoding.

ECC

An error correcting code (ECC) is used to protect the secondary cache tag, the secondary cache data, and the System interface address/data bus. A distinct single-bit error correction and double-bit error detection (SECDED) code is used for each of these three applications.

9.9 Primary Instruction Cache Error Protection and Handling

This section describes error protection and error handling schemes for the primary instruction cache.

Error Protection

The primary instruction cache arrays have the following error protection schemes, as listed in Table 9-2.

Table 9-2 Primary Instruction Cache Array Error Protection

Array	Width	Error Protection
Tag Address	27-bit	Even parity
Tag State	1-bit	Even parity
Data	36-bit	Even parity
LRU	1-bit	None

Error Handling

All primary instruction cache errors are uncorrectable. If an error is detected, the instruction cache unit posts a Cache Error exception and initializes the *D*, *TA*, *TS*, and *PIdx* fields in the local *CacheErr* register (see Chapter 14 for more information). If an error is detected on the tag address or state array, the processor informs the external agent that an uncorrectable tag error was detected by asserting **SysUncErr*** for one **SysClk** cycle.

9.10 Primary Data Cache Error Protection and Handling

This section describes error protection and error handling schemes for the primary data cache.

Error Protection

The primary data cache arrays have the following error protection schemes, as listed in Table 9-3.

Table 9-3 Primary Data Cache Array Error Protection

Array	Width	Error Protection
Tag Address	28-bit	Even parity
Tag State	3-bit	Even parity
Tag Mod	3-bit	Sparse encoding
Data	8-bit	Even parity
LRU	1-bit	None

Error Handling

All primary data cache errors are uncorrectable. If an error is detected, the data cache unit posts a Cache Error exception and initializes the *EE*, *D*, *TA*, *TS*, *TM*, and *PIdx* fields in the local *CacheErr* register (see Chapter 14 for more information). If an error is detected on the tag address, state, or mod array, the processor informs the external agent that an uncorrectable tag error was detected by asserting **SysUncErr*** for one **SysClk** cycle.

9.11 Secondary Cache Error Protection and Handling

This section describes error protection and error handling schemes for the secondary cache.

Error Protection

The secondary cache arrays have the following error protection schemes, as listed in Table 9-4.

Table 9-4 Secondary Cache Array Error Protection

Array	Width	Error Protection
Data	128-bit	9-bit ECC + even parity
Tag	26-bit	7-bit ECC
MRU (Way prediction table)	1-bit	None

Error Handling

This section describes error handling for the data array and the tag array. As shown in Table 9-4, errors are not detected for the way prediction table.

Data Array

The 128-bit wide secondary cache data array is protected by a 9-bit wide ECC. An even parity bit is used for rapid detection of correctable (single-bit) errors; when a correctable parity error is detected, the data is sent through the data corrector. The parity bit does not have any logical effect on the processor's ability to either detect or correct errors.

Whenever the processor writes the secondary cache data array, it drives the proper ECC on **SCDataChk(8:0)** and even parity on **SCDataChk(9)**.

Data Array in Correction Mode

The secondary cache operates in correction mode when the **SCCorEn** mode bit is asserted. Whenever the processor reads the secondary cache data array in correction mode, the data is sent through a data corrector.

If a correctable error is detected, in-line correction is automatically made without affecting latency. The processor informs the external agent that a correctable error was detected and corrected by asserting **SysCorErr*** for one **SysClk** cycle.

If an uncorrectable error is detected, the secondary cache unit posts a Cache Error exception and initializes the *D* and *SIdx* fields in the local *CacheErr* register (see Chapter 14 for more information).

In correction mode, secondary-to-primary cache refill latency is increased by two **PClk** cycles. Multiple processors, operating in a lock-step fashion, remain synchronized in the presence of secondary cache data array correctable errors.

Table 9-5 presents the ECC matrix for the secondary cache data array.

Table 9-5 ECC Matrix for Secondary Cache Data Array

Check Bit		8	7654	3210																																				
Data Bit					1111 2222 7654	1111 2222 3210	1111 1111 9876	1111 1111 5432	1111 1100 1098	1111 0000 7654	1111 0000 3210	9999 9876	9999 5432	9988 1098	8888 7654	8888 3210	7777 9876	7777 5432	7766 1098	6666 7654	6666 3210	5555 9876	5555 5432	5544 1098	4444 7654	4444 3210	3333 9876	3333 5432	3322 1098	2222 7654	2222 3210	1111 9876	1111 5432	11 1098	7654	3210
Number of ones per row	54	1	0000	0000	1111	1111	0000	0011	1111	1111	1110	0011	0000	0110	1111	1111	0000	0111	0010	0001	0110	0010	1101	0000	1011	0000	1101	0000	0000	0000	0010	0000	0010	0000	1000	0000
	53	0	1000	0000	1111	1111	1111	1111	0000	0011	0000	1111	1111	1111	0000	0111	0000	0011	1011	0010	0001	0000	0000	1000	0100	1000	1100	1000	0000	0010	1001	0000	1001	0000	0010	0000
	54	0	0100	0000	1000	0011	1111	1111	1111	1111	0000	0011	0000	0101	0000	0011	1111	1111	0101	1100	1000	0000	1000	0100	0000	0100	0100	0100	0011	1000	0100	1000	0100	1000	0101	0000
	53	0	0010	0000	0100	0010	1000	0000	1000	0000	0001	0010	1111	1100	1111	1011	1111	1100	0011	1111	1111	1101	0110	0010	0000	0010	0010	0010	1000	0101	0000	0100	1100	0100	0000	1000
	53	0	0001	0000	0010	0001	0100	0000	0100	0011	1111	1111	1000	0000	1000	0000	1000	0100	1111	1111	1111	1111	0010	0001	0000	0001	0000	0001	1111	1111	1100	0010	0000	0010	0100	0100
	53	0	0000	1000	0001	0000	0010	0011	0010	0000	1000	0001	0100	0100	0100	0000	0100	0110	1011	1111	1111	1100	0011	1111	1101	0011	1111	0100	1100	0000	0001	0001	0000	0001	0000	0010
	54	0	0000	0100	0000	1010	0001	0010	0001	0010	0001	1000	0010	0011	0010	0000	0010	0001	0000	0000	1011	1010	1111	1111	1100	0000	1010	0000	1101	0000	1111	1111	1111	1111	1100	0001
	53	0	0000	0010	0000	0100	0000	1001	0000	1001	0100	0100	0001	0011	0001	0010	0001	0000	0000	1000	0100	1101	1100	0000	1110	0000	1111	1110	0000	1100	0000	1111	1111	1111	1111	1111
	54	0	0000	0001	0000	0001	0000	0100	0000	0100	0010	0000	0000	1011	0000	1101	0000	0100	0100	0111	0000	0100	1110	0000	1111	1111	0110	0000	1111	0011	1111	1100	0000	0000	0001	1111
Number of ones per column		1	1111	1111	3333	3355	3333	3355	3333	3355	3333	3355	3333	3555	3333	3355	3333	3555	3355	5555	5555	5533	5553	3333	5533	3333	5553	3333	5533	3333	5533	3333	5533	3333	5533	3333

Data Array in Noncorrection Mode

When the **SCCorEn** mode bit is negated, the secondary cache operates in noncorrection mode. Whenever the processor reads the secondary cache data array in noncorrection mode, it checks for even parity on **SCDataChk(9)**. If a parity error is detected, it is assumed that a correctable error has occurred, and the secondary cache block is again read through a data corrector. During this re-read, the processor checks the **SCDataChk(8:0)** bus for the proper ECC.

If a correctable error is detected, correction is automatically performed in-line. To inform the external agent that a correctable error had been detected and corrected, the processor asserts **SysCorErr*** for one **SysClk** cycle.

If an uncorrectable error is detected, the secondary cache unit posts a Cache Error exception and initializes the *D* and *SIdx* fields in the local *CacheErr* register.

Secondary cache data array correctable errors are monitored with Performance Counter 0.

Tag Array

The 26-bit-wide secondary cache tag array is protected by a 7-bit-wide ECC. Table 9-6 presents the ECC matrix for the secondary cache tag array.

Table 9-6 ECC Matrix for Secondary Cache Tag Array

Check Bit			0	12	34	56				
Data Bit		25	22 32	22 10	11 98	11 76	1111 5432	11 1098	7654	3210
Number of ones per row	11	0	0100	1000	1000	0001	1111	1000	1000	1000
	13	0	1000	0100	0100	0010	1111	1111	0000	0100
	11	1	0010	1000	0001	1000	0000	1111	0100	0010
	11	1	0100	0100	0010	0100	1000	0100	1111	0000
	13	0	1000	0001	1000	1000	0100	0000	1111	1111
	12	1	0010	0010	0100	0100	0010	0010	0010	1111
	14	0	1111	1100	1100	1100	0001	0001	0001	0001
Number of ones per column		3	3331	3311	3311	3311	3333	3333	3333	3333

Whenever the processor reads the secondary cache tag array, it checks the **SCTagChk(6:0)** bus for the proper ECC. If a correctable error is detected, correction is automatically performed in-line, without affecting latency. The processor asserts **SysCorErr*** for one **SysClk** cycle to inform the external agent that a correctable error has been detected and corrected. If an uncorrectable error is detected, the secondary cache unit posts a Cache Error exception and initializes the *TA* and *SIdx* fields in the local *CacheErr* register. The processor asserts **SysUncErr*** for one **SysClk** cycle to inform the external agent that an uncorrectable tag error has been detected.

Whenever the processor writes the secondary cache tag array, it drives the proper ECC on the **SCTagChk(6:0)** bus.

9.12 System Interface Error Protection and Handling

This section describes error protection and error handling schemes for the System interface.

Error Protection

The System interface buses have the following error protection schemes, as listed in Table 9-7.

Table 9-7 System Interface Bus Error Protection

Bus	Width	Error Protection
SysCmd	12-bit	Odd parity
SysAD	64-bit	8-bit ECC
SysState	3-bit	Odd parity
SysResp	5-bit	Odd parity

Error Handling

This section describes error handling on the system command bus, system address/data bus, system state bus, and system response bus.

SysCmd(11:0) Bus

The 12-bit wide system command bus, **SysCmd(11:0)**, is protected by odd parity.

Whenever the processor is in master state and it asserts **SysVal*** to indicate that it is driving valid information on the **SysCmd(11:0)** bus, it also drives odd parity on the **SysCmdPar** signal.

Whenever the processor is in slave state and an external agent asserts **SysVal*** to indicate that it is driving valid information on the **SysCmd(11:0)** bus, the processor checks the **SysCmdPar** signal for odd parity. If a parity error is detected, the processor ignores the **SysCmd(11:0)** and **SysAD(63:0)** buses for one **SysClk** cycle. The System interface unit posts a Cache Error exception and sets the SC bit in the local *CacheErr* register.

> **Caution:** By ignoring the **SysCmd(11:0)** and **SysAD(63:0)** buses, the processor to become unsynchronized with other processors or the external agent on the cluster bus.

SysAD(63:0) Bus

The 64-bit wide system address/data bus, **SysAD(63:0)**, is protected by an 8-bit-wide ECC.

Processor in Master State

Whenever the processor is in master state and it asserts **SysVal*** to indicate it is driving valid information on the **SysAD(63:0)** bus, it also drives the proper ECC on the **SysADChk(7:0)** bus.

Processor in Slave State

Whenever the processor is in slave state, error checking is enabled with the assertion of **SysCmd(0)**, and an external agent asserts **SysVal*** to indicate it is driving valid information on the **SysAD(63:0)** bus, the processor checks the **SysADChk(7:0)** bus for the proper ECC.

Correctable Error Detected

If a correctable error is detected during an external address cycle, or during an external data cycle for a processor read or upgrade request originated by the R10000 processor, correction is automatically performed in-line without affecting latency. The processor asserts **SysCorErr*** for one **SysClk** cycle to inform the external agent that a correctable error has been detected and corrected.

Uncorrectable Error Detected

If an uncorrectable error is detected during an external address cycle, the processor ignores the **SysCmd(11:0)** and **SysAD(63:0)** buses for one **SysClk** cycle, and the System interface unit posts a Cache Error exception and sets the *SA* bit in the local *CacheErr* register.

> **Caution:** By ignoring the **SysCmd(11:0)** and **SysAD(63:0)** buses, this processor may become unsynchronized with other processors or the external agent on the cluster bus.

If an uncorrectable error is detected or the data quality indication on **SysCmd(5)** is asserted during an external data cycle for a processor read or upgrade request originated by the processor, the R10000 asserts the corresponding incoming buffer uncorrectable error flag.

When the processor forwards block data from an incoming buffer entry after receiving an external ACK completion response, the associated incoming buffer uncorrectable error flags are checked, and if any are asserted, the System interface unit posts a single Cache Error exception and initializes the *EE*, *D*, and *SIdx* fields in the local *CacheErr* register.

When the processor forwards double/single/partial-word data from an incoming buffer entry after receiving an external ACK completion response, the associated incoming buffer uncorrectable error flag is checked and, if asserted, the System interface unit posts a Bus Error exception.

Table 9-8 presents the ECC matrix for the System interface address/data bus. This ECC matrix is identical to that used by the R4X00 System interface.

Table 9-8 ECC Matrix for System Interface Address/Data Bus

Check Bit			43		52											70		61	
Data Bit		6666 3210	55 98	5555 7654	55 32	5544 1098	4444 7654	4444 3210	3333 9876	3333 5432	3322 1098	2222 7654	2222 3210	1111 9876	1111 5432	11 10	9876	54	3210
Number of ones per row	27	1111	1100	1100	1000	1000	0000	1111	1111	0000	1000	1000	1000	1000	0000	1010	0100	1000	1000
	27	1111	1000	1000	1000	0100	0000	0000	0000	1111	0100	0100	0100	0100	1111	1100	1100	1010	0100
	27	0000	1000	1100	1010	0010	1111	1111	0000	0000	0010	0010	0010	0010	1111	1000	1000	1100	0010
	27	0000	1010	0100	1100	0001	1111	0000	1111	1111	0001	0001	0001	0001	0000	1000	1100	1000	0001
	27	1000	0101	0011	0100	0000	1000	1000	1000	1000	1111	1111	0000	1111	1000	1100	0001	0100	0000
	27	0100	1100	0010	0101	1111	0100	0100	0100	0100	0000	0000	1111	1111	0100	0100	0011	0100	0000
	27	0010	0100	0011	1100	1111	0010	0010	0010	0010	1111	0000	0000	0000	0010	0100	0010	0101	1111
	27	0001	0100	0001	0100	0000	0001	0001	0001	0001	0000	1111	1111	0000	0001	0101	0011	1100	1111
Number of ones per column		3333	5511	3333	5511	3333	3333	3333	3333	3333	3333	3333	3333	3333	3333	5511	3333	5511	3333

SysState(2:0) Bus

The 3-bit wide system state bus, **SysState(2:0)**, is protected by odd parity. The processor drives odd parity on the **SysStatePar** signal.

SysResp(4:0) Bus

The 5-bit wide system response bus, **SysResp(4:0)**, is protected by odd parity.

Whenever an external agent asserts **SysRespVal*** to indicate it is driving valid information on the **SysResp(4:0)** bus, the processor checks the **SysRespPar** signal for odd parity. If a parity error is detected, the processor ignores the **SysResp(4:0)** bus for one **SysClk** cycle. The System interface unit posts a Cache Error exception and sets the *SR* bit in the local *CacheErr* register.

> **Caution:** If the processor ignores the **SysResp(4:0)** bus, it may become unsynchronized with other processors or the external agent on the cluster bus. Also, the processor will "hang" if a parity error is detected on the **SysResp[4:0]** bus during an external completion response cycle for a processor double/single/partial-word read request originated by the processor. The external agent may initiate a Soft Reset sequence to obtain the contents of the *CacheErr* register, and the *CacheErr* register will indicate a System interface uncorrectable system response bus error.

Protocol Observation

The processor continuously observes the protocol on the System interface. Table 9-9 presents the supported protocol observations and the associated error handling sequence.

Table 9-9 Protocol Observation

Protocol Observation	Error Handling
External response data cycle with an unexpected request number during an external block data response for a processor block read or upgrade request originated by the processor.	Ignore the external response data cycle
External block data response specifying a *Reserved* cache block state for a processor block read or upgrade request originated by the processor.	Override the cache block state to *CleanExclusive*
External block data response specifying a *Shared* cache block state for a processor coherent block read exclusive or upgrade request originated by the processor.	Override the cache block state to *CleanExclusive*
External completion response specifying a *Reserved* completion indication.	Ignore the external completion response
External ACK completion response for a processor read request originated by the processor that has not received an external data response.	Override the external ACK completion response to a NACK

CACHE Instructions

10

This chapter describes the CacheOps (CACHE†) used in the R10000 processor.

The format of the CACHE instruction is:

```
CACHE op, offset(base)
```

In a CACHE instruction, the 16-bit offset is sign-extended and added to the contents of the general register *base* to form a Virtual Address (VA). The VA is translated to a Physical Address (PA) using the TLB. The 5-bit sub-opcode specifies a cache instruction variation for that address.

†*CacheOp* and *CACHE instruction* are used interchangeably in this text.

10.1 Notes on CACHE Instruction Operations

This section describes the operations of the CACHE instructions in the R10000 processor.

NOTE: The operation of any operation/cache combination not listed below is undefined, and the operation of this instruction on uncached addresses is also undefined.

Virtual Address

The CACHE instruction uses the following portions of the VA to specify a primary cache block and way:

- **VA[13:5]** defines a 32-byte block in the primary data cache array.
- **VA[13:6]** defines a 64-byte block in the primary instruction cache array.
- In both cases, **VA[0]** defines the way needed by Index operations.

Since **VA[0]** is used to indicate the way, it does not cause alignment errors.

When accessing data in the primary caches, **VA[Blocksize-1]** is also used to read or write a specific word.

Physical Address

The CACHE instruction uses the following portions of the PA to specify a secondary cache block and way**:**

- **PA[Size of secondary cache - 2:Blocksize of secondary cache]** is used to access the secondary cache.
- **PA[0]** is used to specify the way needed by Index operations.

Since **PA[0]** is used to indicate the way during CACHE Index operations, alignment errors are suppressed.

When accessing data in the secondary cache, **PA[Blocksize-1:3]** is also used to read or write a specific doubleword.

CP0 Not Usable

If the CP0 is not usable (if not in Kernel mode, *CU0* must be set in the *Status* register for CP0 to be usable), a Coprocessor Unusable exception is taken.

TLB Refill and TLB Invalid Exceptions on CacheOps

TLB Refill and TLB Invalid exceptions can occur on any operation. For Index operations, where the address (virtual address for the primary caches, physical address for the secondary cache) is used to index the cache but need not match the cache tag, unmapped addresses may be used to avoid TLB exceptions. The operation never causes TLB Modified exceptions.

Hit Operation Accesses

A Hit operation accesses the specified cache as a normal data reference, and performs the specified operation if the cache block contains valid data at the specified physical address (a hit).

The operation is undefined if a CacheOp hit occurs in both ways of the cache.

Watch Exception

There is no Watch exception for CacheOps.

Address Error Exception

During an Index CacheOp, bit 0 is not checked for an Address Error exception since this bit is used as the *Way* indicator bit, and may be non-zero. Bit 1 of an Index CacheOp can still generate an Address Error exception if it is not set to zero.

For all remaining CacheOps, the low-order two bits of the instruction must be set to zero, or else they will generate an Address Error exception.

A CacheOp is never checked for alignment Address Error exceptions, only for privilege-type Address Error exceptions.

Write Back

Write back from the primary data cache goes to the secondary cache. Write back from a secondary cache always goes to the System interface unit.

A secondary write back always writes the most recent data; the primary data cache must be interrogated, and any dirty inconsistent data written back to the secondary cache before the secondary block is written back to the system interface unit. The address to be written is specified by the cache tag and not the translated PA.

Invalidation

When a block is invalidated in the secondary cache, all subset blocks in the primary cache are also invalidated. The *StateMod* bits on invalidated block in the primary data cache are set to "001" (*Normal*) during any invalidation.

CE Bit

The R10000 processor does not support the *CE* bit. The functionality of the *CE* bit has been replaced by the Index Load Data and Index Store Data instructions.

CH Bit

The *CH* bit is supported in the R10000 processor. It is modified by a Hit Invalidate (S) or Hit WriteBack Invalidate (S) CACHE instruction. *CH* is set if there is a hit in the secondary cache, and cleared if there is a miss. The *CH* bit can also be modified by a MTC0 instruction.

Serial Operation of CACHE Instructions

All CACHE instruction variations are performed serially. From the aspect of the primary cache, this means CACHE instructions can impede the instruction stream. For this reason, load/store speculation is not allowed beyond a CACHE instruction until the CACHE instruction has graduated. All load/store accesses, including writebacks to the external agent, must be complete before the CACHE instruction can graduate, and any load/store following a CACHE instruction cannot be issued speculatively until the CACHE instruction graduates. Uncached operations and instruction fetches are not affected.

Instructions Not Supported

The processor does not support the following CACHE instructions:

- Create DirtyExclusive
- Hit WriteBack
- Fill (I)
- Hit Set Virtual variations

Op Field Encoding

Table 10-1 presents the Op field encoding for the CACHE instruction. Encodings not listed in this table are undefined.

Table 10-1 CACHE Instruction Op Field Encoding

Op Field	CACHE Instruction Variation	Target Cache
00000	Index Invalidate	(I)
00100	Index Load Tag	(I)
01000	Index Store Tag	(I)
10000	Hit Invalidate	(I)
10100	Cache Barrier	
11000	Index Load Data	(I)
11100	Index Store Data	(I)
00001	Index WriteBack Invalidate	(D)
00101	Index Load Tag	(D)
01001	Index Store Tag	(D)
10001	Hit Invalidate	(D)
10101	Hit WriteBack Invalidate	(D)
11001	Index Load Data	(D)
11101	Index Store Data	(D)
00011	Index WriteBack Invalidate	(S)
00111	Index Load Tag	(S)
01011	Index Store Tag	(S)
10011	Hit Invalidate	(S)
10111	Hit WriteBack Invalidate	(S)
11011	Index Load Data	(S)
11111	Index Store Data	(S)

10.2 Index Invalidate (I)

Index Invalidate (I) sets a block in the primary instruction cache to *Invalid*. **VA[13:6]** defines the address and **VA[0]** defines the way to be invalidated.

The invalidation takes place by writing the primary instruction cache state bit to 0 (*Invalid*). This also sets the instruction cache state parity bit to 0.

The **LRU** bit does not change.

Parity check is suppressed.

10.3 Index WriteBack Invalidate (D)

Index WriteBack Invalidate (D) sets a block in the primary data cache to *Invalid*. **VA[13:5]** defines the address and **VA[0]** defines the way to be invalidated.

The invalidation takes place by writing the following bits:

- primary data cache state bits are set to 00 (*Invalid*)
- the **SCWay** bit is set to 0
- the **StateMod** bits = 001 (*Normal*)
- the state parity is set to 0.

The **LRU** bit does not change.

If the **StateMod** of the block to be invalidated = 010_2 (*Inconsistent*), the block in the primary data cache must be written back to the secondary cache.

The address and way in the secondary cache to be written back to are read out of the primary data cache tag address and secondary way fields and all 32 bytes are written back.

Only the data field of the secondary cache is modified by this instruction since the processor follows state and data subset rules.

Since the *CE* bit is not defined in the R10000 processor, this instruction no longer has a CP0 *ECC* register mode.

10.4 Index WriteBack Invalidate (S)

The Index WriteBack Invalidate (S) instruction sets a block in the secondary cache to *Invalid* and writes back any dirty data to the System interface unit. This operation extends to any blocks in the primary data or instruction caches which are subsets of the secondary cache block.

The CACHE instruction physical address, **PA[Cachesize-2..Blocksize]**, defines the address and **PA[0]** defines the way to be invalidated.

The invalidation occurs in the following sequence:

1. The processor reads the **STag**, **PIdx**, and **State** bits from the secondary cache tag array. If **State** = 00 (*Invalid*) no further activity takes place. If there is a valid entry, then the STag is used to interrogate the primary instruction and data caches.
2. The processor reads each subset block from the primary instruction cache. If **ITag** = **STag** and **IState** = 1 (*Valid*) then the block is invalidated by writing the **IState** bit to 0 (*Invalid*) and the **IState** parity bit to 0.
3. Read each subset block from the primary data cache. If **DTag** = **STag** and **DState** is not equal to 00 (*Invalid*), then write the **DState** bits = 00 (*Invalid*), the **StateMod** bits = 001 (*Normal*), the **SCWay** bit = 0, and the **DState** parity bit = 0. If the original block is **DState** = 11_2 (*Dirty*) and **StateMod** = 010_2 (*Inconsistent*), also write this block back to the secondary cache using the **DTag** and the **SCWay** bit from the primary data tag array.
4. Set the state of the secondary cache block to 00 (*Invalid*). Since the secondary cache is designed so all tag bits must be written at once, the Tag, VA, and ECC bits are also written. The tag is written with the PA and **VA[13:12]** (virtual index) of the original CACHE instruction address. The ECC is generated.
5. If the secondary cache block's original **State** bits were 11_2 (*Dirty*), the block is written back to the system interface unit. If the block's **State** was *Shared* or *CleanExclusive* the system interface unit is notified with a Tag Invalidation request that the block has been deleted.

The MRU bit is set to point away from the block invalidated unless the line was already invalid.

10.5 Index Load Tag (I)

Index Load Tag (I) reads the primary instruction cache tag fields into the CP0 *TagLo* and *TagHi* registers. **VA[13:6]** defines the address and **VA[0]** defines the way of the tag to be read.

All parity errors caused by Index Load Tag (I) are ignored.

The following mapping defines the operation:

TagLo[0] = Tag parity bit
TagLo[2] = State parity bit
TagLo[3] = LRU bit
TagLo[6] = State bit
TagLo[31:8] = Tag[35:12]
TagHi[3:0] = Tag[39:36]

All other CP0 *TagLo* and *TagHi* bits are set to 0.

10.6 Index Load Tag (D)

Index Load Tag (D) reads the primary data cache tag fields into the CP0 *TagLo* and *TagHi* registers. **VA[13:5]** defines the address and **VA[0]** defines the way of the tag to be read.

All parity errors caused by Index Load Tag (D) are ignored. The following mapping defines the operation:

TagLo[0] = Tag parity bit
TagLo[1] = SCWay
TagLo[2] = State parity bit
TagLo[3] = LRU bit
TagLo[7:6] = State bits
TagLo[31:8] = Tag[35:12]
TagHi[3:0] = Tag[39:36]
TagHi[31:29] = StateMod bits

All other CP0 *TagLo* and *TagHi* bits are set to 0.

10.7 Index Load Tag (S)

Index Load Tag (S) reads the secondary cache tag fields into the CP0 *TagLo* and *TagHi* registers. The **PA[Cachesize-2..Blocksize]** defines the address and **PA[0]** defines the way to be read.

All parity and ECC errors caused by Index Load Tag (D) are ignored.

The following mapping defines the operation:

TagLo[6:0] = Tag ECC bits
TagLo[8:7] = Virtual index bits
TagLo[11:10] = State bits
TagLo[31:14] = Tag[35:18]
TagHi[3:0] = Tag[39:36]
TagHi[31] = MRU Bit

All other CP0 *TagLo* and *TagHi* register bits are set to 0.

10.8 Index Store Tag (I)

Index Store Tag (I) stores the CP0 *TagLo* and *TagHi* registers into the primary instruction cache tag array. **VA[13:6]** defines the address and **VA[0]** defines the way of the tag to be written.

The following mapping defines the operation:

Tag parity bit = **TagLo[0]**
State parity bit = **TagLo[2]**
LRU bit = **TagLo[3]**
State bit = **TagLo[6]**
Tag[35:12] = **TagLo[31:8]**
Tag[39:36] = **TagHi[3:0]**

All the Tag fields, including parity, are directly written.

Parity check is suppressed for all Index Store Tags.

10.9 Index Store Tag (D)

Index Store Tag (D) stores the CP0 *TagLo* and *TagHi* registers into the primary data cache tag array. **VA[13:5]** defines the address and **VA[0]** defines the way of the tag to be written.

The following mapping defines the operation:

Tag parity bit = **TagLo[0]**
SCWay = **TagLo[1]**
State parity bit = **TagLo[2]**
LRU bit = **TagLo[3]**
State bits = **TagLo[7:6]**
Tag[35:12] = **TagLo[31:8]**
Tag[39:36] = **TagHi[3:0]**
StateMod bits = **TagHi[31:29]**

All Tag fields, including parity, are directly written.

Parity check is suppressed for all Index Store Tags.

10.10 Index Store Tag (S)

Index Store Tag (S) stores fields from the CP0 *TagLo* and *TagHi* registers into the secondary cache tag and MRU array fields. The **PA[Cachesize-2..Blocksize]** defines the address and **PA[0]** defines the way to be read.

The following mapping defines the operation:

Tag ECC bits = **TagLo[6:0]**
Virtual index bits = **TagLo[8:7]**
State bits = **TagLo[11:10]**
Tag[35:18] = **TagLo[31:14]**
Tag[39:36] = **TagHi[3:0]**
MRU bit = **TagHi[31]**

All Tag fields, including ECC, are directly written.

Parity check is suppressed for all Index Store Tags.

10.11 Hit Invalidate (I)

Hit Invalidate (I) invalidates an entry in the instruction cache which matches the PA of the CACHE instruction. Both way tags at **VA[13:6]** are read from the instruction cache.

If the **IState** is 1 (*Valid*), and the PA of the CACHE instruction matches the Tag from the instruction cache tag array, the **IState** bit of the entry is written to 0 (*Invalid*) and the **IState** parity bit is written to 0.

The **LRU** bit does not change.

Parity error is checked.

Hit CacheOps can cause cache error exceptions if they check ECC or parity bits.

10.12 Hit Invalidate (D)

Hit Invalidate (D) invalidates an entry in the data cache which matches the PA of the CACHE instruction. Both ways tags at **VA[13:5]** are read from the data cache.

If the **DState** is not equal to 00 (*Invalid*) and the PA of the CACHE instruction matches the DTag from the data cache tag array, then the **State** bits are written to 00 (*Invalid*), the **SCWay** bit = 0, the **StateMod** bits = 001_2 (*Normal*), and the **DState** parity = 0.

The **LRU** bit is left unchanged.

Parity check is enabled.

Hit CacheOps can cause cache error exceptions if they check ECC or parity bits.

10.13 Hit Invalidate (S)

Hit Invalidate (S) invalidates all entries in the secondary, primary instruction, and primary data caches which match the PA of the CACHE instruction. The following sequence takes place:

1. The processor reads the Tags from both ways of the secondary cache at the address pointed to by the PA of the CACHE instruction. If the tag entry's STag matches the CACHE instruction PA, and the **State** of the entry is not equal to 00 (*Invalid*), then a Hit has occurred in that entry. If there is no Hit, the CACHE instruction completes.
2. The processor checks each entry in the primary caches to determine which corresponds to the CACHE instruction PA and the *PIdx* read from the secondary cache tag array. Any entry which matches is invalidated. No write back is required by Hit Invalidate (S).
3. The processor sets the tag array entry of the secondary cache block which was hit to **State** = 00 (*Invalid*), **Tag** = PA of CACHE instruction, and **PIdx** = **VA[13:12]** of CACHE instruction.
4. ECC is generated.
5. The **MRU** bit is written to point to the way opposite to that being invalidated.
6. If the processor Eliminate Request mode bit, **PrcElmReq**, is set, a processor eliminate request is sent to notify the external agent that a block in the secondary cache has been invalidated.
7. Hit Invalidate (S) sets the **CH** bit if it hits in the secondary cache.
8. Once the **CH** bit is set it stays set until cleared by a MTC0 instruction, or the next CacheOp that can change the **CH** bit.

Hit CacheOps can cause cache error exceptions if they check ECC or parity bits.

10.14 Cache Barrier

Cache Barrier does not change any cache fields. It is used when serialization of a CACHE instruction is needed without unwanted side effects. For more information, see the section titled Serial Operation of CACHE Instructions, in this chapter.

10.15 Hit Writeback Invalidate (D)

Hit Writeback Invalidate (D) invalidates an entry in the primary data cache which matches the PA of the CACHE instruction. In addition, it writes back to the secondary cache any *DirtyExclusive* or *Inconsistent* data found in the primary data cache. Both way DTags at **VA[13:5]** are read from the data cache.

If the **DState** is not equal to 00 (*Invalid*) and PA of the CACHE instruction matches the DTag, then the **DState** bits of the entry are set to 00 (*Invalid*), the **SCWay** is set to 0, the **DState** parity is set to 0, and the **StateMod** bits are set to 001_2 (*Normal*).

The **LRU** bit is left unchanged.

If the state of the block to be invalidated was found to be **StateMod** = 010_2 (*Inconsistent*), the block in the primary data cache must be written back to the secondary cache. The address and way in the secondary cache to be written back to are read out of the primary data cache Tag Address and secondary way fields, and all 32 bytes are written back.

Only the data field of the secondary cache is modified by this instruction since the processor obeys State and data subset rules.

Since the *CE* bit is not defined in the R10000 processor, this instruction no longer has an *ECC* register mode.

Hit CacheOps can cause cache error exceptions if they check ECC or parity bits.

10.16 Hit WriteBack Invalidate (S)

Hit Writeback Invalidate (S) checks for a block which matches the CACHE instruction PA in the secondary cache, invalidates it, and writes back any dirty data to the System interface unit. This operation extends to any blocks in the primary data or instruction caches which are subsets of the secondary cache block. The operation takes place in the following sequence:

1. The processor reads the **STag**, **PIdx**, and **State** bits from both ways of the secondary tag array.
2. If the PA of the CACHE instruction matches the **STag**, and the **State** does not equal 00 (*Invalid*), a hit has occurred. If there is a hit, the **STag** is used to interrogate the primary caches. If there is not a hit, the instruction ends.
3. The processor reads each subset block from the primary instruction cache. If there is a match then invalidate the block by writing the **IState** bit to 0 (*Invalid*) and the **IState** parity bit to 0.
4. Read each subset block from the primary data cache. If there is a match then write the **DState** bits = 00 (*Invalid*), the **StateMod** bits = 001 (*Normal*), the **SCWay** bit = 0, and the **DState** parity bit = 0. If the original State of any subset block is **StateMod** = 010_2 (*Inconsistent*), also write it back to the secondary cache using the DTag and the secondary way bit from the primary data tag array.
5. Write the **State** of the secondary cache block = 00 (*Invalid*). Since the secondary cache is designed so all tag bits must be written at once, the STag, PIdx, and ECC bits are also written. The STag is written with whatever the PA and **VA[13:12]** of the original CACHE instruction were. The Tag ECC is generated.
6. If the secondary block's original **State** bits were 11_2 (*Dirty*) then the block is written back to the system interface unit. If the block's State was *Shared* or *CleanExclusive* the system interface unit is simply notified that the block has been deleted with a "Tag Invalidation" request.
7. The **MRU** bit is set to point away from the block invalidated.

Hit WriteBack Invalidate (S) set the *CH* bit if it hits in the secondary cache. Once the *CH* bit is set it stays set until cleared by a MTC0 Instruction.

Hit CacheOps can cause cache error exceptions if they check ECC or parity bits.

10.17 Index Load Data (I)

Index Load Data (I) loads a single instruction from the primary instruction cache into the CP0 *TagHi*, *TagLo*, and *ECC* registers. A predecoded instruction in R10000 is 36 bits of data and one bit of parity. The address of the target instruction is **VA[13:2]** of the CACHE instruction. The way of the target instruction is **VA[0]** of the CACHE instruction. The instruction itself is loaded into CP0 *TagHi[3:0]* and *TagLo[31:0]*. The parity bit is loaded into CP0 *ECC[0]*. The tag field is not read.

Parity checking is suppressed during operation of Index Load Data (I).

10.18 Index Load Data (D)

Index Load Data (D) loads a singleword of data and the corresponding four bits of byte parity into CP0 *TagLo* and *ECC*. The address of the target singleword is **VA[13:2]** of the CACHE instruction. The way of the target singleword is **VA[0]** of the CACHE instruction. The singleword of data will be loaded into the CP0 *TagLo* register. The byte parity will be loaded into CP0 *ECC*[3:0] register. The tag field is not read.

Parity checking is suppressed during operation of Index Load Data (D).

10.19 Index Load Data (S)

Index Load Data (S) loads a doubleword of data and all 10 check bits into the CP0 *TagHi*, *TagLo*, and *ECC* registers. The address of the target doublewords comes from the PA of the CACHE instruction. The way comes from **PA[0]** of the CACHE instruction. The high word will be loaded into CP0 *TagHi* and the low word of data will be loaded into CP0 *TagLo*. The check bits will be loaded into CP0 *ECC[9:0]*. The MRU field is unmodified.

ECC correction and checking is suppressed during Index Load Data (S).

10.20 Index Store Data (I)

Index Store Data (I) stores a single instruction into the primary instruction cache. The address where this instruction will be written comes from **VA[13:2]** of the CACHE instruction. The way where the data will be written comes from **VA[0]** of the CACHE instruction. The instruction itself comes from CP0 *TagHi[3:0*] and *TagLo[31:0*]. The parity bit is also stored. This comes from CP0 *ECC[0]*. The data to be stored bypasses the predecode and is written directly into the instruction cache. The tag field is unmodified.

10.21 Index Store Data (D)

Index Store Data (D) stores a word of data and its byte parity into the data cache from the CP0 *TagLo* and *ECC* registers. The address where this word will be written is defined by **VA[13:2]** of the CACHE instruction. The way is defined by **VA[0]**. The data word comes from CP0 *TagLo*. The parity bits come from CP0 *ECC[3:0]*. The data cache tag array including the LRU bit is left unchanged.

10.22 Index Store Data (S)

Index Store Data (S) stores a quadword of data and 10 check bits into the secondary cache data array. It stores a doubleword of data from CP0 *TagHi* and *TagLo* and pads the remaining doubleword with zeroes. This allows the ECC and parity, which are based on the quadword, to be valid for the doubleword of data stored. The address of the quadword stored is defined by the PA of the CACHE instruction, and the way is defined by **PA[0]**. The data stored in the non-padded doubleword comes from CP0 *TagHi* and *TagLo*. The check bits are stored from *ECC[9:0]*. The tag array including the MRU bit is left unchanged.

JTAG Interface Operation

11

The JTAG interface is implemented according to the standard IEEE 1149.1 test access port protocol specifications.

The JTAG interface accesses the JTAG controller and instruction register as well as a boundary scan register.

11.1 Test Access Port (TAP)

The test access port (TAP) consists of four interface signals. These signals are used to control the serial loading and unloading of instructions and test data, as well as to execute tests.

The TAP consists of the following signals:

JTDI:	Serial data input	(Input signal)
JTDO:	Serial data output	(Output signal)
JTMS:	Mode select	(Input signal)
JTCK:	Clock	(Input signal)

The timing and the relationship of the TAP signals follows the IEEE 1149.1 standard protocol.

TAP Controller (Input)

The R10000 processor implements the 16-state TAP controller specified by the IEEE 1149.1 standard in the following manner:

- The **JTMS** signal operates the state machine synchronized by the **JTCK** signal.
- The TAP controller is reset by keeping the **JTMS** signal asserted through five consecutive edges of **JTCK.** This reset condition sets the reset state of the controller. The TAP controller is also reset by asserting **SysReset***. This pin must not be asserted while using the boundary scan register.

11.2 Instruction Register

The JTAG instruction register is four bits wide, permitting a total of 16 instructions to control the selection of the bypass register, the boundary scan register, and other data registers.

The encoding of the instruction register is given in Table 11-1:

Table 11-1 JTAG Instruction Register Encoding

MSB...LSB	Selected Data Register
0000 0001	Boundary Scan Register Sample - Preload
0010 to 1110	Data Register (not used)
1111	Bypass Register

The 0001 value is provided to represent sample-preload, but also selects the boundary scan register.

During a reset of the TAP controller, the value 1111 is loaded into the parallel output of the instruction register, thus selecting the bypass register as the default.

During the Shift-IR state of the TAP controller, data is shifted serially into the instruction register from **JTDI**, and the LSB of the instruction register is shifted out onto **JTDO**.

During the Update-IR state, the current state of the instruction register is shifted to its parallel output for decoding.

11.3 Bypass Register

The bypass register is 1 bit wide.

When the bypass register is selected and the TAP controller is in the Shift-DR state, data on **JTDI** is shifted into the bypass register and the output of the bypass register is shifted out onto **JTDO**.

11.4 Boundary Scan Register

The bypass register is 1 bit wide.

The boundary scan data register is selected by loading 0000 into the instruction register. The Shift-DR, Update-DR, and Capture-DR states of the TAP controller are used to operate the boundary scan register according to the IEEE 1149.1 standard specifications.

The boundary scan register provides serial access to each of the processor interface pins, as shown in Figure 11-1. Hence, the boundary scan register can be used to load and observe specific logic values on the processor pins.

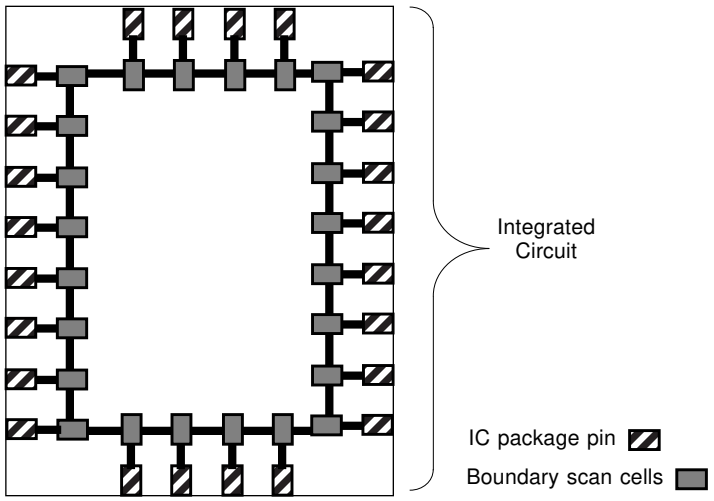

Figure 11-1 JTAG Boundary Scan Cells

The main application of the boundary scan register is board-level interconnect testing.

The use of the boundary scan register for applying data to and capturing data from the internal microprocessor circuitry is not supported.

The boundary scan register list for rev 1.2 of the fab is given in Table 11-2. The **TriState** signal will be eliminated from the BSR in rev 2.0 of the fab, and beyond.

An additional bit is provided in the boundary scan register to control the direction of bidirectional pins. As it is loaded through JTDI, this bit is the first bit in the boundary scan chain. The logic value of this bit is latched during the Update-DR state, and sets the direction of all bidirectional pins as follows:

Value	Direction
0	Input
1	Output

The value is set to 0 during reset, setting all bidirectional pins to *input* prior to any boundary scan operations.

Table 11-2 Boundary Scan Register Pinlist, rev 1.2

Signal	Signal	Signal	Signal	Signal	Signal
1. SCDataChk[1]	2. SCData[63]	3. SCData[62]	4. SCData[61]	5. SCData[60]	6. SCData[59]
7. SCData[58]	8. SCData[57]	9. SCData[56]	10. SCData[55]	11. SCData[54]	12. SCData[53]
13. SCData[52]	14. SCData[51]	15. SCData[50]	16. SCData[49]	17. SCData[48]	18. SCData[47]
19. SCData[46]	20. SCData[45]	21. SCData[44]	22. SCData[43]	23. SCData[42]	24. SCData[41]
25. SCData[40]	26. SCData[39]	27. SCData[38]	28. SCData[37]	29. SCData[36]	30. SCData[35]
31. SCData[34]	32. SCData[33]	33. SCData[32]	34. SysAD[0]	35. SysAD[1]	36. SysAD[2]
37. SysAD[3]	38. SysAD[4]	39. SysAD[5]	40. SysAD[6]	41. SysAD[7]	42. SysAD[8]
43. SysAD[9]	44. SysAD[10]	45. SysAD[11]	46. SysAD[12]	47. SysAD[13]	48. SysAD[14]
49. SysAD[15]	50. SCData[0]	51. SCData[1]	52. SCData[2]	53. SCData[3]	54. SCData[4]
55. SCData[5]	56. SCData[6]	57. SCData[7]	58. SCData[8]	59. SCData[9]	60. SCData[10]
61. SCData[11]	62. SCData[12]	63. SCData[13]	64. SCData[14]	65. SCData[15]	66. SCData[16]
67. SCData[17]	68. SCData[18]	69. SCData[19]	70. SCData[20]	71. SCData[21]	72. SCData[22]
73. SCData[23]	74. SCData[24]	75. SCData[25]	76. SCData[26]	77. SCData[27]	78. SCData[28]
79. SCData[29]	80. SCData[30]	81. SCData[31]	82. SCDataChk[0]	83. SCAAddr[18]	84. SCAAddr[17]
85. SCAAddr[16]	86. SCAAddr[15]	87. SCAAddr[14]	88. SCAAddr[13]	89. SCAAddr[12]	90. SCAAddr[11]
91. SCAAddr[10]	92. SCAAddr[9]	93. SCDataChk[2]	94. SCDataChk[4]	95. SCData[64]	96. SCData[65]
97. SCData[66]	98. SCData[67]	99. SCData[68]	100. SCData[69]	101. SCData[70]	102. SCData[71]
103. SCDataChk[9]	104. SysCyc*	105. SysAD[32]	106. SysAD[33]	107. SysAD[34]	108. SysAD[35]
109. SysAD[36]	110. SysAD[37]	111. SysAD[38]	112. SysAD[39]	113. SysAD[40]	114. SysAD[41]
115. SysAD[42]	116. SysAD[43]	117. SysAD[44]	118. SysAD[45]	119. SysAD[46]	120. SysAD[47]
121. SCData[72]	122. SCData[73]	123. SCData[74]	124. SCData[75]	125. SCData[76]	126. SCData[77]
127. SCData[78]	128. SCData[79]	129. SCAAddr[0]	130. SCAAddr[1]	131. SCAAddr[2]	132. SCAAddr[3]
133. SCAAddr[4]	134. SCAAddr[5]	135. SCAAddr[6]	136. SCAAddr[7]	137. SCAAddr[8]	138. SCADWay
139. SCADCS*	140. SCADOE*	141. SCADWr*	142. SCData[80]	143. SCData[81]	144. SCData[82]
145. SCData[83]	146. SCData[84]	147. SCData[85]	148. SCData[86]	149. SCData[87]	150. SCData[88]
151. SCData[89]	152. SCData[90]	153. SCData[91]	154. SCData[92]	155. SCData[93]	156. SCData[94]
157. SCData[95]	158. SCDataChk[6]	159. SCDataChk[8]	160. Spare1	161. SCTCS*	162. SCTOE*
163. SCTWr*	164. SCTag[25]	165. SCTag[24]	166. SCTag[23]	167. SCTag[22]	168. SCTag[21]
169. SCTag[20]	170. SCTag[19]	171. SCTag[18]	172. SCTag[17]	173. SCTag[16]	174. SCTag[15]
175. SCTag[14]	176. SCTag[13]	177. SCTag[12]	178. SCTag[11]	179. SCTag[10]	180. SCTag[9]
181. SCTag[8]	182. SCTag[7]	183. SCTag[6]	184. SCTag[5]	185. SCTag[4]	186. SCTag[3]
187. SCTag[2]	188. SCTag[1]	189. SCTag[0]	190. SCTagLSBAddr	191. TriState†	192. SCTWay
193. SCTagChk[6]	194. SCTagChk[5]	195. SCTagChk[4]	196. SCTagChk[3]	197. SCTagChk[2]	198. SCTagChk[1]
199. SCTagChk[0]	200. SysCmd[0]	201. SysCmd[1]	202. SysCmd[2]	203. SysCmd[3]	204. SysCmd[4]
205. SysCmd[5]	206. SysCmd[6]	207. SysCmd[7]	208. SysCmd[8]	209. SysCmd[9]	210. SysCmd[10]
211. SysCmd[11]	212. SysCmdPar	213. SysVal*	214. SysReq*	215. SysRel*	216. SysGnt*
217. SysReset*	218. SysRespVal*	219. SysRespPar	220. SysResp[4]	221. SysResp[3]	222. SysResp[2]
223. SysResp[1]	224. SysResp[0]	225. SysGblPerf*	226. SysRdRdy*	227. SysWrRdy*	228. SysStateVal*
229. SysStatePar	230. SysState[2]	231. SysState[1]	232. SysState[0]	233. SysCorErr*	234. SysUncErr*
235. SysNMI*	236. SCDataChk[7]	237. SCDataChk[5]	238. SCData[127]	239. SCData[126]	240. SCData[125]
241. SCData[124]	242. SCData[123]	243. SCData[122]	244. SCData[121]	245. SCData[120]	246. SCData[119]
247. SCData[118]	248. SCData[117]	249. SCData[116]	250. SCData[115]	251. SCData[114]	252. SCData[113]
253. SCData[112]	254. SCBDWr*	255. SCBDOE*	256. SCBDCS*	257. SCBDWay	258. SCBAddr[8]
259. SCBAddr[7]	260. SCBAddr[6]	261. SCBAddr[5]	262. SCBAddr[4]	263. SCBAddr[3]	264. SCBAddr[2]
265. SCBAddr[1]	266. SCBAddr[0]	267. SCData[111]	268. SCData[110]	269. SCData[109]	270. SCData[108]

†. Will be eliminated after rev. 1.2.

Table 11-2 (cont.) Boundary Scan Register Pinlist, rev 1.2

Signal	Signal	Signal	Signal	Signal	Signal
271. SCData[107]	272. SCData[106]	273. SCData[105]	274. SCData[104]	275. SysAD[63]	276. SysAD[62]
277. SysAD[61]	278. SysAD[60]	279. SysAD[59]	280. SysAD[58]	281. SysAD[57]	282. SysAD[56]
283. SysAD[55]	284. SysAD[54]	285. SysAD[53]	286. SysAD[52]	287. SysAD[51]	288. SysAD[50]
289. SysAD[49]	290. SysAD[48]	291. SysADChk[7]	292. SysADChk[6]	293. SysADChk[5]	294. SysADChk[4]
295. SysADChk[3]	296. SysADChk[2]	297. SysADChk[1]	298. SysADChk[0]	299. SysAD[31]	300. SysAD[30]
301. SysAD[29]	302. SysAD[28]	303. SysAD[27]	304. SysAD[26]	305. SysAD[25]	306. SysAD[24]
307. SysAD[23]	308. SysAD[22]	309. SysAD[21]	310. SysAD[20]	311. SysAD[19]	312. SysAD[18]
313. SysAD[17]	314. SysAD[16]	315. SCData[103]	316. SCData[102]	317. SCData[101]	318. SCData[100]
319. SCData[99]	320. SCData[98]	321. SCData[97]	322. SCData[96]	323. SCDataChk[3]	324. SCBAddr[9]
325. SCBAddr[10]	326. SCBAddr[11]	327. SCBAddr[12]	328. SCBAddr[13]	329. SCBAddr[14]	330. SCBAddr[15]
331. SCBAddr[16]	332. SCBAddr[17]	333. SCBAddr[18]			

Electrical Specifications

12

This chapter contains the following electrical and signal information about the R10000 processor:

- DC electrical specification
- AC electrical specification
- signal integrity issues

12.1 DC Electrical Specification

This section describes the following DC electrical characteristics of the R10000 processor:

- DC power supply levels
- **DCOk** and power supply sequencing
- maximum operating conditions
- input signal level sensing
- mode definitions
- **Vref[SC,Sys]**
- unused inputs
- DC input/output specifications

DC Power Supply Levels

The processor core is powered by a +3.3V (+/- 5%) supply. The processor output drivers are powered from a separate supply, dependent on the output logic family used in the application system:

- For JEDEC-compatible HSTL operation, the nominal value for **VccQSC** and **VccQSys** are in the 1.5V (+/- 100 millivolt) range.
- For CMOS/TTL compatible systems, **VccQSC** and **VccQSys** can be externally tied to the same **Vcc** as the core power supply.

NOTE: The I/O pins of the R10000 processor may not be driven higher than 4.0V by any device in the system until the **Vcc** and **VccQ** inputs are stable.

DCOk and Power Supply Sequencing

The following guidelines are designed to protect the processor from damage or latch-up:

- With respect to the **Vcc** (3.3V) (supply to the core), **VccQ[SC,Sys]** (either 1.5V or 3.3V) must not be driven more than a diode threshold voltage.
- **Vref** should not go higher than **VccQ[SC,Sys]**. Generally, **Vref** is derived from **VccQ** through a resistor divider, and therefore cannot rise above **VccQ**.
- The power to termination resistors must not arrive before **Vcc** and **VccQ[SC,Sys]** arrive at the processor.
- None of the supplies can float or be driven negative.

One method of protecting the processor from excessive input voltage is to sequence the power supplies for the entire system, ensuring that the power to the processor is stable before any components drive signals to the processor. Another method to tristate all external drivers to the processor with the **DCOk** pin, until the processor has stabilized.

NOTE: The input voltage required for the **DCOk** is 3.3V in either the CMOS/TTL or the HSTL configuration. Both **DCOk** pins must be tied together externally.

Maximum Operating Conditions

Table 12-1 shows the maximum conditions under which the processor operates.

Table 12-1 Maximum Operating Conditions

Parameter	Symbol	Value
Core Supply Voltage	Vcc	3.6 volts
Output Supply Voltage	VccQ (HTSL) VccQ (CMOS/TTL)	1.6 volts 3.6 volts
Junction Temperature	T_j	20° to 85° C
Applied Input Voltage:	Vin	-0.5 to Vcc+0.5 volts
Maximum Power	P_{R10000}	30 watts
PClk Frequency	f	200 MHz

Input Signal Level Sensing

The processor input signals are all received by CMOS receivers that are compatible with either HSTL or CMOS/TTL logic levels. The I/O levels are defined by **VrefSC** and **VrefSys**, according to the appropriate logic family (HSTL or CMOS/TTL).

Mode Definitions

The mode bit, **ODrainSys**, is provided to select the characteristics of the pad ring.

When asserted, this mode bit tristates the PMOS pullup devices to select system interface output drivers. This mode is included to allow for multiprocessor systems to use a GTL-like open drain configuration with external load/termination resistors providing logic high levels.

Vref[SC,Sys]

The **Vref[SC,Sys]** pins must be connected to a stable reference voltage source. This reference point is used in the input sense amp current mirror to provide the switch point for the logic levels.

Inside the processor, the **Vref[SC,Sys]** signals have a large capacitance, and a low-pass filter at each receiver. The **DCOk** pins must not be asserted until there has been sufficient time for **Vref[SC,Sys]** to stabilize at each of the receivers inside the processor.

A typical **Vref[SC,Sys]** generator is two resistors which provide the **Vref[SC,Sys]** level associated with the chosen logic family, and a 10µF tantalum capacitor connected to the processor's **Vref[SC,Sys]** pin to provide stability.

Unused Inputs

Several input pins are unused during normal system operation, and should be tied to **V**cc through resistors:

- **JTDI**
- **JTCK**
- **JTMS**

Several input pins are unused during normal system operation, and should be tied to **Vss** through 100 ohm resistors:

- **TCA**, **TCB**
- **PLLDis**
- **Spare1**, **Spare3**

Several input pins are unused during normal system operation, and should be tied to **Vss**:

- **PLLSpare1**, **PLLSpare2**, **PLLSpare3**, **PLLSpare4**
- **SelDVCO**

DC Input/Output Specifications

All processor output drivers are CMOS push-pull, and the signals swing between **VccQ** and **Vss**. In open drain mode, the gates of the PMOS pullup devices are disabled. Input-only pins include a disabled output buffer for implicit ESD protection.

Tables 12-2 and 12-3 describe the DC characteristics of the I/O signals for the HSTL and CMOS/TLL configurations.

NOTE: As the JEDEC Standard 8-x evolves, the HSTL specifications will also change, and the processor will remain compliant with these standards.

Table 12-2 DC Characteristics for HSTL Configuration

Symbol	Parameter	Minimum	Maximum	Units	Conditions
VOH	Output high voltage	VccQ /2 + 0.3V	N/A	V	N/A
VOL	Output low voltage		VccQ /2 - 0.3V	V	N/A
VIH	Input high voltage	Vref + 100mV	Vcc + 300mV	V	N/A
VIL	Input low voltage	-300mV	Vref - 100mV	V	N/A
ILeak	I/O leakage current	-TBD	TBD	µA	N/A

Table 12-3 DC Characteristics for CMOS/TTL Configuration

Symbol	Parameter	Minimum	Maximum	Units	Conditions
VOH	Output high voltage	2.4	N/A	V	Vcc = VccQ = min
VOL	Output low voltage	N/A	0.4	V	Vcc = VccQ = min
VIH	Input high voltage	2.0	N/A	V	N/A
VIL	Input low voltage	N/A	0.8	V	N/A
ILeak	I/O leakage current	-TBD	TBD	µA	N/A

12.2 AC Electrical Specification

This section describes the following AC electrical characteristics of the R10000 processor:

- maximum operating conditions
- test specification
- secondary cache and system interface timing
- enable/output delay, setup, hold time
- asynchronous inputs

Maximum Operating Conditions

The R10000 chip clamps signals that overshoot the DC limits established for input logic levels. These limits are published as part of the fabrication process characterization.

The R10000 chip provides silicon diode clamps on all signal pins.

Test Specification

HSTL test conditions are based on the JEDEC Standard conditions.

Secondary Cache and System Interface Timing

Timing measurements are referenced from the mid-swing point of the input signal to the crossing point of the **SysClk** and **SysClk*** input clocks. All input signals maintain a 1 V/ns edge rate in the 20% to 80% range of the input signal swing.

Enable/Output Delay, Setup, Hold Time

Table 12-4 lists the delay, setup, and hold times for the HTSL version of the processor.

Table 12-4 AC Characteristics for HSTL Configuration

HSTL	Minimum	Maximum
Output delay	0.5 ns	1.5 ns
Setup	1.0 ns	
Hold	1.0 ns	

Table 12-5 lists the delay, setup, and hold times for the CMOS/TTL version of the processor.

Table 12-5 AC Characteristics for CMOS/TTL Configuration

LVCMOS	Minimum	Maximum
Output delay	0.5 ns	2.0 ns
Setup	1.0 ns	
Hold	1.0 ns	

Asynchronous Inputs

The **SysReset*** input can be asserted asynchronously to **SysClk**, but must be negated synchronously with **SysClk**, adhering to the AC electrical specifications listed above.

12.3 Signal Integrity Issues

In this section, the following signal integrity considerations are described for a R10000-based system:

- Power Supply Regulation
- Decoupling Capacitance
- Reference Voltage
- Maximum Input Voltage Levels
- Output I-V Curves
- Switching and Slew Rate Characteristics

Reference Voltage

Most input pins on the processor use a current-mirror sense amp with **Vref[SC,Sys]** supplied to the negative input to provide a single rail input receiver. The following input pins are exceptions to this rule:

- **SysClk** and **SysClk***
- **DCOk**

All other inputs require a stable **Vref[SC,Sys]** supply for proper operation.

The **Vref[SC,Sys]** source can be a simple voltage divider; the actual impedance of this source is not critical, since the **Vref[SC,Sys]** signals are sampled through a low-pass filter on the processor.

Power Supply Regulation

The system must provide connections to all of the **Vcc**, **VccQ[SC,Sys]**, and **Vss** pins on the processor package. The power supply voltages must be held to 5% tolerance at the processor pin connection.

Maximum Input Voltage Levels

Maximum excursion of the input signal due to ringing may reach Vcc+0.5V or Vss-0.5V for periods of less than 10% of the total driven waveform period. The R10000 processor includes overshoot clamps by silicon diode protection which limit the overshoot to approximately 500 mV beyond each supply rail.

Decoupling Capacitance

In order to regulate the transient current requirements of a R10000-based system, it is suggested that explicit decoupling capacitors be used. The R10000 package allows for the following capacitors:

- eight **Vcc-Vss**
- six **VccQSC-Vss**
- two **VccQSys-Vss**

The package also provides six connections for the PLL power supplies and loop capacitors.

VccPa (**VccPd**) is connected to **VssPa** (**VssPd**) through three decoupling capacitors, as shown in Figures 12-1 and 12-2. The 0.1µF and 1 nF low-inductance capacitors are placed in parallel with the 10 µF capacitor, as close to the R10000 package as possible.†

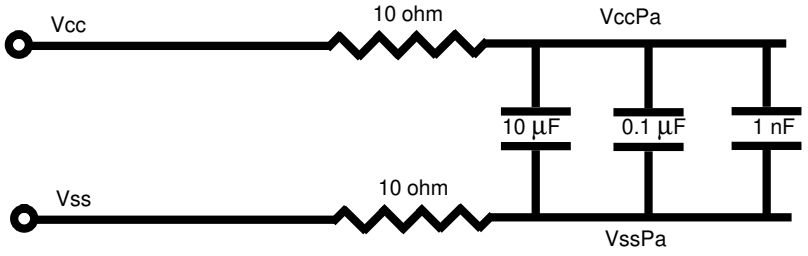

Figure 12-1 Decoupling VccPa and VssPa

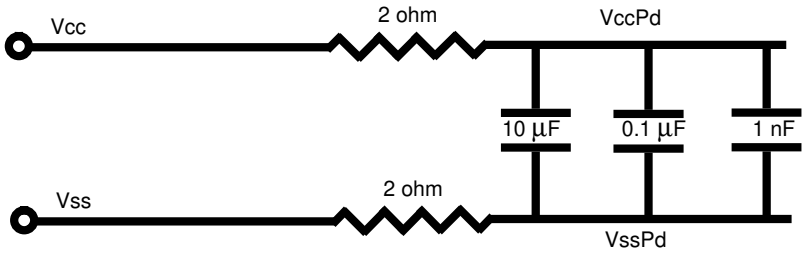

Figure 12-2 Decoupling VccPd and VssPd

Output I-V Curves

Not yet available

Switching and Slew Rate Characteristics

Not yet available

† Decoupling between VccPa and VssPa is far more important than decoupling between VccPd and VssPd, if both are not possible.

Packaging

13

The R10000 microprocessor is presently supplied in one standard package configuration:

- a single-chip 599 ceramic LGA (Land Grid Array)

MIPS Licensees are encouraged to develop package solutions with MIPS Semiconductor Partners to meet specific requirements.

13.1 R10000 Single-Chip Package, 599CLGA

The standard single-chip R10000 package is a **599CLGA** (ceramic land grid array), as shown in Figure 13-1.

The 599CLGA package minimizes output switching noise by reducing the inductance of the power and ground paths leading into the package. Much of the decrease in power/ground inductance is accomplished by shortening the wire bonds running from the die pads to the package inner leads. The 599CLGA is designed with its cavity-side down, and the die is connected directly to a thermal slug.

Mechanical Characteristics

The 599CLGA has lands on a straight 1.27mm (.050inch) grid. It is a cavity-down, multi-layer ceramic package with an integral copper-tungsten slug, and is designed for use with a socket. Preliminary information suggests that the 599CLGA can withstand a force of 100 kilograms applied to the CuW slug, without damage, and a PWB assembly should insure that this force is not exceeded. Drawings for a reference LGA-PWB assembly are included in this chapter.

Electrical Characteristics

The 599CLGA uses multilayer construction, incorporating stripline configuration for signals. Multiple planes distribute power and ground throughout the package and provide built-in distributed bypass/coupling capacitance between the primary power supplies: **Vcc**, **VccQSC**, **VccQSys**, and **Vss**.

Pads are present on the package body for attaching chip-capacitors to provide additional bypass capacitance between the primary power supplies and the PLL power supply (**VccPa** and **VssPa**), and to provide an additional PLL loop filter capacitor (**PLLRC**). Chip-capacitors on the R10000 are assembled by the chip manufacturer.

Detailed electrical package characteristics will be provided by the MIPS Semiconductor Partners as they become available. The data in Table 13-1 is provided as an estimate of the package parasitics. These estimates include the effects of bondwires, package traces and vias, but not the sockets.

Table 13-1 R10000 599CLGA Electrical Characteristics

Parameter	Description	Minimum	Typical	Maximum
L_{sig}	Effective signal inductance	4.0nH		8.4nH
M_{sig}	Signal-to-signal mutual inductance			1.3nH
C_{sig}	Signal loading capacitance	3.0pF		5.6pF
C_m	Signal-to-signal mutual capacitance			0.5pF
R_{sig}	Signal resistance	400mΩ		1300mΩ
Z_0	Characteristic impedance		40Ω	
T_{pd}	Propagation delay			200ps

The copper-tungsten slug (provided for thermal performance) is hard-connected to **Vss** to minimize EMI radiation from the package.

Thermal Characteristics

The 599CLGA incorporates a copper-tungsten slug to provide an efficient thermal path from the processor to the heatsink.

The thermal analysis listed in Table 13-2 gives a preliminary indication of heatsink requirements for the 599CLGA.

Table 13-2 R10000 599CLGA Thermal Characteristics - Preliminary

Parameter	Description	Value
T_j	Maximum junction temperature	85° C
T_a	Maximum ambient temperature	40° C
P_{R10000}	Maximum power dissipation	30 watts
T_{ja}	Minimum temperature differential	45° C
Θ_{ja}	Minimum total thermal resistance	1.5° C/W
Θ_{jc}	Estimated junction to case thermal resistance	0.3° C/W
Θ_{ca}	Required case to ambient thermal resistance	1.2° C/W

System designers must take care, especially in desktop applications, to ensure sufficient airflow and heat-dissipation surface area to meet the required case-to-ambient thermal resistance, Θ_{ca}=1.2C/W.

The thermal interface between the package and heatsink is very important. Typically, grease or compliant material is inserted between the package and heatsink to increase the contact area between their surfaces.

Assembly Drawings and Pinout List

The following pages contain a pinout list (Table 13-3), and drawings of an example R10000 LGA-PWB assembly, including details of the PWB, heatsink, and bolster plate. Actual hardware specifications are dependent on the user.

An assembly drawing of the 599LGA is also shown in Figure 13-2. Note that hardware specifications given in this drawing will require modifications to accommodate the actual dimensions of the socket, PWB, heatsink, bolster, etc.

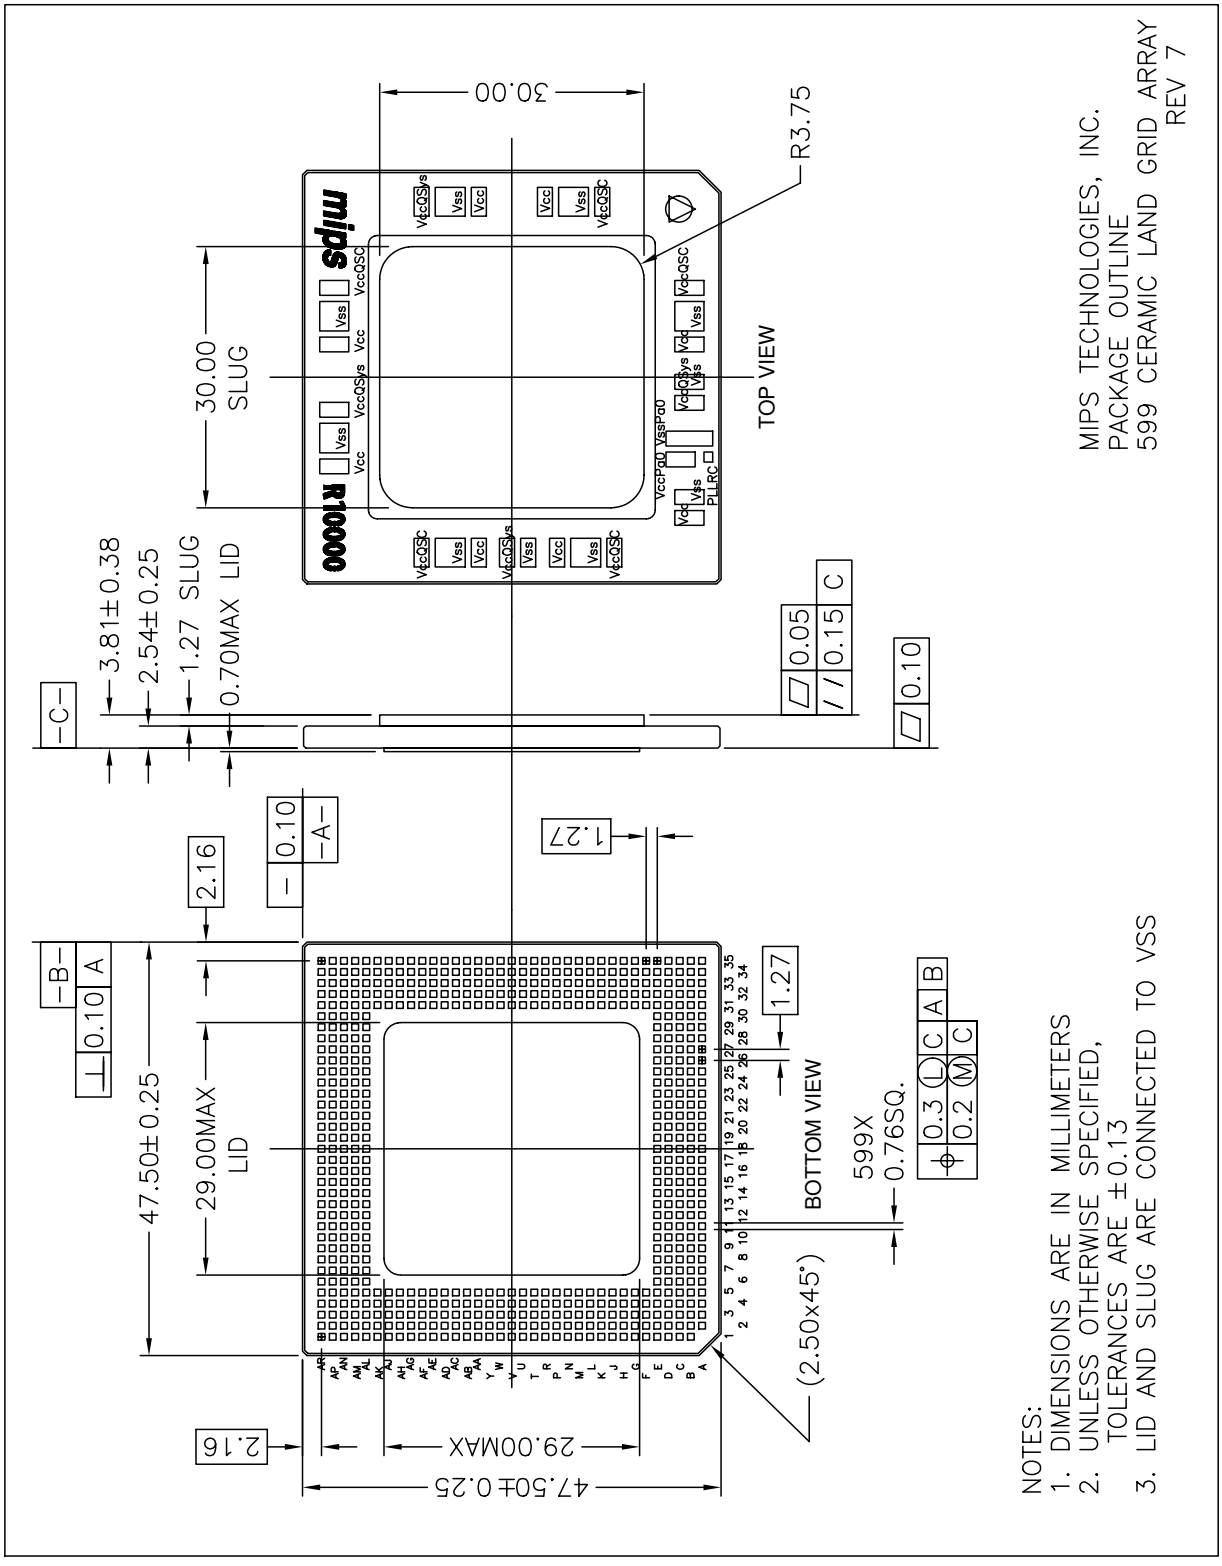

Figure 13-1 R10000 599CLGA Package Outline

599CLGA Pinout

Table 13-3 599CLGA Pinout

Signal	Location	Signal	Location	Signal	Location
DCOk	AF.....2	DCOk	B.....22	JTCK	W..... 33
JTDI	W 35	JTDO	Y 31	JTMS	AA 34
PLLDis	E.....24	PLLRC	A 25	PLLSpare1	C 21
PLLSpare2	A 21	PLLSpare3	D 21	PLLSpare4	B 21
SCAAddr<0>	E.....13	SCAAddr<1>	A 11	SCAAddr<2>	D..... 12
SCAAddr<3>	C.....11	SCAAddr<4>	E.....12	SCAAddr<5>	B 10
SCAAddr<6>	D 11	SCAAddr<7>	C 10	SCAAddr<8>	A..... 9
SCAAddr<9>	B.....30	SCAAddr<10>	E.....29	SCAAddr<11>	A..... 31
SCAAddr<12>	D 30	SCAAddr<13>	C 31	SCAAddr<14>	E 30
SCAAddr<15>	B.....32	SCAAddr<16>	D.....31	SCAAddr<17>	B 33
SCAAddr<18>	C.....32	SCADCS*	B.....9	SCADOE*	D..... 9
SCADWay	E.....10	SCADWr*	A 8	SCBAddr<0>	AL 13
SCBAddr<1>	AP.....12	SCBAddr<2>	AM 12	SCBAddr<3>	AR..... 11
SCBAddr<4>	AL.....12	SCBAddr<5>	AN.....11	SCBAddr<6>	AM..... 11
SCBAddr<7>	AP.....10	SCBAddr<8>	AL 11	SCBAddr<9>	AL 29
SCBAddr<10>	AP.....30	SCBAddr<11>	AM 30	SCBAddr<12>	AR..... 31
SCBAddr<13>	AL.....30	SCBAddr<14>	AN.....31	SCBAddr<15>	AM..... 31
SCBAddr<16>	AP.....32	SCBAddr<17>	AP.....33	SCBAddr<18>	AN 32
SCBDCS*	AN.....10	SCBDOE*	AL 10	SCBDWay	AR..... 9
SCBDWr*	AP.....9	SCClk<0>	B.....13	SCClk<1>	A..... 26
SCClk<2>	AA 31	SCClk<3>	AM 15	SCClk<4>	W..... 1
SCClk<5>	E.....1	SCClk<0>*	E.....15	SCClk<1>*	B 26
SCClk<2>*	AB.....33	SCClk<3>*	AR 14	SCClk<4>*	W..... 4
SCClk<5>*	F 4	SCData<0>	R 31	SCData<1>	N 34
SCData<2>	P 33	SCData<3>	M 35	SCData<4>	P 32
SCData<5>	M.....34	SCData<6>	N.....33	SCData<7>	L 35
SCData<8>	N 31	SCData<9>	L.....33	SCData<10>	M..... 32
SCData<11>	K 34	SCData<12>	M 31	SCData<13>	J 35
SCData<14>	L.....32	SCData<15>	J.....34	SCData<16>	K..... 33
SCData<17>	H 35	SCData<18>	K 31	SCData<19>	G..... 34
SCData<20>	J 32	SCData<21>	G 33	SCData<22>	J 31
SCData<23>	F 35	SCData<24>	H.....32	SCData<25>	F 34
SCData<26>	G 31	SCData<27>	E.....35	SCData<28>	F 32
SCData<29>	D 34	SCData<30>	F.....31	SCData<31>	E 32
SCData<32>	AA 32	SCData<33>	AB 35	SCData<34>	AC..... 34
SCData<35>	AB.....32	SCData<36>	AD.....35	SCData<37>	AC..... 33
SCData<38>	AD 34	SCData<39>	AC 31	SCData<40>	AE 35
SCData<41>	AD 32	SCData<42>	AE 33	SCData<43>	AD 31
SCData<44>	AF.....34	SCData<45>	AE 32	SCData<46>	AG..... 35
SCData<47>	AF.....33	SCData<48>	AG.....34	SCData<49>	AF 31
SCData<50>	AH.....35	SCData<51>	AG.....32	SCData<52>	AJ 34
SCData<53>	AG 31	SCData<54>	AJ.....33	SCData<55>	AH 32
SCData<56>	AK 35	SCData<57>	AJ.....31	SCData<58>	AK..... 34

Signal	Location	Signal	Location	Signal	Location
SCData<59>	AK..... 32	SCData<60>	AL.....35	SCData<61>	AK 31
SCData<62>	AM..... 34	SCData<63>	AM.....33	SCData<64>	D 28
SCData<65>	B 29	SCData<66>	E 27	SCData<67>	C 28
SCData<68>	D..... 27	SCData<69>	E 26	SCData<70>	A 28
SCData<71>	C 26	SCData<72>	B 15	SCData<73>	D 15
SCData<74>	A..... 14	SCData<75>	C.....14	SCData<76>	A 12
SCData<77>	D..... 14	SCData<78>	B 12	SCData<79>	C 13
SCData<80>	E 9	SCData<81>	C.....8	SCData<82>	D 8
SCData<83>	B 7	SCData<84>	C.....7	SCData<85>	A 6
SCData<86>	E 7	SCData<87>	B 6	SCData<88>	D 6
SCData<89>	A..... 5	SCData<90>	E 6	SCData<91>	C 5
SCData<92>	D..... 5	SCData<93>	B 4	SCData<94>	C 4
SCData<95>	B 3	SCData<96>	AN 29	SCData<97>	AP.....29
SCData<98>	AM..... 28	SCData<99>	AN 28	SCData<100>	AL.....27
SCData<101>	AR..... 28	SCData<102>	AM.....27	SCData<103>	AP.....27
SCData<104>	AL 16	SCData<105>	AP 15	SCData<106>	AL.....15
SCData<107>	AP 13	SCData<108>	AN 14	SCData<109>	AN.....13
SCData<110>	AM..... 14	SCData<111>	AR.....12	SCData<112>	AM 9
SCData<113>	AR..... 8	SCData<114>	AL.....9	SCData<115>	AN.....8
SCData<116>	AM..... 8	SCData<117>	AP 7	SCData<118>	AN.....7
SCData<119>	AR..... 6	SCData<120>	AL.....7	SCData<121>	AP.....6
SCData<122>	AM..... 6	SCData<123>	AR.....5	SCData<124>	AL.....6
SCData<125>	AN 5	SCData<126>	AM.....5	SCData<127>	AP.....4
SCDataChk<0>	D..... 33	SCDataChk<1>	AL.....32	SCDataChk<2>	C 29
SCDataChk<3>	AR..... 30	SCDataChk<4>	A 30	SCDataChk<5>	AP.....3
SCDataChk<6>	E 4	SCDataChk<7>	AN 4	SCDataChk<8>	D 3
SCDataChk<9>	B 27	SCTCS*	D 2	SCTag<0>	R 1
SCTag<1>	R 4	SCTag<2>	P 1	SCTag<3>	R 5
SCTag<4>	P 3	SCTag<5>	N 2	SCTag<6>	P.....4
SCTag<7>	M..... 1	SCTag<8>	N 3	SCTag<9>	M 2
SCTag<10>	N 5	SCTag<11>	M.....4	SCTag<12>	L.....1
SCTag<13>	M..... 5	SCTag<14>	K.....2	SCTag<15>	L.....4
SCTag<16>	J 1	SCTag<17>	K.....3	SCTag<18>	J.....2
SCTag<19>	K..... 5	SCTag<20>	H 1	SCTag<21>	J.....4
SCTag<22>	G..... 2	SCTag<23>	J 5	SCTag<24>	G 3
SCTag<25>	H 4	SCTagChk<0>	V.....4	SCTagChk<1>	W 3
SCTagChk<2>	V..... 2	SCTagChk<3>	V.....5	SCTagChk<4>	V 1
SCTagChk<5>	U..... 3	SCTagChk<6>	U 1	SCTOE*	G 5
SCTWay	T 3	SCTWr*	F 2	SCTagLSBAddr	T.....5
SelDVCO	E 21	Spare1	F 5	Spare3	U 4
SysAD<0>	Y 34	SysAD<1>	W 32	SysAD<2>	V 35
SysAD<3>	V..... 31	SysAD<4>	V.....34	SysAD<5>	U 33
SysAD<6>	V..... 32	SysAD<7>	U 32	SysAD<8>	U 35
SysAD<9>	T 33	SysAD<10>	T 34	SysAD<11>	T.....31
SysAD<12>	R 35	SysAD<13>	R.....32	SysAD<14>	R 34
SysAD<15>	P 35	SysAD<16>	AL.....26	SysAD<17>	AR 27
SysAD<18>	AN 26	SysAD<19>	AP 26	SysAD<20>	AL.....25

Signal	Location	Signal	Location	Signal	Location
SysAD<21>	AN.....25	SysAD<22>	AM..... 25	SysAD<23>	AR..... 25
SysAD<24>	AL.....24	SysAD<25>	AP..... 24	SysAD<26>	AM..... 24
SysAD<27>	AR.....24	SysAD<28>	AL 23	SysAD<29>	AN 23
SysAD<30>	AM 22	SysAD<31>	AP..... 23	SysAD<32>	C..... 20
SysAD<33>	B.....20	SysAD<34>	D..... 19	SysAD<35>	A..... 19
SysAD<36>	C.....19	SysAD<37>	A..... 18	SysAD<38>	D..... 18
SysAD<39>	E.....18	SysAD<40>	B..... 18	SysAD<41>	C..... 17
SysAD<42>	A 17	SysAD<43>	D..... 17	SysAD<44>	B 16
SysAD<45>	C.....16	SysAD<46>	A..... 15	SysAD<47>	E 16
SysAD<48>	AN.....20	SysAD<49>	AR 19	SysAD<50>	AL 19
SysAD<51>	AN.....19	SysAD<52>	AM 19	SysAD<53>	AP 18
SysAD<54>	AM 18	SysAD<55>	AR 18	SysAD<56>	AL 18
SysAD<57>	AR.....17	SysAD<58>	AM 17	SysAD<59>	AN 17
SysAD<60>	AL.....17	SysAD<61>	AP..... 16	SysAD<62>	AN 16
SysAD<63>	AR.....15	SysADChk<0>	AN..... 22	SysADChk<1>	AR..... 22
SysADChk<2>	AL.....21	SysADChk<3>	AP..... 21	SysADChk<4>	AM..... 21
SysADChk<5>	AR.....21	SysADChk<6>	AL 20	SysADChk<7>	AP 20
SysClk	A 22	SysClk*	A..... 23	SysClkRet*	C..... 23
SysClkRet	B.....23	SysCmd<0>	Y 2	SysCmd<1>	Y 3
SysCmd<2>	AA 1	SysCmd<3>	Y 5	SysCmd<4>	AA 2
SysCmd<5>	AA 4	SysCmd<6>	AB 1	SysCmd<7>	AA 5
SysCmd<8>	AB.....3	SysCmd<9>	AC 2	SysCmd<10>	AB 4
SysCmd<11>	AD 1	SysCmdPar	AC 3	SysCorErr*	AK..... 4
SysCyc*	E.....20	SysGblPerf*	AG..... 4	SysGnt*	AD 4
SysNMI*	AK 5	SysRdRdy*	AJ..... 3	SysRel*	AE 1
SysReq*	AC 5	SysReset*	AD..... 5	SysResp<0>	AJ 2
SysResp<1>	AF.....5	SysResp<2>	AH..... 1	SysResp<3>	AF 3
SysResp<4>	AG 2	SysRespPar	AE 4	SysRespVal*	AG..... 1
SysState<0>	AL.....1	SysState<1>	AJ..... 5	SysState<2>	AK..... 2
SysStatePar	AH.....4	SysStateVal	AK 1	SysUncErr*	AM..... 2
SysVal*	AD 2	SysWrRdy*	AG..... 5	TCA	AM..... 3
TCB	AL.....4	TriState	T..... 2	VccPa	B 25
VccPa	C.....25	VccPd	E..... 22	VrefByp	C..... 22
VssPa	A 24	VssPa	B..... 24	VssPd	D..... 22
Vcc	A 2	Vcc	A..... 34	Vcc	AB 2
Vcc	AB.....34	Vcc	AE 3	Vcc	AF 32
Vcc	AF.....4	Vcc	AH..... 2	Vcc	AH 34
Vcc	AL.....3	Vcc	AL 31	Vcc	AL 33
Vcc	AL.....5	Vcc	AM 10	Vcc	AM..... 16
Vcc	AM 20	Vcc	AM 26	Vcc	AN 18
Vcc	AN.....2	Vcc	AN..... 34	Vcc	AP 1
Vcc	AP.....14	Vcc	AP..... 22	Vcc	AP 28
Vcc	AP.....35	Vcc	AP..... 8	Vcc	AR..... 2
Vcc	AR.....34	Vcc	B..... 1	Vcc	B 14
Vcc	B.....28	Vcc	B..... 35	Vcc	B 8
Vcc	C.....18	Vcc	C 2	Vcc	C..... 34
Vcc	D 10	Vcc	D..... 16	Vcc	D..... 20

Signal	Location	Signal	Location	Signal	Location
Vcc	D..... 26	Vcc	E.....3	Vcc	E.....31
Vcc	E 33	Vcc	E.....5	Vcc	F.....1
Vcc	H 2	Vcc	H 34	Vcc	K 32
Vcc	K..... 4	Vcc	L 3	Vcc	P.....2
Vcc	P 34	Vcc	T 32	Vcc	T.....4
Vcc	V 3	Vcc	V.....33	Vcc	Y 32
Vcc	Y 4	VccQSC	A 10	VccQSC	A 32
VccQSC	A..... 4	VccQSC	AB.....31	VccQSC	AD.....33
VccQSC	AF 35	VccQSC	AH.....31	VccQSC	AH.....33
VccQSC	AK..... 33	VccQSC	AL.....14	VccQSC	AL.....28
VccQSC	AL 8	VccQSC	AM.....35	VccQSC	AN.....12
VccQSC	AN 3	VccQSC	AN 30	VccQSC	AN.....33
VccQSC	AN 6	VccQSC	AR.....10	VccQSC	AR 32
VccQSC	AR..... 4	VccQSC	C.....12	VccQSC	C 3
VccQSC	C 30	VccQSC	C.....33	VccQSC	C 6
VccQSC	D..... 1	VccQSC	D 35	VccQSC	E.....14
VccQSC	E 28	VccQSC	E.....8	VccQSC	F.....3
VccQSC	F..... 33	VccQSC	H 3	VccQSC	H.....31
VccQSC	H 33	VccQSC	H 5	VccQSC	K 1
VccQSC	K..... 35	VccQSC	M.....3	VccQSC	M 33
VccQSC	P 31	VccQSC	P 5	VccQSC	R 2
VccQSC	T 1	VccQSys	A 16	VccQSys	A 20
VccQSys	AB 5	VccQSys	AD 3	VccQSys	AF.....1
VccQSys	AH 3	VccQSys	AH.....5	VccQSys	AK 3
VccQSys	AL 22	VccQSys	AM.....1	VccQSys	AN.....24
VccQSys	AR..... 16	VccQSys	AR.....20	VccQSys	AR 26
VccQSys	T 35	VccQSys	Y.....1	VccQSys	Y 35
VrefSC	AA..... 35	VrefSys	Y.....33	Vss	A 13
Vss	A..... 27	Vss	A 29	Vss	A 3
Vss	A..... 33	Vss	A 35	Vss	A 7
Vss	AA..... 3	Vss	AA 33	Vss	AC 1
Vss	AC..... 32	Vss	AC 35	Vss	AC 4
Vss	AE 2	Vss	AE.....31	Vss	AE.....34
Vss	AE 5	Vss	AG 3	Vss	AG.....33
Vss	AJ 1	Vss	AJ 32	Vss	AJ.....35
Vss	AJ 4	Vss	AL.....2	Vss	AL.....34
Vss	AM..... 13	Vss	AM.....23	Vss	AM 29
Vss	AM..... 32	Vss	AM.....4	Vss	AM 7
Vss	AN 1	Vss	AN.....15	Vss	AN.....21
Vss	AN 27	Vss	AN.....35	Vss	AN.....9
Vss	AP 11	Vss	AP.....17	Vss	AP.....19
Vss	AP 2	Vss	AP.....25	Vss	AP.....31
Vss	AP 34	Vss	AP.....5	Vss	AR 1
Vss	AR..... 13	Vss	AR.....23	Vss	AR 29
Vss	AR..... 3	Vss	AR.....33	Vss	AR 35
Vss	AR..... 7	Vss	B 11	Vss	B.....17
Vss	B 19	Vss	B 2	Vss	B.....31

Signal	Location	Signal	Location	Signal	Location
Vss	B.....34	Vss	B.....5	Vss	C..... 1
Vss	C.....15	Vss	C 24	Vss	C..... 27
Vss	C.....35	Vss	C 9	Vss	D..... 13
Vss	D 23	Vss	D..... 24	Vss	D..... 25
Vss	D 29	Vss	D..... 32	Vss	D..... 4
Vss	D 7	Vss	E..... 11	Vss	E 17
Vss	E.....19	Vss	E.....2	Vss	E 23
Vss	E.....25	Vss	E.....34	Vss	G..... 1
Vss	G 32	Vss	G..... 35	Vss	G..... 4
Vss	J.....3	Vss	J.....33	Vss	L 2
Vss	L.....31	Vss	L.....34	Vss	L 5
Vss	N.....1	Vss	N.....32	Vss	N 35
Vss	N.....4	Vss	R 3	Vss	R..... 33
Vss	U 2	Vss	U..... 31	Vss	U..... 34
Vss	U 5	Vss	W 2	Vss	W..... 31
Vss	W 34	Vss	W 5		

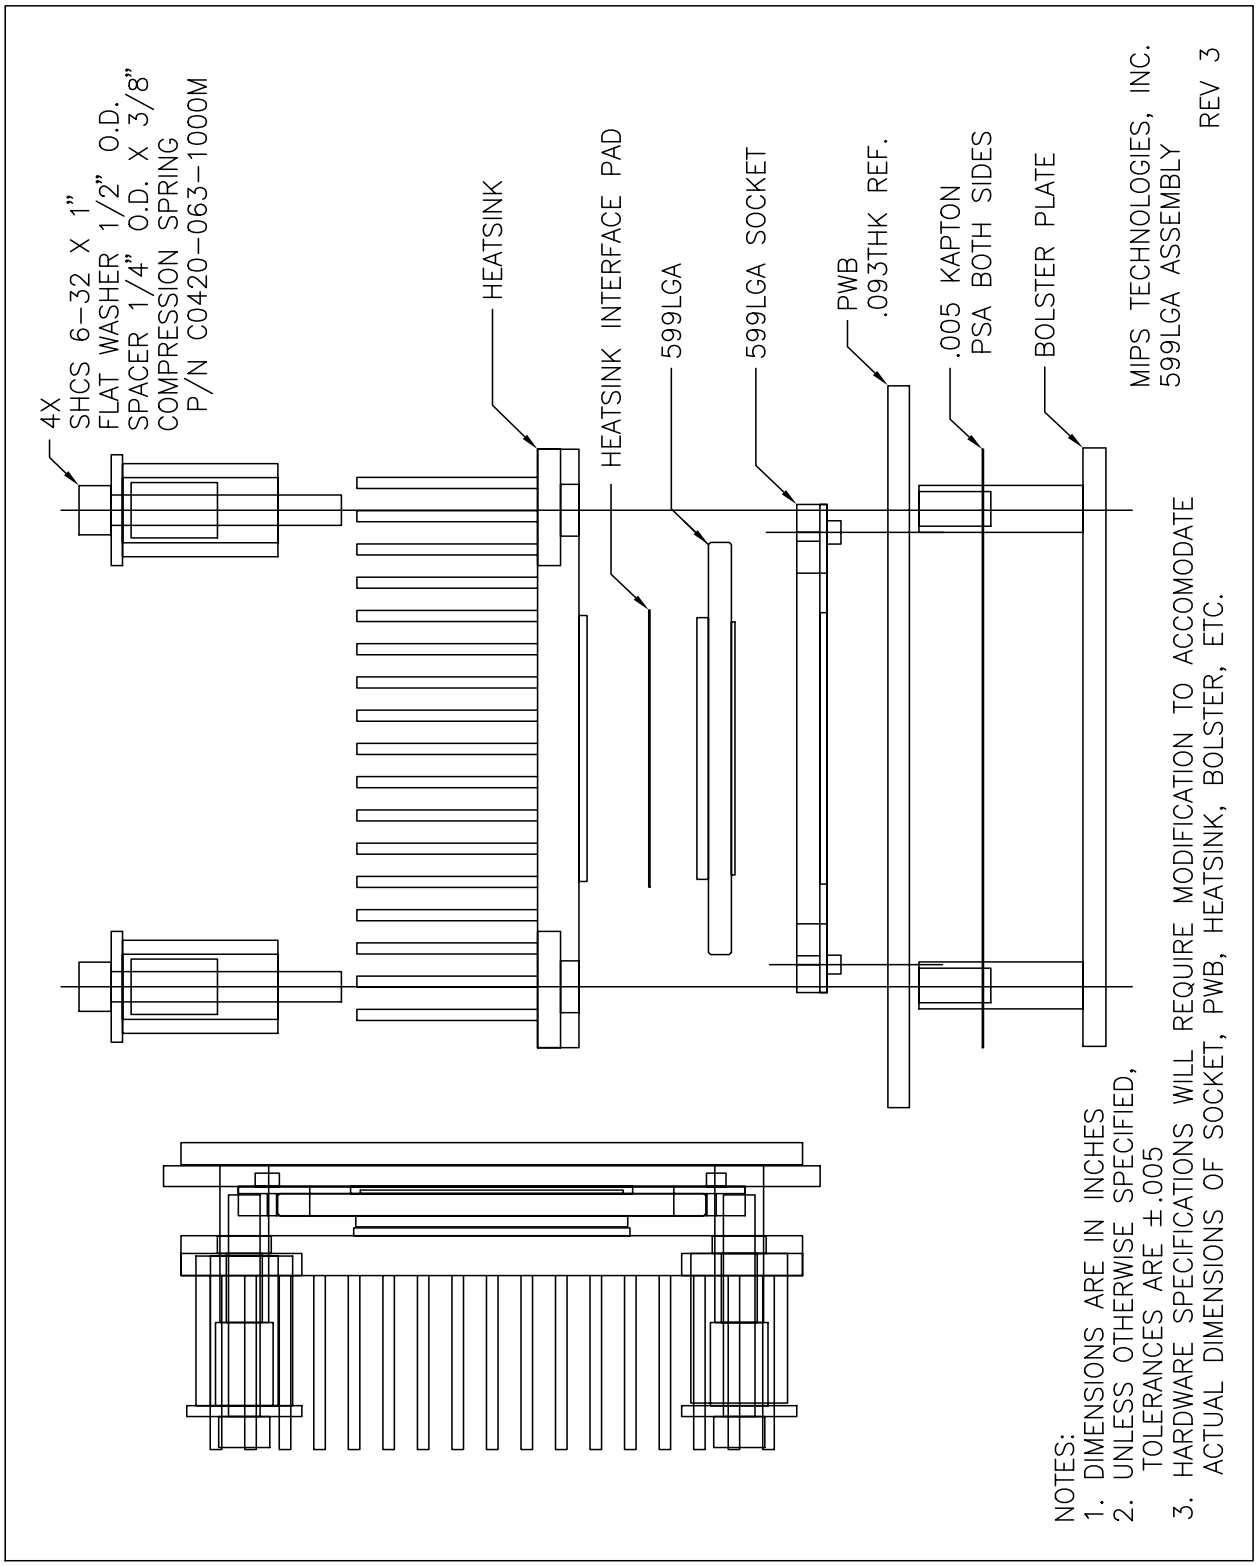

Figure 13-2 599LGA Assembly Drawing

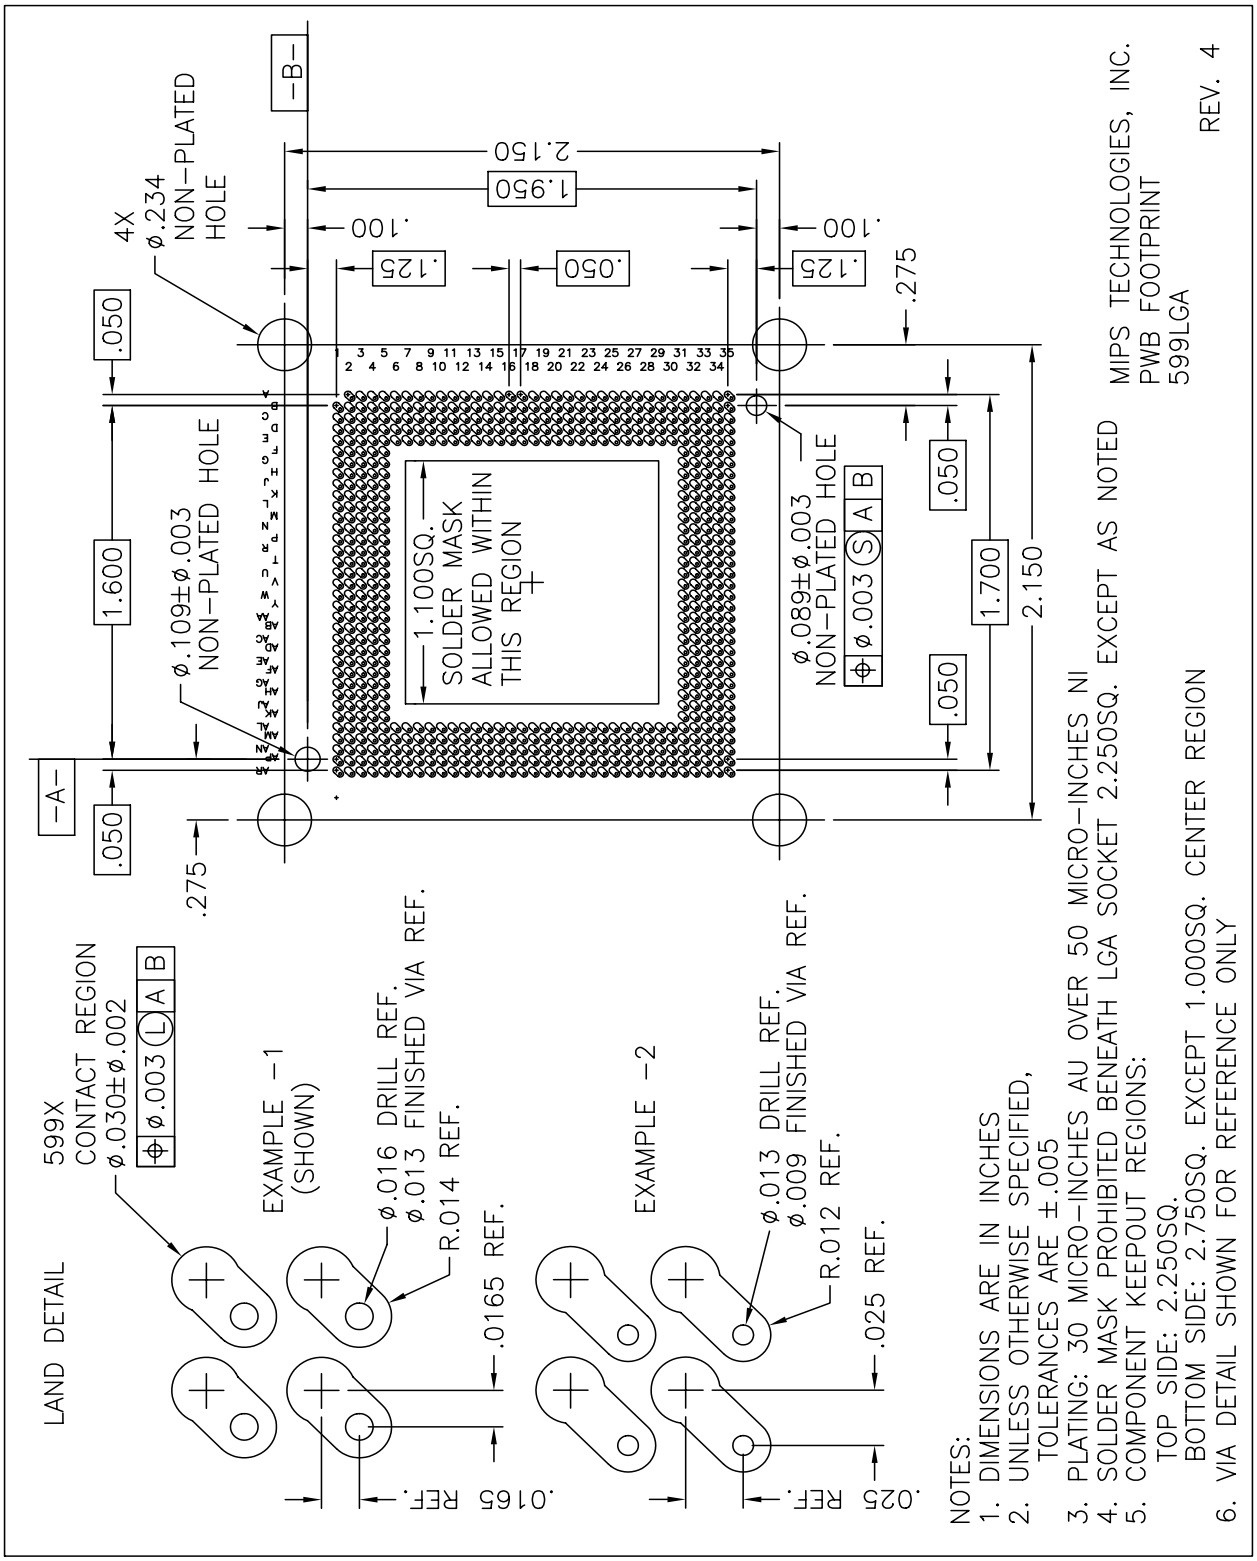

Figure 13-3 599LGA PWB Footprint

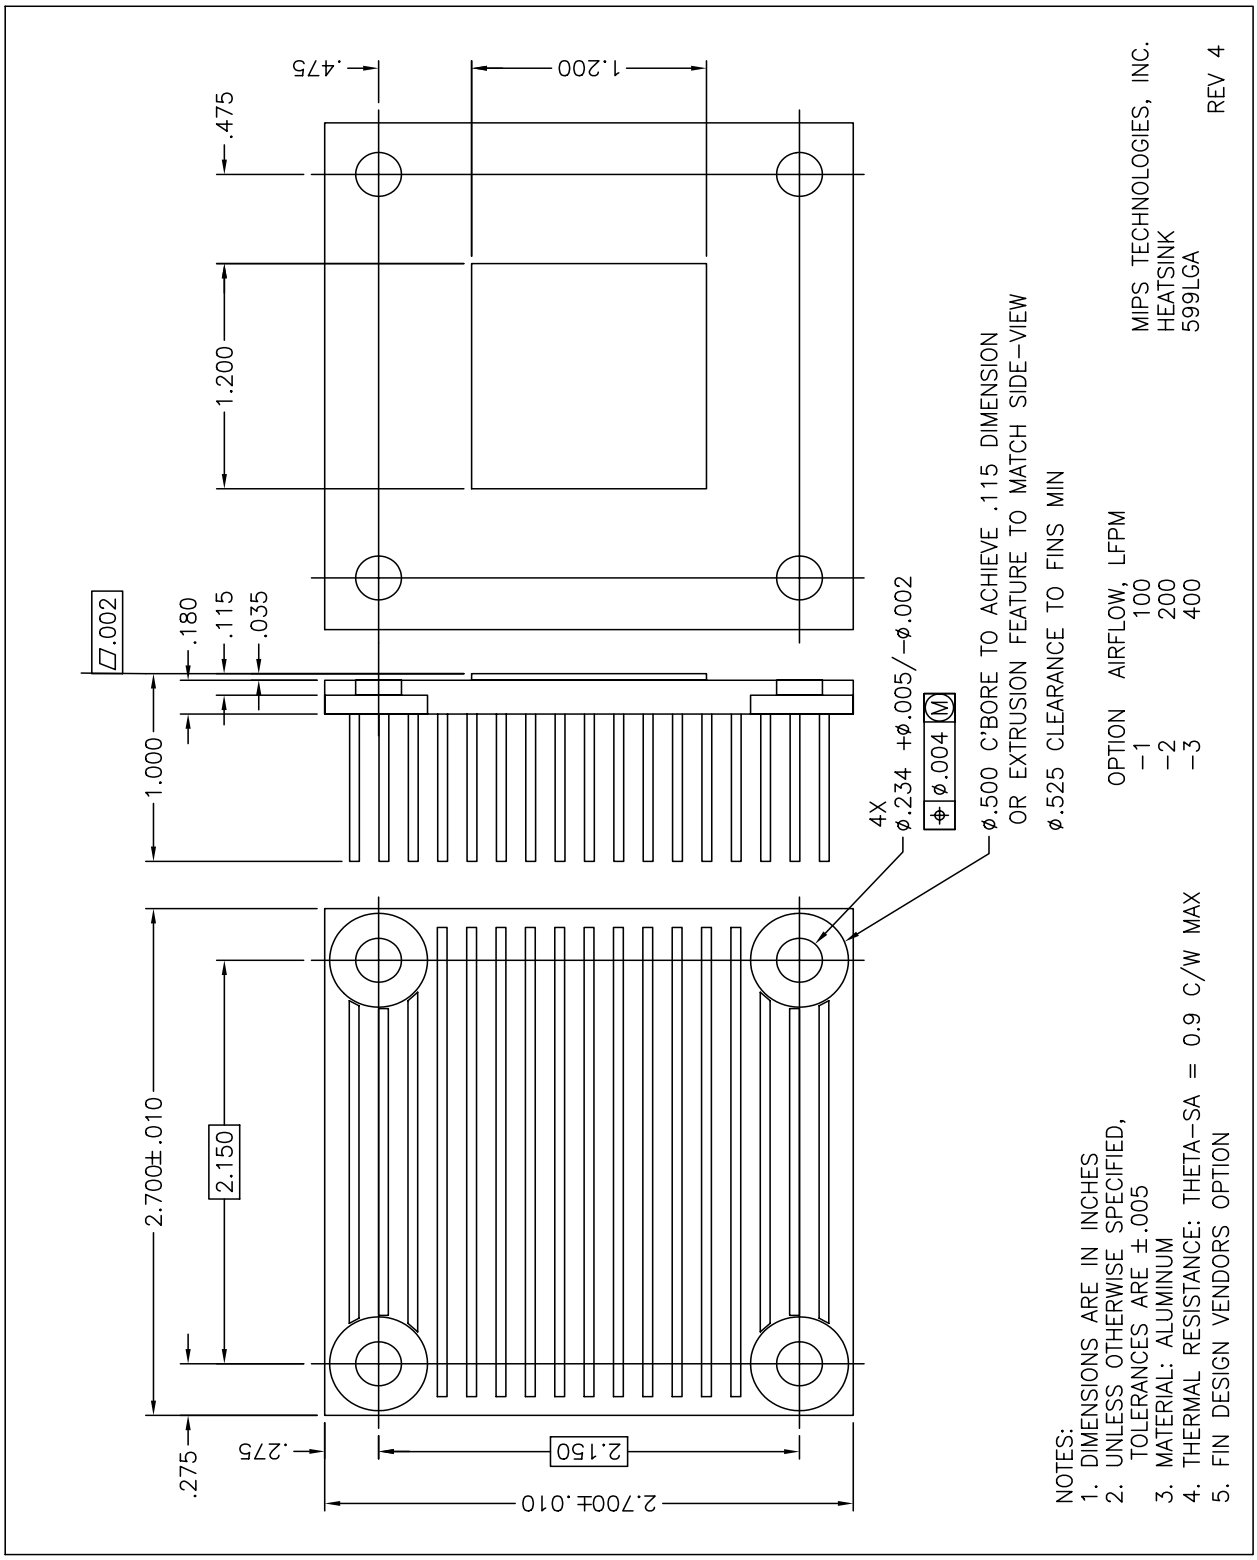

Figure 13-4 599LGA Heatsink

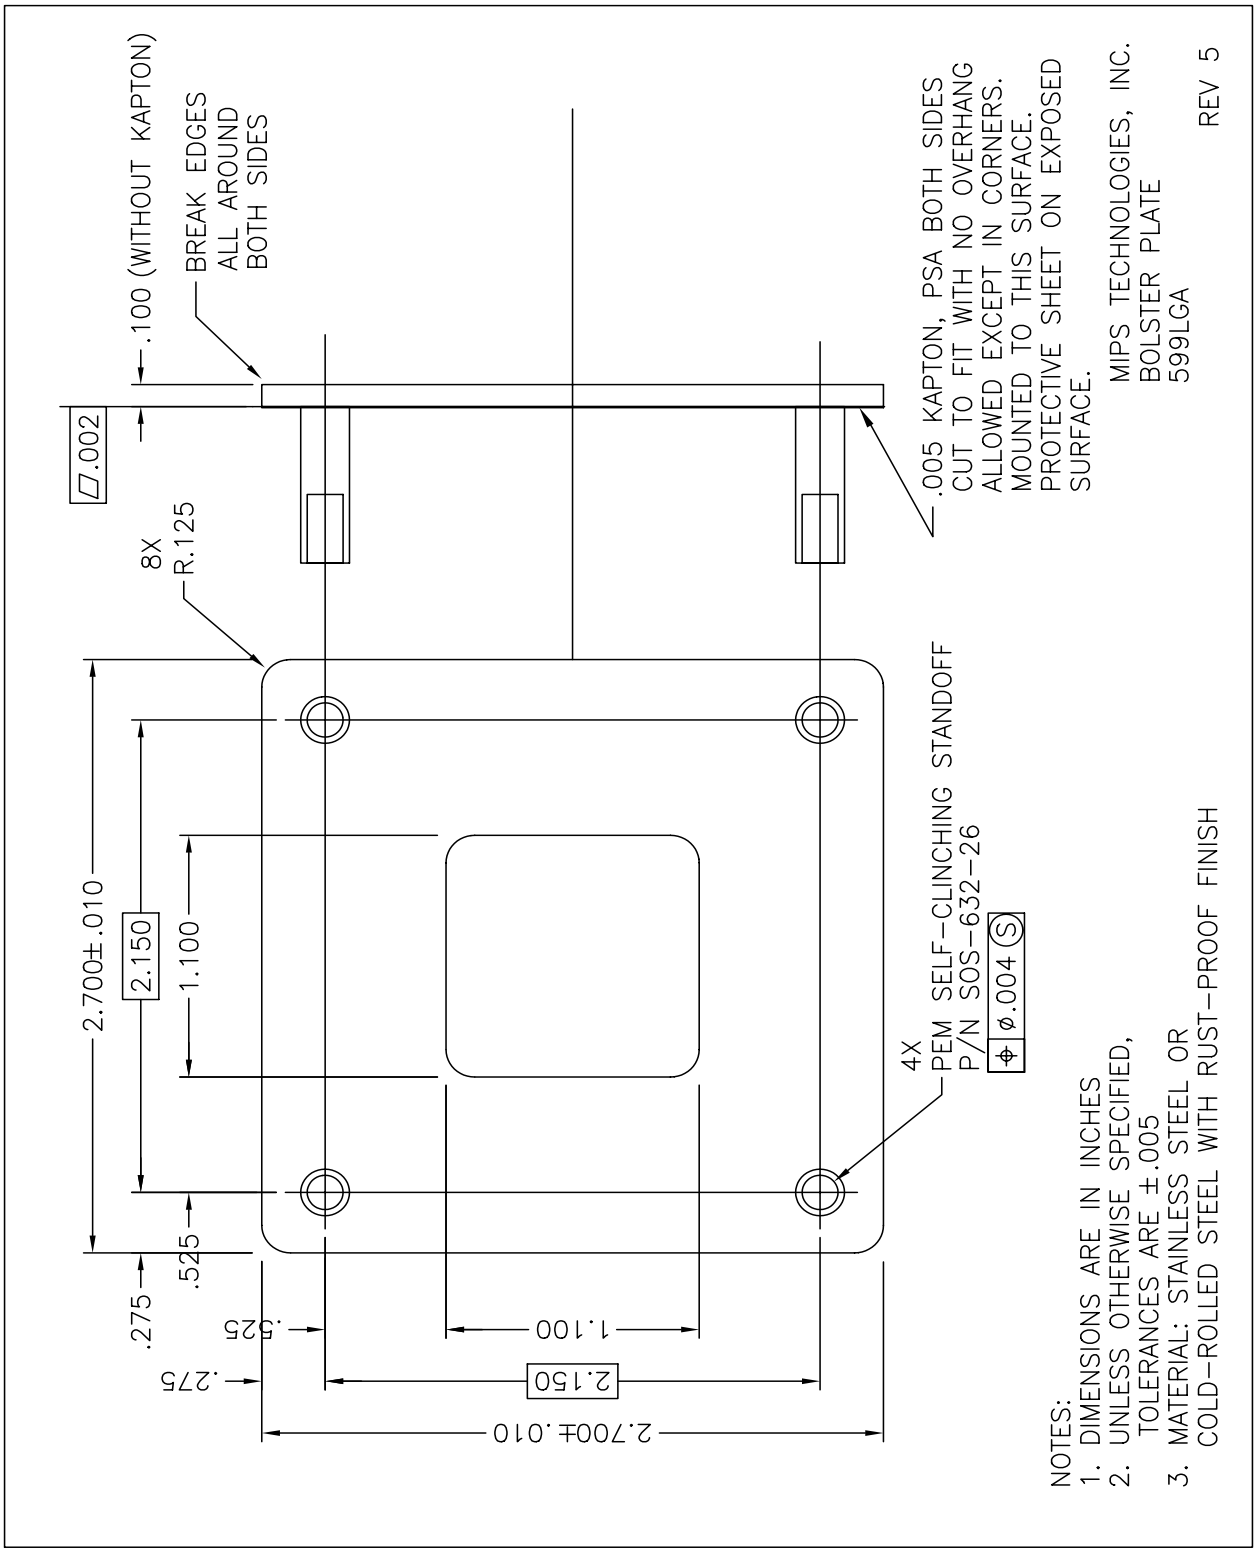

Figure 13-5 599LGA Bolster Plate

Coprocessor 0

14

This chapter describes the Coprocessor 0 operation, concentrating on the CP0 register definitions and the R10000 processor implementation of CP0 instructions.

The Coprocessor 0 (CP0) registers control the processor state and report its status. These registers can be read using MFC0 instructions and written using MTC0 instructions. CP0 registers are listed in Table 14-1.

Table 14-1 Coprocessor 0 Registers

Register No.	Register Name	Description
0	Index	Programmable register to select TLB entry for reading or writing
1	Random	Pseudo-random counter for TLB replacement
2	EntryLo0	Low half of TLB entry for even VPN (Physical page number)
3	EntryLo1	Low half of TLB entry for odd VPN (Physical page number)
4	Context	Pointer to kernel virtual PTE table in 32-bit addressing mode
5	Page Mask	Mask that sets the TLB page size
6	Wired	Number of wired TLB entries (lowest TLB entries not used for random replacement)
7	*Undefined*	*Undefined*
8	BadVAddr	Bad virtual address
9	Count	Timer count
10	EntryHi	High half of TLB entry (Virtual page number and ASID)
11	Compare	Timer compare
12	Status	Processor Status Register
13	Cause	Cause of the last exception taken
14	EPC	Exception Program Counter
15	PRId	Processor Revision Identifier
16	Config	Configuration Register (secondary cache size, etc.)
17	LLAddr	Load Linked memory address
18	WatchLo	Memory reference trap address (low bits Adr[39:32])
19	WatchHi	Memory reference trap address (high bits Adr[31:3])
20	XContext	Pointer to kernel virtual PTE table in 64-bit addressing mode
21	FrameMask	Mask the physical addresses of entries which are written into the TLB
22	BrDiag	Branch Diagnostic register
23	*Undefined*	*Undefined*
24	*Undefined*	*Undefined*
25	PC	Performance Counters
26	ECC	Secondary cache ECC and primary cache parity
27	CacheErr	Cache Error and Status register
28	TagLo	Cache Tag register - low bits
29	TagHi	Cache Tag register - high bits
30	ErrorEPC	Error Exception Program Counter

Coprocessor 0 instructions are enabled if the processor is in Kernel mode, or if bit 28 (*CU0*) is set in the *Status* register. Otherwise, executing one of these instructions generates a Coprocessor 0 Unusable exception.

14.1 Index Register (0)

The *Index* register is a 32-bit, read/write register containing six bits to index an entry in the TLB. The high-order bit of the register shows the success or failure of a TLB Probe (TLBP) instruction.

The *Index* register also specifies the TLB entry affected by TLB Read (TLBR) or TLB Write Index (TLBWI) instructions.

Figure 14-1 shows the format of the *Index* register; Table 14-2 describes the *Index* register fields.

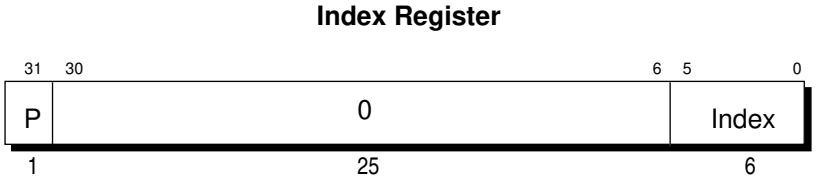

Figure 14-1 Index Register

Table 14-2 Index Register Field Descriptions

Field	Description
P	Probe failure. Set to 1 when the previous TLBProbe (TLBP) instruction was unsuccessful.
Index	Index to the TLB entry affected by the TLBRead and TLBWrite instructions
0	Reserved. Must be written as zeroes, and returns zeroes when read.

14.2 Random Register (1)

The *Random* register is a read-only register of which six bits index an entry in the TLB. This register decrements when any instruction graduates at that particular cycle, and its values range between an upper and a lower bound, as follows:

- The lower bound is set by the number of TLB entries reserved for exclusive use by the operating system (the contents of the *Wired* register).
- The upper bound is set by the total number of TLB entries minus 1 (64 – 1 maximum).

The *Random* register specifies the entry in the TLB that is affected by the TLB Write Random instruction. The register does not need to be read for this purpose; however, the register is readable to verify proper operation of the processor.

To simplify testing, the *Random* register is set to the value of the upper bound upon system reset. This register is also set to the upper bound when the *Wired* register is written.

Figure 14-2 shows the format of the *Random* register; Table 14-3 describes the *Random* register fields.

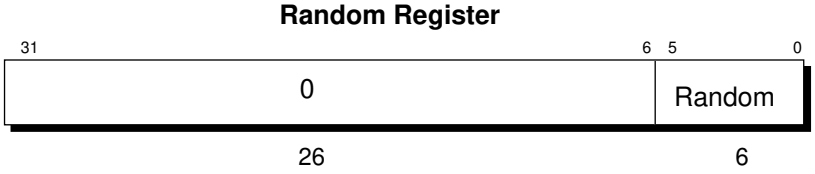

Figure 14-2 Random Register

Table 14-3 Random Register Field Descriptions

Field	Description
Random	TLB Random index
0	Reserved. Must be written as zeroes, and returns zeroes when read.

14.3 EntryLo0 (2), and EntryLo1 (3) Registers

The *EntryLo* register consists of two registers with identical formats:

- *EntryLo0* is used for even virtual pages.
- *EntryLo1* is used for odd virtual pages.

The *EntryLo0* and *EntryLo1* registers are read/write registers. They hold the physical page frame number (PFN) of the TLB entry for even and odd pages, respectively, when performing TLB read and write operations. Figure 14-3 shows the format of these registers.

EntryLo0 and EntryLo1 Registers

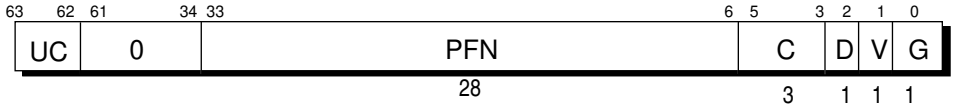

UC..... Uncached attribute.
PFN..... Page frame number; the upper bits of the physical address.
C ..... Specifies the TLB page coherency attribute.
D ..... Dirty. If this bit is set, the page is marked as dirty and, therefore, writable. This bit is actually a write-protect bit that software can use to prevent alteration of data.
V..... Valid. If this bit is set, it indicates that the TLB entry is valid; otherwise, a TLBL or TLBS invalid exception occurs.
G ..... Global. If this bit is set in both Lo0 and Lo1, then the processor ignores the ASID during TLB lookup.
0..... Reserved. Must be written as zeroes, and returns zeroes when read.

Figure 14-3 Fields of the EntryLo0 and EntryLo1 Registers

The *PFN* fields of the *EntryLo0* and *EntryLo1* registers span bits 33:6 of the 40-bit physical address.

Two additional bits for the mapped space's *uncached attribute* can be loaded into bits 63:62 of the *EntryLo* register, which are then written into the TLB with a TLB Write. During the address cycle of processor double/single/partial-word read and write requests, and during the address cycle of processor *uncached accelerated* block write requests, the processor drives the uncached attribute on **SysAD[59:58]**. The same *EntryLo* registers are used for the 64-bit and 32-bit addressing modes. In both modes the registers are 64 bits wide, however when the MIPS III ISA is not enabled (32-bit User and Supervisor modes) only the lower 32 bits of the *EntryLo* registers are accessible.

MIPS III is disabled when the processor is in 32-bit Supervisor or User mode. Loading of the integer registers is limited to bits *31:0*, sign-extended through bits *63:32*. *EntryLo[33:31]* or *PFN[39:38]* can only be set to all zeroes or all ones. In 32- and 64-bit modes, the *UC* and *PFN* bits of both *EntryLo* registers are written into the TLB. The *PFN* bits can be masked by setting bits in the *FrameMask* register (described in this chapter) but the *UC* bits cannot be masked or initialized in 32-bit User or Supervisor modes. In 32-bit Kernel mode, MIPS III is enabled and 64-bit operations are always available to program the *UC* bits.

There is only one *G* bit per TLB entry, and it is written with *EntryLo0[0]* and *EntryLo1[0]* on a TLB write.

14.4 Context (4)

The *Context* register is a read/write register containing the pointer to an entry in the page table entry (PTE) array; this array is an operating system data structure that stores virtual-to-physical address translations.

When there is a TLB miss, the CPU loads the TLB with the missing translation from the PTE array. Normally, the operating system uses the *Context* register to address the current page map which resides in the kernel-mapped segment, *kseg3*. The *Context* register duplicates some of the information provided in the *BadVAddr* register, but the information is arranged in a form that is more useful for a software TLB exception handler.

Figure 14-4 shows the format of the *Context* register; Table 14-4 describes the *Context* register fields.

Context Register

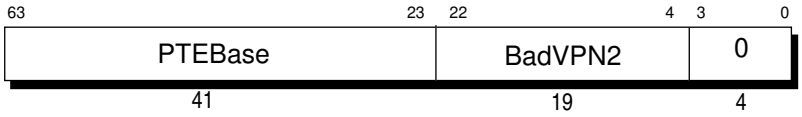

Figure 14-4 Context Register Format

Table 14-4 Context Register Fields

Field	Description
BadVPN2	This field is written by hardware on a miss. It contains the virtual page number (VPN) of the most recent virtual address that did not have a valid translation.
PTEBase	This field is a read/write field for use by the operating system. It is normally written with a value that allows the operating system to use the *Context* register as a pointer into the current PTE array in memory.

The 19-bit *BadVPN2* field contains bits 31:13 of the virtual address that caused the TLB miss; bit 12 is excluded because a single TLB entry maps to an even-odd page pair. For a 4-Kbyte page size, this format can directly address the pair-table of 8-byte PTEs. For other page and PTE sizes, shifting and masking this value produces the appropriate address.

14.5 PageMask Register (5)

The *PageMask* register is a read/write register used for reading from or writing to the TLB; it holds a comparison mask that sets the variable page size for each TLB entry, as shown in Table 14-5. Format of the register is shown in Figure 14-5.

TLB read and write operations use this register as either a source or a destination; when virtual addresses are presented for translation into physical address, the corresponding bits in the TLB identify which virtual address bits among bits 24:13 are used in the comparison. When the *Mask* field is not one of the values shown in Table 14-5, the operation of the TLB is undefined.

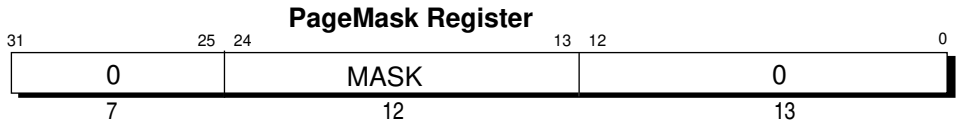

Mask.....Page comparison mask.
0.....Reserved. Must be written as zeroes, and returns zeroes when read.

Figure 14-5 PageMask Register

Table 14-5 Mask Field Values for Page Sizes

Page Size (Mask)	Bit											
	24	23	22	21	20	19	18	17	16	15	14	13
4 Kbytes	0	0	0	0	0	0	0	0	0	0	0	0
16 Kbytes	0	0	0	0	0	0	0	0	0	0	1	1
64 Kbytes	0	0	0	0	0	0	0	0	1	1	1	1
256 Kbytes	0	0	0	0	0	0	1	1	1	1	1	1
1 Mbyte	0	0	0	0	1	1	1	1	1	1	1	1
4 Mbytes	0	0	1	1	1	1	1	1	1	1	1	1
16 Mbytes	1	1	1	1	1	1	1	1	1	1	1	1

14.6 Wired Register (6)

The *Wired* register is a read/write register that specifies the boundary between the *wired* and *random* entries of the TLB as shown in Figure 14-6. Wired entries are fixed, nonreplaceable entries, which cannot be overwritten by a TLB write operation. Random entries can be overwritten.

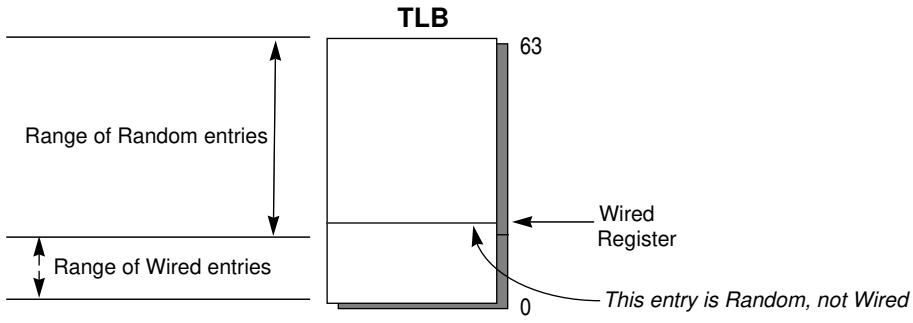

Figure 14-6 Wired Register Boundary

The *Wired* register is set to 0 upon system reset. Writing this register also sets the *Random* register to the value of its upper bound (see *Random* register, above). Figure 14-7 shows the format of the *Wired* register; Table 14-6 describes the register fields.

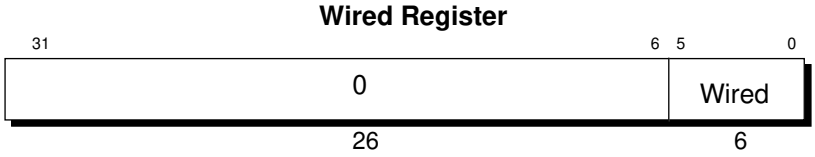

Figure 14-7 Wired Register

Table 14-6 Wired Register Field Descriptions

Field	Description
Wired	TLB Wired boundary
0	Reserved. Must be written as zeroes, and returns zeroes when read.

14.7 BadVAddr Register (8)

The Bad Virtual Address register (*BadVAddr*) is a read-only register that displays the most recent virtual address that caused either a TLB or Address Error exception. The *BadVAddr* register remains unchanged during Soft Reset, NMI, or Cache Error exceptions. Otherwise, the architecture leaves this register undefined.

Figure 14-8 shows the format of the *BadVAddr* register.

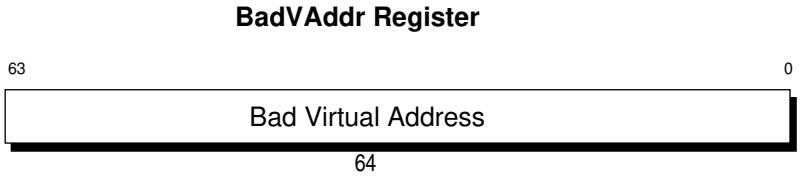

Figure 14-8 BadVAddr Register Format

14.8 Count and Compare Registers (9 and 11)

The *Count* and *Compare* registers are 32-bit read/write registers whose formats are shown in Figure 14-9.

The *Count* register acts as a real-time timer. Like the R4400 implementation, the R10000 *Count* register is incremented every *other* **PClk** cycle. However, unlike the R4400, the R10000 processor has no Timer Interrupt Enable boot-mode bit, so the only way to disable the timer interrupt is to negate the interrupt mask bit, *IM[7]*, in the Status register. This means the timer interrupt cannot be disabled without also disabling the *Performance Counter* interrupt, since they share *IM[7]*.

The *Compare* register can be programmed to generate an interrupt at a particular time, and is continually compared to the *Count* register. Whenever their values equal, the interrupt bit *IP[7]* in the *Cause* register is set. This interrupt bit is reset whenever the *Compare* register is written.

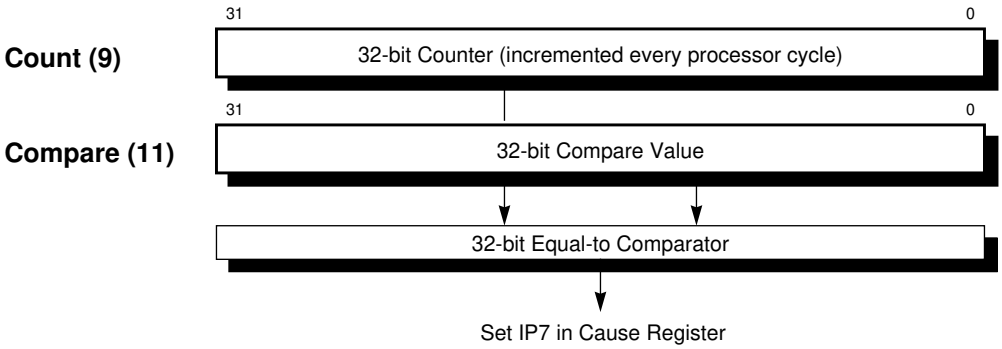

Figure 14-9 Count and Compare Registers

14.9 EntryHi Register (10)

The *EntryHi* register holds the high-order bits of a TLB entry for TLB read and write operations.

The *EntryHi* register is accessed by the TLB Probe, TLB Write Random, TLB Write Indexed, and TLB Read Indexed instructions.

Figure 14-10 shows the format of this register.

EntryHi Register

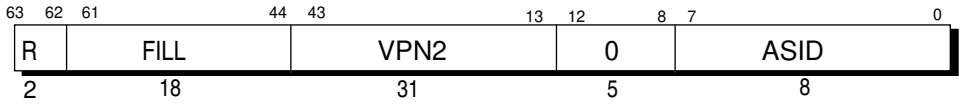

VPN2....Virtual page number divided by two (maps to two pages); upper bits of the virtual address
ASID.....Address space ID field. An 8-bit field that lets multiple processes share the TLB; each process has a distinct mapping of otherwise identical virtual page numbers.
R.....Region. (00 → user, 01 → supervisor, 11 → kernel) used to match $vAddr_{63...62}$
Fill.....Reserved. 0 on read; ignored on write.
0 .....Reserved. Must be written as zeroes, and returns zeroes when read.

Figure 14-10 EntryHi Register

In 64-bit addressing mode, the *VPN2* field contains bits 43:13 of the 44-bit virtual address.

In 32-bit addressing mode only the lower 32 bits of the *EntryHi* register are used, so the format remains the same as in the R4400 processor's 32-bit addressing mode. The *FILL* field is ignored on write and read as zeroes, as it was in the R4400 implementation.

When either a TLB refill, TLB invalid, or TLB modified exception occurs, the *EntryHi* register is loaded with the virtual page number (VPN2) and the ASID of the virtual address that did not have a matching TLB entry.

14.10 Status Register (12)

The *Status* register (SR) is a read/write register that contains the operating mode, interrupt enabling, and the diagnostic states of the processor. The following list describes the more important *Status* register fields; Figure 14-11 shows the format of the entire register, and Table 14-8 describes the *Status* register fields.

Some of the important fields include:

- The 4-bit *Coprocessor Usability* (*CU*) field controls the usability of 4 possible coprocessors. Regardless of the *CU0* bit setting, CP0 is always usable in Kernel mode. The *XX* bit enables the MIPS IV ISA in User mode.
- By default, the R10000 processor implements the same user instruction set as the R4400 processor. To enable execution of the MIPS IV instructions in User mode, the *MIPS IV User Mode* bit, *(XX*) of the CP0 *Status* register must be set.

  The MIPS IV instruction extension uses COP1X as the opcode; this designation was COP3 in the R4400 processor. For this reason the *CU3* bit is omitted in the R10000 processor, and is used as the *XX* bit. In *Kernel* and *Supervisor* modes, the state of the *XX* bit is ignored, and MIPS IV instructions are always available.

  Mode bit settings are shown in Table 14-7; dashes in the table represent *don't cares*.

Table 14-7 ISA and Status Register Settings for User, Supervisor and Kernel Mode Operations

Mode	UX	SX	KX	XX	MIPS II	MIPS III	MIPS IV
User	0	-	-	0	Yes	No	No
	0	-	-	1	Yes	No	Yes
	1	-	-	0	Yes	Yes	No
	1	-	-	1	Yes	Yes	Yes
Supervisor	-	0	-	-	Yes	No	Yes
	-	1	-	-	Yes	Yes	Yes
Kernel	-	-	-	-	Yes	Yes	Yes

NOTE: Operation with the MIPS IV ISA does not assume or require that the MIPS III instruction set or 64-bit addressing be enabled — *KX*, *SX* and *UX* may all be set to zero.

- The *Reduced Power (RP*) bit is reserved and should be zero. The R10000 processor does not define a reduced power mode.
- The *Reverse-Endian (RE)* bit, bit 25, reverses the endianness of the machine. The processor can be configured as either little-endian or big-endian at system reset; reverse-endian selection is available in Kernel and Supervisor modes, and in the User mode when the *RE* bit is 0. Setting the *RE* bit to 1 inverts the User mode endianness.
- The 9-bit *Diagnostic Status* (*DS*) field is used for self-testing, and checks the cache and virtual memory system. This field is described in Table 14-9 and Figure 14-12.
- The 8-bit *Interrupt Mask* (*IM*) field controls the enabling of eight interrupt conditions. Interrupts must be enabled before they can be asserted, and the corresponding bits are set in both the *Interrupt Mask* field of the *Status* register and the *Interrupt Pending* field of the *Cause* register.
- The processor mode is undefined if the *KSU* field is set to 3 (11_2). The R10000 processor implements this as User mode.

Status Register

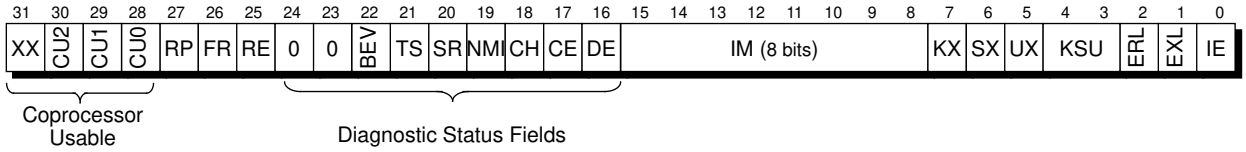

Figure 14-11 Status Register

Status Register Fields

Table 14-8 describes the *Status* register fields.

Table 14-8 Status Register Fields

Field	Description
XX	Enables execution of MIPS IV instructions in User mode. 1 → MIPS IV instructions usable 0 → MIPS IV instructions unusable
CU	Controls the usability of each of the four coprocessor unit numbers. CP0 is always usable when in Kernel mode, regardless of the setting of the CU_0 bit. 1 → usable 0 → unusable
RP	In the R4400 processor, this bit enables reduced-power operation by reducing the internal clock frequency. In the R10000 processor, this bit should be set to zero.
FR	Enables additional floating-point registers 0 → 16 registers 1 → 32 registers
RE	*Reverse-Endian* bit, valid in User mode.
DS	*Diagnostic Status* field (see Figure 14-12).
IM	*Interrupt Mask*: controls the enabling of each of the external, internal, and software interrupts. An interrupt is taken if interrupts are enabled, and the corresponding bits are set in both the *Interrupt Mask* field of the *Status* register and the *Interrupt Pending* field of the *Cause* register. 0 → disabled 1→ enabled
KX	Enables 64-bit addressing in Kernel mode. The extended-addressing TLB refill exception is used for TLB misses on kernel addresses. 0 → 32–bit 1 → 64–bit
SX	Enables 64-bit addressing and operations in Supervisor mode. The extended-addressing TLB refill exception is used for TLB misses on supervisor addresses. 0 → 32–bit 1 → 64–bit
UX	Enables 64-bit addressing and operations in User mode. The extended-addressing TLB refill exception is used for TLB misses on user addresses. 0 → 32–bit 1 → 64–bit

Table 14-8 (cont.) Status Register Fields

Field	Description
KSU	Mode bits $11_2 \rightarrow$ Undefined (implemented as User mode) $10_2 \rightarrow$ User $01_2 \rightarrow$ Supervisor $00_2 \rightarrow$ Kernel
ERL	Error Level; set by the processor when Reset, Soft Reset, NMI, or Cache Error exception are taken. 0 $\rightarrow$ normal 1 $\rightarrow$ error
EXL	Exception Level; set by the processor when any exception other than Reset, Soft Reset, NMI, or Cache Error exception are taken. 0 $\rightarrow$ normal 1 $\rightarrow$ exception
IE	Interrupt Enable 0 $\rightarrow$ disable all interrupts 1 $\rightarrow$ enables all interrupts

Diagnostic Status Field

The 9-bit *Diagnostic Status* (*DS*) field is used for self-testing, and checks the cache and virtual memory system. This field is described in Table 14-9 and shown Figure 14-12.

Some of the important *DS* fields include:

- In the R4400, the *TS* bit of the diagnostic field indicates a TLB *shutdown* has occurred due to matching of multiple virtual page entries during address translation. In the R10000 processor, the *TS* bit indicates a TLB write has introduced an entry that would allow matching of more than one virtual page entry during translation. In this case, the TLB entries that allow the multiple matches, even in the *Wired* area, are invalidated before the new TLB entry is written. This prevents multiple matches during address translation.

  The *TS* bit is updated for each TLB write. It can also be read and written by software (in the R4400, the *TS* bit is read-only); to clear the *TS* bit one needs to write a 0 into it. As in the R4400, Reset/Soft Reset/NMI exceptions also clear the *TS* bit.
- The *NMI* bit is new to the R10000 processor; it distinguishes between Soft Reset and NMI exceptions. Both exceptions set the *SR* bit to 1; the NMI exception sets the *NMI* bit to 1, whereas the Soft Reset exception sets it to 0.
- The *CE* bit is reserved in the R10000 processor and should be a 0.

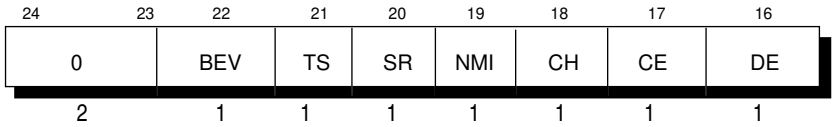

Figure 14-12 Diagnostic Status Field

Table 14-9 Status Register Diagnostic Status Bits

Bit	Description
BEV	Controls the location of TLB refill and general exception vectors. 0 → normal 1→ bootstrap
TS	This bit is set when a TLB write presents an entry that matches any other virtual page entry in the TLB. Should this occur, any TLB entries that allow multiple matches, even in the *Wired* area, are invalidated before this new entry can be written into the TLB. This prevents multiple matches during address translation. 0 → normal 1→ TLB shutdown has occurred.
SR	1→ Indicates a Soft Reset or NMI exception.
NMI	1→ Indicates a nonmaskable interrupt has occurred. Used to distinguish between a Soft Reset and a nonmaskable interrupt in a Soft Reset exception.
CH	Hit (tag match and valid state) or miss indication for last CACHE Hit Invalidate, Hit Write Back Invalidate for a secondary cache. 0 → miss 1 → hit
CE	Reserved in the R10000, and should be set to 0.
DE	Specifies that cache parity or ECC errors cannot cause exceptions. 0 → parity/ECC remain enabled 1 → disables parity/ECC
0	Reserved. Must be written as zeroes, and returns zeroes when read.

Coprocessor Accessibility

Three *Status* register *CU* bits control coprocessor accessibility: *CU0*, *CU1*, and *CU2* enable coprocessors 0, 1, and 2, respectively. If a coprocessor is unusable, any instruction that accesses it generates an exception.

The following describes the coprocessor implementations and operations on the R10000:

- Coprocessor 0 is always enabled in kernel mode, regardless of the *CU0* bit.
- Coprocessor 1 is the floating-point coprocessor. If *CU1* is 0 (disabled), all floating-point instructions generate a Coprocessor Unusable exception. In MIPS IV, the COP3 instruction is replaced with a second floating-point instruction, COP1X. In addition, new functions are added to COP1 (see Chapter 15). The floating-point branch conditional and compare instructions are expanded to use the eight Floating-Point *Status* register condition bits, instead of the original single bit. If any of these extra bits are referenced (cc > 0) when not using the MIPS IV ISA, an Unimplemented Instruction exception is taken. The integer conditional move (MOVC) instruction tests a floating-point condition bit; it causes a coprocessor unusable exception if coprocessor 1 is disabled.
- Coprocessor 2 is defined, but does not exist in the R10000; its instructions (COP2, LWC2, LDC2, SWC2, SDC2) always cause an exception, but the exception code depends upon whether the coprocessor, as indicated by *CU2*, is enabled.
- Coprocessor 3 has been removed from the MIPS III ISA, and is no longer defined. If MIPS IV is disabled, the coprocessor 3 instruction (COP3) always causes a Reserved Instruction exception.

14.11 Cause Register (13)

The 32-bit read/write *Cause* register describes the cause of the most recent exception.

Figure 14-13 shows the fields of this register; Table 14-10 describes the *Cause* register fields. A 5-bit exception code (*ExcCode*) indicates one of the causes, as listed in Table 14-11.

All bits in the *Cause* register, with the exception of the *IP[1:0]* bits, are read-only; *IP[1:0]* are used for software interrupts.

Table 14-10 Cause Register Fields

Field	Description
BD	Indicates whether the last exception taken occurred in a branch delay slot. 1 → delay slot 0 → normal
CE	Coprocessor unit number referenced when a Coprocessor Unusable exception is taken. This bit is undefined for any other exception.
IP	Indicates an interrupt is pending. This bit remains unchanged for NMI, Soft Reset, and Cache Error exceptions. 1 → interrupt pending 0 → no interrupt
ExcCode	Exception code field (see Table 14-11)
0	Reserved. Must be written as zeroes, and returns zeroes when read.

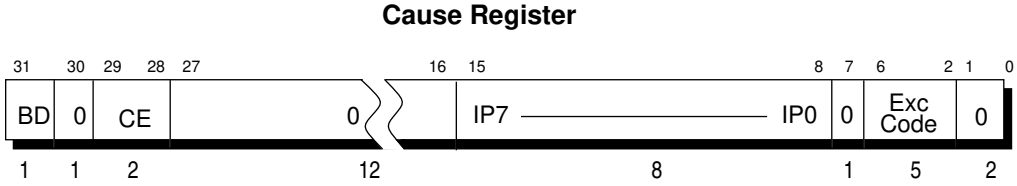

Figure 14-13 Cause Register Format

Table 14-11 Cause Register ExcCode Field

Exception Code Value	Mnemonic	Description
0	Int	Interrupt
1	Mod	TLB modification exception
2	TLBL	TLB exception (load or instruction fetch)
3	TLBS	TLB exception (store)
4	AdEL	Address error exception (load or instruction fetch)
5	AdES	Address error exception (store)
6	IBE	Bus error exception (instruction fetch)
7	DBE	Bus error exception (data reference: load or store)
8	Sys	Syscall exception
9	Bp	Breakpoint exception
10	RI	Reserved instruction exception
11	CpU	Coprocessor Unusable exception
12	Ov	Arithmetic Overflow exception
13	Tr	Trap exception
14	–	Reserved
15	FPE	Floating-Point exception
16–22	–	Reserved
23	WATCH	Reference to *WatchHi/WatchLo* address
24–30	–	Reserved
31	–	Reserved

14.12 Exception Program Counter (14)

The Exception Program Counter (*EPC*)† is a read/write register that contains the address at which processing resumes after an exception has been serviced.

For synchronous exceptions, the *EPC* register contains either:

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

The processor does not write to the *EPC* register when the *EXL* bit in the *Status* register is set to a 1.

Figure 14-14 shows the format of the *EPC* register.

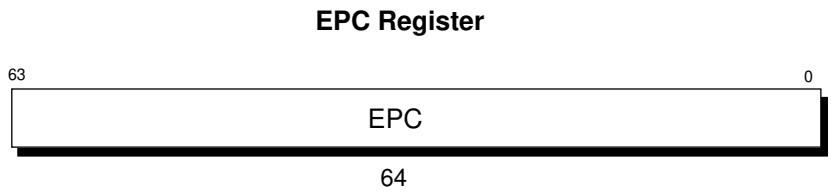

Figure 14-14 EPC Register Format

† The *ErrorEPC* register provides a similar capability, described later in this chapter.

14.13 Processor Revision Identifier (PRId) Register (15)

The 32-bit, read-only *Processor Revision Identifier* (*PRId*) register contains information identifying the implementation and revision level of the CPU and CP0. Figure 14-15 shows the format of the *PRId* register; Table 14-12 describes the *PRId* register fields.

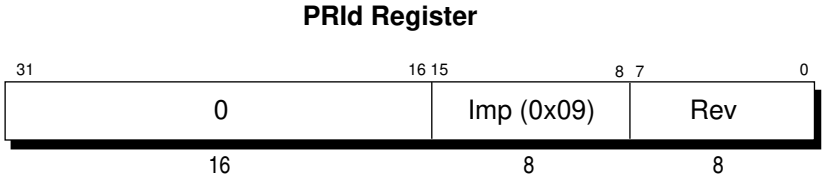

Figure 14-15 Processor Revision Identifier Register Format

Table 14-12 PRId Register Fields

Field	Description
Imp	Implementation number
Rev	Revision number
0	Reserved. Must be written as zeroes, and returns zeroes when read.

The low-order byte (bits 7:0) of the *PRId* register is interpreted as a revision number, and the high-order byte (bits 15:8) is interpreted as an implementation number. The implementation number of the R10000 processor is 0x09. The content of the high-order halfword (bits 31:16) of the register are reserved.

The revision number is stored as a value in the form *y.x*, where *y* is a major revision number in bits 7:4 and *x* is a minor revision number in bits 3:0.

The revision number can distinguish some chip revisions, however there is no guarantee that changes to the chip will necessarily be reflected in the *PRId* register, or that changes to the revision number necessarily reflect real chip changes. For this reason, software should not rely on the revision number in the *PRId* register to characterize the chip.

14.14 Config Register (16)

The R10000 processor's *Config* register has a different format from that of the R4400, since the R10000 processor has different mode bits and configurations, however some fields are still compatible: *K0*, *DC*, *IC*, and *BE*. The value of bits 24:0 are taken directly from the Mode bit settings during a reset sequence; refer to Table 8-1 for these bit definitions. Table 14-13 shows the R10000 *Config* register fields, along with values which are hardwired into the register at boot time; Figure 14-16 shows the *Config* register format.

Table 14-13 Config Register Field Definitions

Field	Bits	Name	Hardwired Values
K0	2:0	Coherency algorithm 000_2 → reserved 001_2 → reserved 010_2 → uncached 011_2 → cacheable noncoherent 100_2 → cacheable coherent exclusive 101_2 → cacheable coherent exclusive on write 110_2 → reserved 111_2 → uncached accelerated	
DN	4:3	Device number	
CT	5	CohPrcReqTar	
PE	6	PrcElmReq	
PM	8:7	PrcReqMax	
EC	12:9	SysClkDiv	
SB	13	SCBlkSize	
SK	14	SCCorEn	
BE	15	MemEnd	
SS	18:16	SCSize	
SC	21:19	SCClkDiv	
	25:22	Reserved	0
DC	28:26	Primary data cache size (hardwired to 011_2)	32 Kbytes
IC	31:29	Primary instruction cache size (hardwired to 011_2)	32 Kbytes

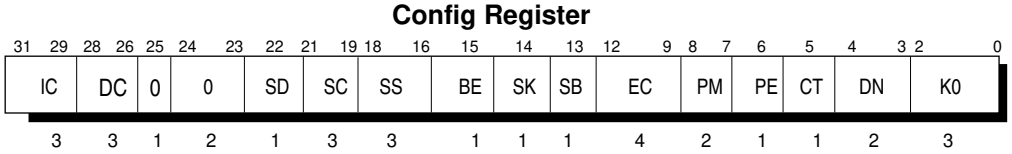

Figure 14-16 Config Register Format

14.15 Load Linked Address (LLAddr) Register (17)

Physical addresses for Load Link instructions are no longer written into this register. *LLAddr* is implemented as a read/write scratch register used for NT compatibility.

Figure 14-17 shows the format of the *LLAddr* register.

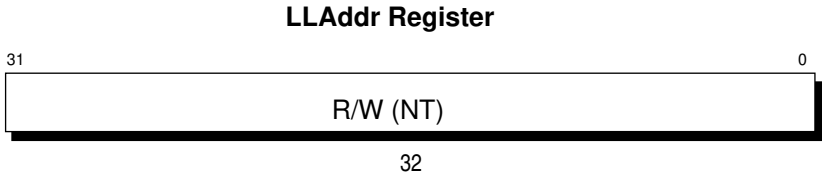

Figure 14-17 LLAddr Register Format

14.16 WatchLo (18) and WatchHi (19) Registers

WatchHi and *WatchLo* are 32-bit read/write registers which contain a physical address of a doubleword location in main memory. If enabled, any attempt to read or write this location causes a Watch exception. This feature is used for debugging.

Bits 7:0 of the *WatchHi* register contain bits 39:32 of the trap physical address, shown in Figure 14-18. The WatchLo register contains physical address bits 31:3. The remaining bits of the register are ignored on write and read as zero.

Table 14-14 describes the *WatchLo* and *WatchHi* register fields.

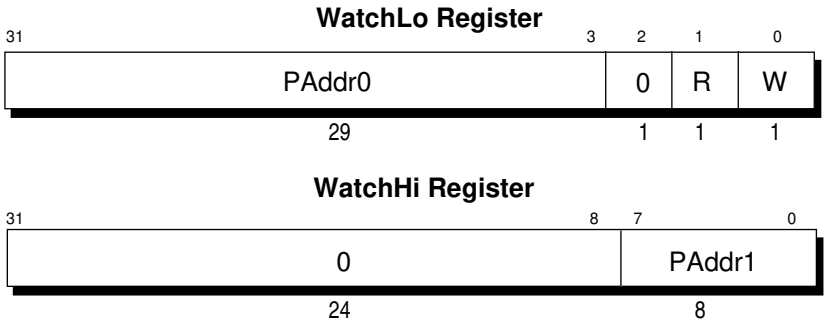

Figure 14-18 WatchLo and WatchHi Register Formats

Table 14-14 WatchHi and WatchLo Register Fields

Field	Description
PAddr1	Bits 39:32 of the physical address
PAddr0	Bits 31:3 of the physical address
R	Trap on load references if set to 1
W	Trap on store references if set to 1
0	Ignored on write and read as zero.

14.17 XContext Register (20)

The read/write *XContext* register contains a pointer to an entry in the page table entry (PTE) array, an operating system data structure that stores virtual-to-physical address translations. When there is a TLB miss, the operating system software loads the TLB with the missing translation from the PTE array. The *XContext* register no longer shares the information provided in the *BadVAddr* register, as it did in the R4400.

The *XContext* register is for use with the XTLB refill handler, which loads TLB entries for references to a 64-bit address space, and is included solely for operating system use. The operating system sets the PTE base field in the register, as needed. Normally, the operating system uses the *Context* register to address the current page map, which resides in the kernel-mapped segment *kseg3*.

Figure 14-19 shows the format of the *XContext* register; Table 14-15 describes the *XContext* register fields.

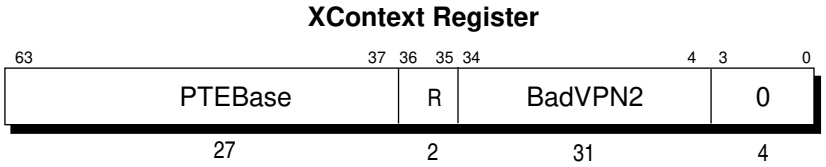

Figure 14-19 XContext Register Format

The 31-bit *BadVPN2* field holds bits 43:13 of the virtual address that caused the TLB miss; bit 12 is excluded because a single TLB entry maps to an even-odd page pair. For a 4-Kbyte page size, this format may be used directly to address the pair-table of 8-byte PTEs. For other page and PTE sizes, shifting and masking this value produces the appropriate address.

Table 14-15 XContext Register Fields

Field	Description
BadVPN2	The *Bad Virtual Page Number*/2 field is written by hardware on a miss. It contains the VPN of the most recent invalidly translated virtual address.
R	The *Region* field contains bits 63:62 of the virtual address. 00_2 = user 01_2 = supervisor 11_2 = kernel.
PTEBase	The *Page Table Entry Base* read/write field is normally written with a value that allows the operating system to use the *Context* register as a pointer into the current PTE array in memory.

14.18 FrameMask Register (21)

The *FrameMask* register is new with the R10000 processor. It masks bits of the *EntryLo0* and *EntryLo1* registers so that these masked bits are not passed to the TLB while doing a TLB write (either TLBWI or TLBWR).

A zero in the *FrameMask* register allows its corresponding bit in the *EntryLo*[*1,0]* registers to pass to the TLB; a one in the *FrameMask* register masks off its corresponding bit in the *EntryLo* registers and passes a zero to the TLB. Bits 15:0 of the *FrameMask* register control bits 33:18 of the *EntryLo* registers.

The remaining bits of this register are ignored on write and read as zeroes. The content of this register is set to zero after a processor reset or a power-up event.

Figure 14-20 shows the *FrameMask* register format.

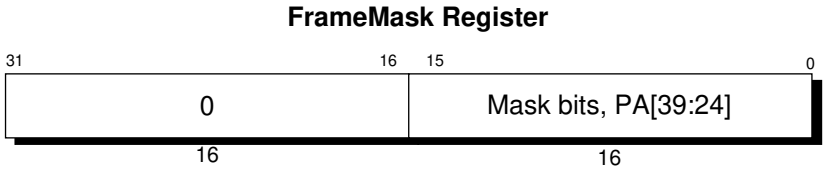

Figure 14-20 Framemask Register Format

14.19 Diagnostic Register (22)

CP0 register 22, the *Diagnostic* register, is a new 64-bit register for processor-specific diagnostic functions. (Since this register is designed for local use, the diagnostic functions are subject to change without notice.) Currently, this register helps test the ITLB, branch caches, and the branch prediction scheme. In addition, it provides choices for branch prediction algorithms, to help diagnostic program writing.

The ten fields of the *Diagnostic* register, shown in Figure 14-21, are described below. All fields are read-only (all writes are ignored).

ITLBM: this field is a 4-bit read-only counter. This field is incremented by one for each ITLB miss, and any overflow is ignored. Its value is undefined during reset, and its value is meaningless when used in an unmapped space.

BSIdx: this field defines the entry in the branch stack to be used for the latest conditional branch decoded. Its value is meaningless if the latest branch was an unconditional branch.

DBRC: this field disables the use of the branch return cache (BRC).

BRCV: this field indicates whether or not the branch return cache (BRC) is valid. BRC has only one entry (four instructions).

BRCW: this field indicates whether or not the latest branch (JAL, JALR RX, BGEZAL, BGEZALL, BLTZAL, or BLTZALL) caused a write into BRC. It is not affected by any other type of branch.

BRCH: this field indicates whether or not the latest branch (JR r31 or JALR rx,r31) has a BRC hit. It not affected by any other type of branch.

MP: this field indicates whether or not the latest conditional branch verified was mispredicted.

BPMode: this is a read-write field for branch prediction algorithm control.

00_2: 2-bit counter scheme

01_2: all conditional branches are predicted not taken

10_2: all conditional branches are predicted taken

11_2: forward conditional branches are predicted not taken and backward conditional branches are predicted taken.

The default mode is 00 on processor reset.

BPState: this field contains the new 2-bit state for a conditional branch after it is verified. It is also used to hold the 2-bit state to read/write when a branch prediction table read/write operation is executed.

BPIdx: this field contains the index to the *BPT* for *BPT* read/write/initialization operations, and should contain *VA[11:3]* of the branch for BPT read/write operations. The upper six bits of the *BPIdx* field contain the line address for BPT line initialization operations; the lower three bits of *BPIdx* are ignored.

BPOp: this field indicates the following BPT operations:

00_2: BPT read

01_2: BPT write

10_2: initializes BPT line to all zeroes (strongly not taken)

11_2: initializes BPT line to all ones (strongly taken).

Figure 14-21 shows the format of the *Diagnostic* register.

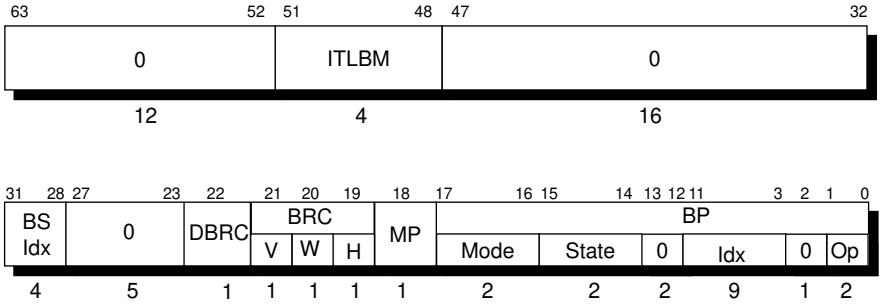

Figure 14-21 Diagnostic Register Format

14.20 Performance Counter Registers (25)

The R10000 processor defines two performance counters and two associated control registers, which are mapped into CP0 register 25. An encoding in the MTC0/MFC0 instructions on register 25 indicates which counter or control register is used.

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

The counter asserts an interrupt, *IP[7]*, when its most significant bit (bit 31) becomes one (the counter overflows) and the associated performance control register enables the interrupt.

The counting continues after counter overflow whether or not an interrupt is signalled.

The format of the control registers are shown in Figure 14-22.

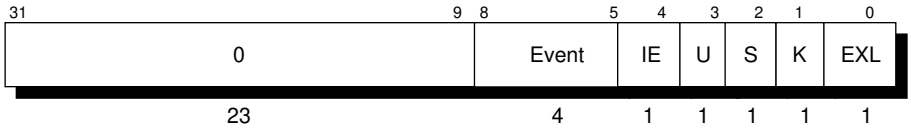

Figure 14-22 Control Register Format

The fields of the *Control* register are:

- The *Event* field specifies the event to be counted, listed in Table 14-16.

Table 14-16 Counter Events

Event	Counter 0	Counter 1
0	Cycles	Cycles
1	Instructions issued	Instructions graduated
2	Load/prefetch/sync/CacheOp issued	Load/prefetch/sync/CacheOp graduated
3	Stores issued	Stores graduated
4	Store conditional issued	Store conditional graduated
5	Failed store conditional	Floating-point instructions graduated
6	Branches decoded	Quadwords written back from primary data cache
7	Quadwords written back from secondary cache	TLB refill exceptions
8	Correctable ECC errors on secondary cache data	Branches mispredicted
9	Instruction cache misses	Data cache misses
10	Secondary cache misses (instruction)	Secondary cache misses (data)
11	Secondary cache way mispredicted (instruction)	Secondary cache way mispredicted (data)
12	External intervention requests	External intervention request hits in secondary cache
13	External invalidate requests	External invalidate request hits in secondary cache
14	Virtual coherency condition	Stores or prefetches with store hint to *CleanExclusive* secondary cache blocks
15	Instructions graduated	Stores or prefetches with store hint to *Shared* secondary cache blocks

- The *IE* bit enables the assertion of *IP[7]* when the associated counter overflows.
- The *U*, *S*, *K*, and *EXL* bits indicate the processor modes in which the event is counted: U is user mode; S is supervisor mode; K is kernel mode when *EXL* and *ERL* both are set to 0; the system is in kernel mode and handling an exception when *EXL* is set to 1, as shown in Table 14-20.

These modes can be set individually; for example, one could set all four bits to count a certain event in all processor modes except during a cache error exception.

The following rules apply:

- R10000 defines "issue" as the point in time that an instruction becomes active and is issued to a functional unit for execution.
- R10000 defines "graduation" as the point in time that the functional unit successfully completes the execution of an instruction. It is the oldest active instruction.
- The number of executed branches is the difference between decoded branches and mispredicted branches. Although newer branches may be speculatively decoded before an older branch is resolved as mispredicted, these newer branch instructions, nullified by the older misprediction, are not counted as mispredicted.
- When an integer multiply or divide instruction graduates, it is counted as two graduated instructions.
- Any instruction that sets the FP *Status* register bits (EVZOUI) is counted as a graduated floating point instruction.
- TLB misses are counted as either nascent or within the TLB handler.

The performance counters and associated control registers are written by using an MTC0 instruction, as shown in Table 14-17.

Table 14-17 Writing Performance Registers Using MTC0

Opcode[15:11]	Opcode[1:0]	Operation
11001	00	Move to Performance Control 0
11001	01	Move to Performance Counter 0
11001	10	Move to Performance Control 1
11001	11	Move to Performance Counter 1

The performance counters and associated control registers are read by using a MFC0 instruction, as shown in Table 14-18.

Table 14-18 Reading Performance Registers Using MFC0

Opcode[15:11]	Opcode[1:0]	Operation
11001	00	Move from Performance Control 0
11001	01	Move from Performance Counter 0
11001	10	Move from Performance Control 1
11001	11	Move from Performance Counter 1

The format of the performance control registers are shown in Table 14-19.

Table 14-19 Performance Control Register Format

[8:5]	[4]	[3:0]
Event select	IP[7] interrupt enable	Count enable (U/S/K/EXL)

The count enable field specifies whether counting is to be enabled during User, Supervisor, Kernel, and/or Exception level mode. Any combination of count enable bits may be asserted.

All unused bits in the performance control registers are reserved.

All counting is disabled when the *ERL* bit of the CP0 *Status* register is asserted.

Table 14-20 defines the operation of the count enable bits of the performance control registers.

Table 14-20 Count Enable Bit Definition

Count Enable Bit	Count Qualifier (CP0 *Status* Register Fields)
U	KSU = 2 (User mode), EXL = 0, ERL = 0
S	KSU = 1 (Supervisor mode), EXL = 0, ERL = 0
K	KSU = 0 (Kernel mode), EXL = 0, ERL = 0
EXL	EXL = 1, ERL = 0

The following rules apply:

- The performance counter registers may be preloaded with an MTC0 instruction, and counting is enabled by asserting one or more of the count enable bits in the performance control registers.
- The interrupt enable bit must be asserted to cause *IP[7]*.
- To determine the cause of the interrupt, the interrupt handler routine must query the following:
  - the performance counter register
  - the interrupt enable bit of the associated performance control register of both counters
- If neither of the counters caused the interrupt, *IP[7]* must be the result of the CP0 *Count* register matching the CP0 *Compare* register.

14.21 ECC Register (26)

The R10000 processor implements a 10-bit read/write *ECC* register which is used to read and write the secondary cache data ECC or the primary cache data parity bits. (Tag ECC and parity are loaded to and stored from the *TagLo* register.) Unlike the R4400, the only CacheOps that use *ECC* register are *Index Load Data* and *Index Store Data*.

In the R4400, both the primary instruction and data caches are parity byte-protected.

In the R10000 processor, the following protection schemes are used:

- The primary instruction cache is word-protected (where one word contains 36 bits), and one parity bit is used for each instruction word (*IP* in Figure 14-23).
- The primary data cache is byte-protected, with four bits used for each 32-bit data word (*DP* in Figure 14-23).
- Each quadword of the secondary cache data uses nine bits of ECC and one bit of parity (*SP* and *ECC* in Figure 14-23).

The primary instruction CacheOps load or store one instruction word at a time; therefore, one bit is used in the ECC register. The primary data CacheOps load or store four bytes at a time; therefore, four bits are used in the ECC register. The secondary CacheOps use **ECC[9]** as the parity bit and **ECC[8:0]** as the 9-bit ECC. For the *Index Store Data* cacheops, the unused bits are ignored. For *Index Load Data* CacheOps, the unused a bits are with zeroes.

Figure 14-23 shows the format of the *ECC* register; Table 14-21 describes the register fields.

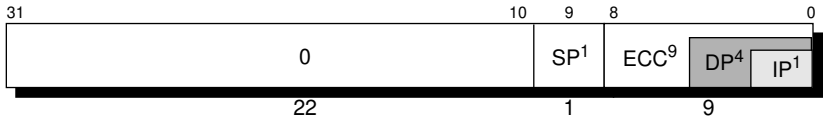

Figure 14-23 ECC Register Format

Table 14-21 ECC Register Fields

Field	Description
SP	A 1-bit field specifying the parity bit read from or written to a secondary cache.
ECC	An 9-bit field specifying the ECC bits read from or written to a secondary cache.
DP	An 4-bit field specifying the parity bits read from or written to a primary data cache.
IP	An 1-bit field specifying the parity bit read from or written to a primary instruction cache.
0	Reserved. Must be written as zeroes, and returns zeroes when read.

14.22 CacheErr Register (27)

The *CacheErr* register is a 32-bit read-only register that handles ECC errors in the secondary cache or system interface, and parity errors in the primary caches.

R10000 processor correction policy is as follows:

- Parity errors cannot be corrected.
- Single-bit ECC errors can be corrected by hardware without taking a Cache Error exception.
- Double-bit ECC errors can be detected but not corrected by hardware.
- All uncorrectable errors take Cache Error exceptions unless the *DE* bit of the *Status* register is set.
- As in the R4400, cache errors are imprecise.

The *CacheErr* register provides cache index and status bits which indicate the source and nature of the error; it is loaded when a Cache Error exception is taken.

CacheErr Register Format for Primary Instruction Cache Errors

Figure 14-24 shows the format of the *CacheErr* register when a primary instruction cache error occurs.

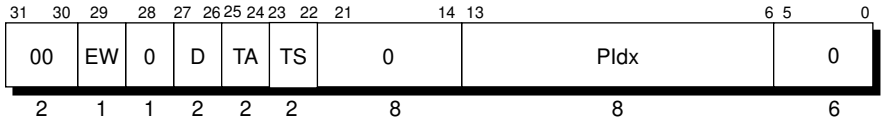

Figure 14-24 CacheErr Register Format for Primary Instruction Cache Errors

EW: set when *CacheErr* register is already holding the values of a previous error

D: data array error (way1 || way0)

TA: tag address array error (way1 || way0)

TS: tag state array error (way1 || way0)

PIdx: primary cache virtual block index, **VA[13:6]**

CacheErr Register Format for Primary Data Cache Errors

Figure 14-25 shows the format of the *CacheErr* register when a primary data cache error occurs.

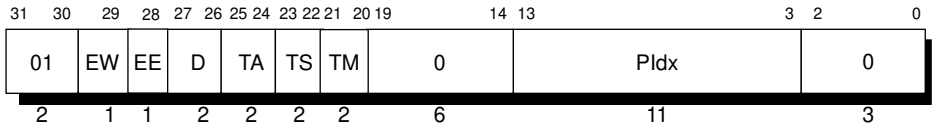

Figure 14-25 CacheErr Register Format for Primary Data Cache Errors

EW: set when *CacheErr* register is already holding the values of a previous error

EE: tag error on an inconsistent block

D: data array error (way1 || way0)

TA: tag address array error (way1 || way0)

TS: tag state array error (way1 || way0)

TM: tag mod array error (way1 || way0)

PIdx: primary cache virtual double word index, **VA[13:6]**

CacheErr Register Format for Secondary Cache Errors

Figure 14-26 shows the format of the *CacheErr* register when a secondary cache error occurs.

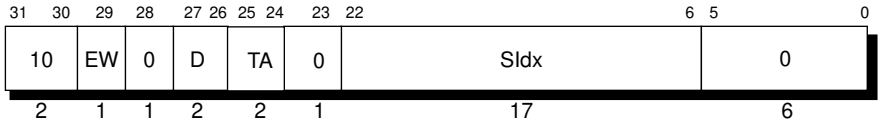

Figure 14-26 CacheErr Register Format for Secondary Cache Errors

EW: set when *CacheErr* register is already holding the values of a previous error

D: uncorrectable data array error (way1 || way0)

TA: uncorrectable tag array error (way1 || way0)

SIdx: secondary cache physical block index (**PA[22:6]** for 16-word block size or **PA[22:7]** for 32-word block size)

CacheErr Register Format for System Interface Errors

Figure 14-27 shows the format of the *CacheErr* register when a System interface error occurs.

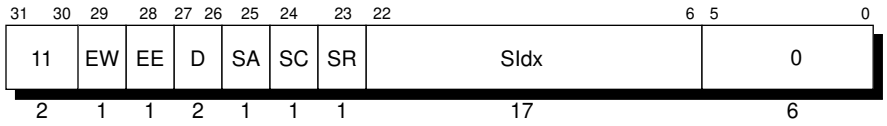

Figure 14-27 CacheErr Register Format for System Interface Errors

EW: set when *CacheErr* register is already holding the values of a previous error

EE: data error on a *CleanExclusive* or *DirtyExclusive*

D: uncorrectable system block data response error (way1 || way0)

SA: uncorrectable system address bus error

*S*C: uncorrectable system command bus error

*S*R: uncorrectable system response bus error

SIdx: secondary cache physical block index

14.23 TagLo (28) and TagHi (29) Registers

The *TagHi* and *TagLo* registers are 32-bit read/write registers used to hold the following:†

- the primary cache tag and parity
- the secondary cache tag and ECC
- the data in primary or secondary caches for certain CacheOps

TagHi/Lo formats in the R10000 processor differ from those in the R4400 due to changes in CacheOps and cache architecture. R10000 formats depend on the type of CacheOp executed and the cache to which it is applied. The reserved fields are read as zeroes after executing an *Index Load Tag* or an *Index Load Data* CacheOp and ignored when executing an *Index Store Tag* or an *Index Store Data* CacheOp.

To ensure NT kernel compatibility, the *TagLo* register is implemented as a 32-bit read/write register. The value written by an MTC0 instruction can be retrieved by a MFC0 instruction, unless an intervening CACHE instruction has modified the content.

This section gives the TagLo and TagHi register formats for the following CacheOp and cache combinations:

- CacheOp is Index Load/Store Tag
  - primary instruction cache operation
  - primary data cache operation
  - secondary cache operation
- CacheOp is Index Load/Store Data
  - primary instruction cache operation
  - primary data cache operation
  - secondary cache operation

† To ensure NT kernel compatibility, the TagLo register is implemented as a 32-bit read/write register. The value written by a MTC0 instruction can be retrieved by a MFC0 instruction, unless intervening CACHE instructions modify the content.

CacheOp is Index Load/Store Tag

This section describes the three states of the *TagLo* and *TagHi* registers, when the CacheOp is an *Index Load/Store Tag* for the following operations:

- primary instruction cache operation
- primary data cache operation
- secondary cache operation

Primary Instruction Cache Operation

If the CacheOp is an *Index Load/Store Tag* for a primary instruction cache operation, the fields of the *TagHi* and *TagLo* registers are defined as follows:

PTag0: contains physical address bits [35:12] stored in the cache tag

PState: contains the primary instruction cache state for the line, as follows:

1 = *Valid*

0 = *Invalid*

LRU: indicates whether or not this line is the least recently used of the set

SP: state even parity bit for the *PState* field

TP: tag even parity bit.

PTag1: contains physical address bits [39:36] stored in the cache tag

Figure 14-28 shows the fields of the *TagHi* and *TagLo* registers.

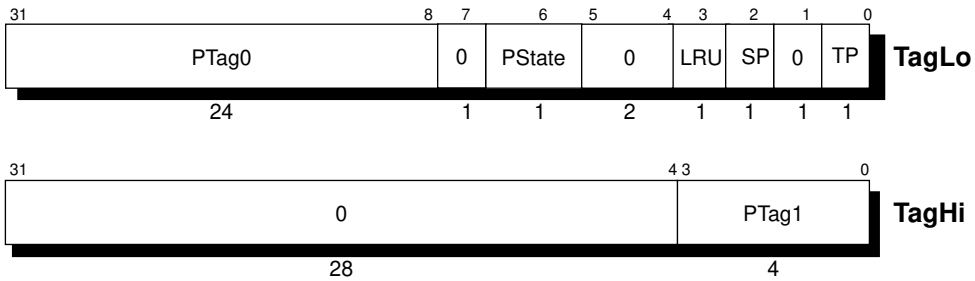

Figure 14-28 TagHi/Lo Register Fields in Primary Instruction Cache When CacheOp is Index Load/Store Tag

Primary Data Cache Operation

If the CacheOp is an *Index Load/Store Tag* for primary data cache operations, the fields of the *TagHi* and *TagLo* registers are defined as follows:

State Modifier: holds the status of the line, as follows:

001_2 = neither refilled or written

010_2 = this line may have been written and inconsistent from the secondary cache (*W* bit)

100_2 = this line is being refilled (*Refill* bit).

PTag1: contains physical address bits [39:36] stored in the cache tag

PTag0: contains physical address bits [35:12] stored in the cache tag

PState: together with the *Refill* bit of the *State Modifier* in the *TagHi* register, *PState* determines the state of the cache block in the primary data cache, as shown in Table 14-22.

Table 14-22 PState Field Definition in TagHi/Lo Registers, For Primary Data Cache Operation When CacheOp is Index Load/Store Tag

PState	Refill=0	Refill=1
00_2	Invalid	Refill *clean* (block is being refilled)
01_2	Shared	Upgrade Share (converting *shared* to *dirty*)
10_2	Clean Exclusive	Upgrade Clean (converting *clean* to *dirty*).
11_2	Dirty Exclusive	Refill *dirty* (block is being refilled for a store)

LRU: indicates whether this line was the least recently used of the set

SP: state even parity bit for the *PState* field and the *Way* bit

Way: indicates which secondary cache set contains the primary cache line for this tag

TP: tag even parity bit.

Figure 14-29 shows the fields of the *TagHi* and *TagLo* registers.

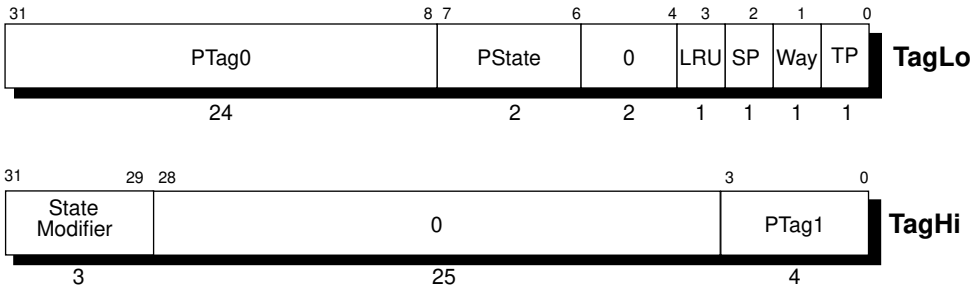

Figure 14-29 TagHi/Lo Register Fields in Primary Data Cache When CacheOp is Index Load/Store Tag

Secondary Cache Operation

If the CacheOp is an *Index Load/Store Tag* for secondary cache operations, the fields of the *TagHi* and *TagLo* registers are defined as follows:

STag0: contains physical address bits [35:18] stored in the cache tag

SState: contains the secondary cache state of the line, *as follows:*

00_2 = *Invalid*

01_2 = *Shared*

10_2 = *Clean Exclusive*

11_2 = *Dirty Exclusive*

VIndex (virtual index): contains only two bits of significance since the32 Kbyte 2-way set associative primary caches are addressed using only two untranslated address bits (**VA[13:12]**) plus the offset within the virtual page.

ECC: contains the ECC for the *STag*, *SState* and *VIndex* fields.

MRU: indicates whether this line was the most recently used in the set

STag1: contains the physical address bits [39:36] stored in the cache tag.

Figure 14-30 shows the fields of the *TagHi* and *TagLo* registers.

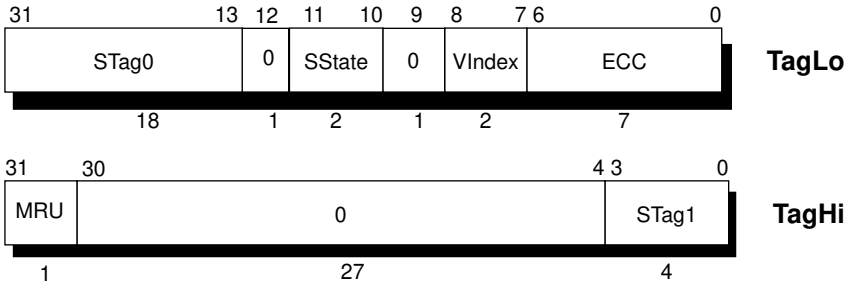

Figure 14-30 TagHi/Lo Register Fields in Secondary Cache When CacheOp is Index Load/Store Tag

CacheOp is Index Load/Store Data

This section describes the following three states of the *TagLo* and *TagHi* registers, when the CacheOp is an *Index Load/Store Data*:

- primary instruction cache operation
- primary data cache operation
- secondary cache operation

Primary Instruction Cache Operation

If the CacheOp is an *Index Load/Store Data* for the primary instruction cache, the *TagHi* register stores the most significant four bits of a 36-bit instruction, as shown in Figure 14-31; the rest of the instruction is stored in the *TagLo* register.

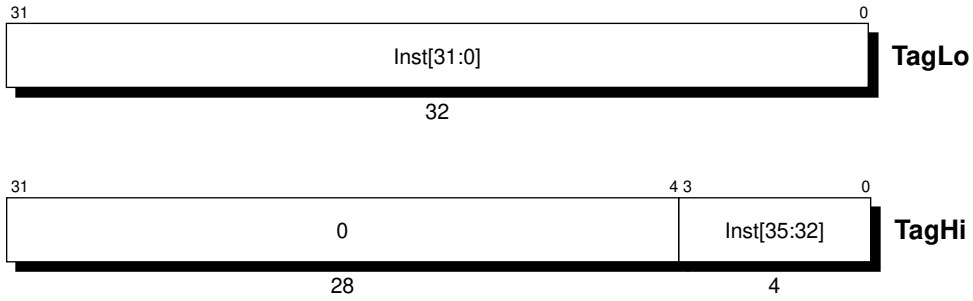

Figure 14-31 TagHi/Lo Register Fields in Primary Instruction Cache When CacheOp is Index Load/Store Data

Primary Data Cache Operation

If the CacheOp is *Index Load/Store Data* for primary data cache, the *TagHi* register is not used. The *TagLo* registers contains a 32-bit data word for the cache operation, as shown in Figure 14-32.

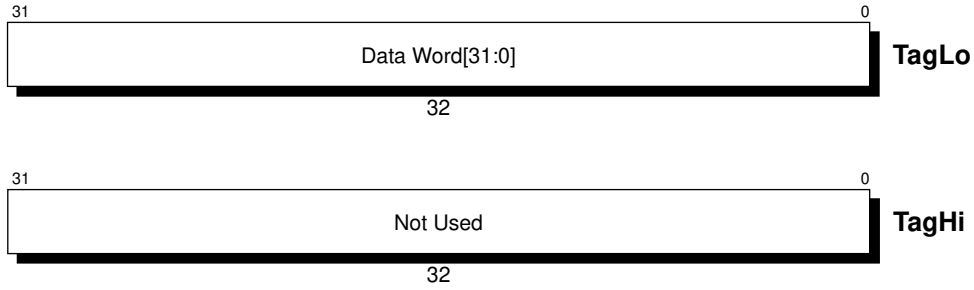

Figure 14-32 TagHi/Lo Register Fields in Primary Data Cache When CacheOp is Index Load/Store Data

Secondary Cache Operation

If the CacheOp is *Index Load/Store Data* for the secondary cache, a doubleword of data is required for the CacheOp. The *TagHi* register stores the upper 32 bits of the doubleword and the *TagLo* register stores the lower 32 bits, as shown below in Figure 14-33.

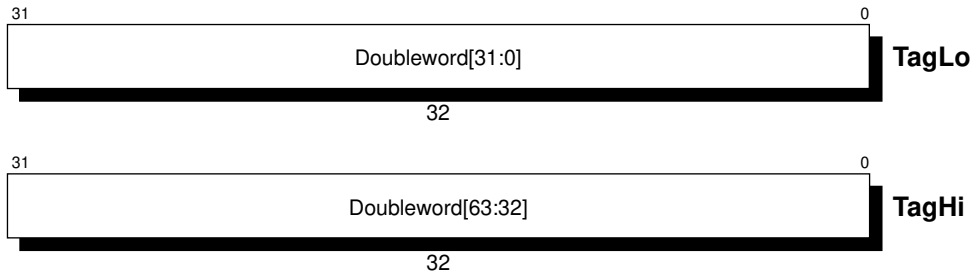

Figure 14-33 TagHi/Lo Register Fields in Secondary Cache When CacheOp is Index Load/Store Data

14.24 ErrorEPC Register (30)

The *ErrorEPC* register is similar to the *EPC* register, except that *ErrorEPC* is used on ECC and parity error exceptions. It is also used to store the program counter (PC) on Reset, Soft Reset, and nonmaskable interrupt (NMI) exceptions.

The read/write *ErrorEPC* register contains the virtual address at which instruction processing can resume after servicing an error. Figure 14-34 shows the format of the *ErrorEPC* register.

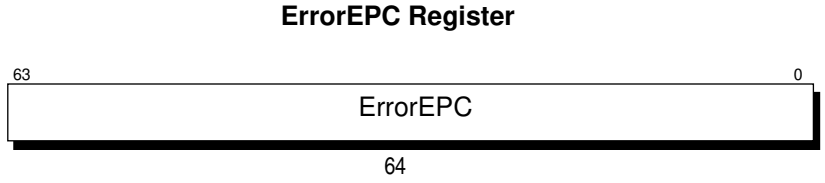

Figure 14-34 ErrorEPC Register Format

14.25 CP0 Instructions

Table 14-23 lists the CP0 instructions defined for the R10000 processor. Since they are implementation dependent, they are included here and not in the MIPS ISA manual.

Table 14-23 CP0 Instructions

OpCode	Description	ISA
CACHE	Cache Operation	III
DMFC0	Doubleword Move From CP0	III
DMTC0	Doubleword Move To CP0	III
ERET	Exception Return	III
MFC0	Move from CP0	I
MTC0	Move to CP0	I
TLBP	Probe TLB for Matching Entry	I
TLBR	Read Indexed TLB Entry	I
TLBWI	Write Indexed TLB Entry	I
TLBWR	Write Random TLB Entry	I

Hazards

The processor detects most of the pipeline hazards in hardware, including CP0 hazards and load hazards. No NOP instructions are required to correct instruction sequences.

Branch on Coprocessor 0

On the R4400 processor, CacheOps that hit in the specified cache set the *CH* bit in the Diagnostic field of the CP0 *Status* register (bit 18). Though it was undocumented, this bit could be tested by the *Branch on Coprocessor 0* instructions (*bc0t*, *bc0f*, *bc0tl*, *bc0fl*).

The R10000 processor also implements the *CH* bit but it is not associated with a Coprocessor 0 condition. Instead, execution of a branch on Coprocessor 0 instruction takes a Reserved Instruction exception.

CP0 Move Instructions

The R10000 processor implements Coprocessor 0 move instructions, MTC0, MFC0, DMTC0, and DMFC0, exactly the same as in the R4400 processor, even though some operations are undefined during certain conditions. The exact operations of CP0 move instructions on 32/64-bit CP0 registers are summarized Table 14-24.

Table 14-24 CP0 Move Instructions

Instruction	CP0 Register Size	MIPS 3 Enable?	Operation
MFC0 rt,rd	32 or 64	Don't care	rt <- $rd_{31}^{32} \| rd_{31..0}$
MTC0 rt,rd	32	Don't care	rd <- $rt_{31..0}$
	64	Don't care	rd <- $rt_{63..0}$
DMFC0 rt,rd	32	Yes	undefined (rt <- $0^{32} \| rd_{31..0}$)
	64	Yes	rt <- $rd_{63..0}$
	32 or 64	No	Reserved Instruction exception
DMTC0 rt,rd	32	Yes	undefined (rd <- $rt_{31..0}$)
	64	Yes	rd <- $rt_{63..0}$
	32 or 64	No	Reserved Instruction exception.

The returned value of MFC0/DMFC0 from a non-existing CP0 register is undefined.

CACHE Instruction

CACHE **Cache** **CACHE**

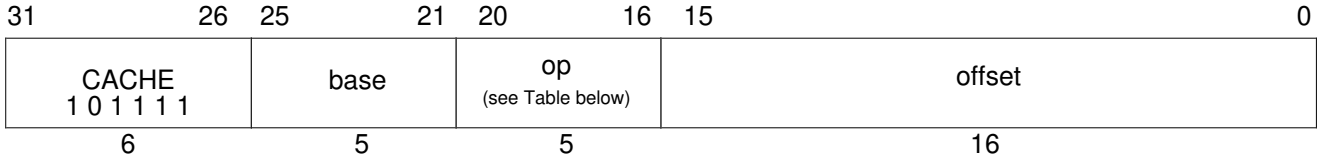

Format: CACHE op, offset(base)

Description:

The 16 bit *offset* is sign-extended and added to the contents of general register *base* to form a CacheOp virtual address (VA). The VA is translated to a physical address (PA) through the TLB, and the 5-bit opcode (decoded inTable 14-25) specifies a cache operation for that address, together with the affected cache. Operation of this instruction on any combination not listed in the tables below is undefined. The operation of this instruction on uncached addresses is also undefined.

More detailed descriptions of the CacheOps listed below are given separately, in Chapter 10.

Table 14-25 CACHE Instruction Op Field Encoding

Op Field	CACHE Instruction Variation	Target Cache
00000	Index Invalidate	(I)
00100	Index Load Tag	(I)
01000	Index Store Tag	(I)
10000	Hit Invalidate	(I)
10100	Cache Barrier	
11000	Index Load Data	(I)
11100	Index Store Data	(I)
00001	Index WriteBack Invalidate	(D)
00101	Index Load Tag	(D)
01001	Index Store Tag	(D)
10001	Hit Invalidate	(D)
10101	Hit WriteBack Invalidate	(D)
11001	Index Load Data	(D)
11101	Index Store Data	(D)
00011	Index WriteBack Invalidate	(S)
00111	Index Load Tag	(S)
01011	Index Store Tag	(S)
10011	Hit Invalidate	(S)
10111	Hit WriteBack Invalidate	(S)
11011	Index Load Data	(S)
11111	Index Store Data	(S)

CACHE Cache (continued) CACHE

Fill, *Create Dirty*, *Hit WriteBack* and *Hit Set Virtual* are not supported in the R10000 processor.

The R10000 processor adds two new CacheOps: *Index Load Data* (110_2) and *Index Store Data* (111_2). These changes are also reflected in the CP0 *TagHi*, *TagLo* and *ECC* registers.

The primary instruction and data caches have a block size of 16 words and 32 bytes (8 data words), respectively.

> **NOTE:** A 32-bit instruction is predecoded into a 36-bit instruction word before entering the primary instruction cache. The instruction fetch addresses remain the same and are not affected by the predecode.

The secondary cache, a unified cache, has a block size of either 64 or 128 bytes, configurated during reset. For a cache of $2^{CACHESIZE}$ bytes with $2^{BLOCKSIZE}$ bytes per tag,

$$VA_{CACHESIZE-2..BLOCKSIZE}$$

specifies the block for the primary cache, and

$$PA_{CACHESIZE-2..BLOCKSIZE}$$

specifies the block for the secondary cache.

For the Index CacheOps, address bit 0 is used to specify the way, 0 or 1, for the CacheOp. For this reason, bit 0 is not checked for alignment-type Address Error exception for the Index CacheOps. For CacheOps that access data in caches,

$$VA_{BLOCKSIZE-1..2}$$

specifies a word within a block for primary caches, and

$$PA_{BLOCKSIZE-1..3}$$

specifies a doubleword in the secondary cache.

A cache *hit* accesses the specified cache as normal data references, and performs the specified operation if the cache block contains valid data at the specified physical address. If the cache line is invalid or contains a differing physical address (a cache *miss*), no operation is performed. Since the R10000 processor uses 2-way set associative caches, the Hit operation performs tag comparison in both ways of the cache. No index needs to be provided for such CacheOps. If both ways register a hit, the execution of the CacheOp is undefined.

CACHE — Cache (continued) — CACHE

Write back from the primary data cache goes to the secondary cache, and write back from the secondary cache goes to the system interface. The primary data cache is written back to the secondary cache before the secondary cache is written back to the system interface; the address to be written is based by the cache tag, rather than the translated PA from the CacheOp instruction. A secondary cache write back also interrogates the primary data cache for any dirty inconsistent data.

When a line is invalidated in the secondary cache, all subset lines in the primary caches are also invalidated.

CacheOps are serialized with respect to cached loads/stores and CP0 instructions. Therefore, in general, there are no hazards for CacheOps. However, if the CacheOps modify the current instruction fetching stream, they may not work properly since the instruction fetch pipeline usually prefetches and buffers instructions and CacheOps are not serialized with respect to the instruction fetch pipeline. Programmers should be aware of such potential hazards; one solution is to put a COP0 instruction after the CacheOp to prevent the speculative execution and force the CacheOp to complete, and then use a Jump Register instruction to flush the instruction fetch pipeline. Succeeding instructions will then be re-fetched from caches.

If CP0 is not usable, a Coprocessor Unusable exception is taken. CacheOps may induce Address Error or TLBL exceptions (Refill or Invalid) during address translation, but never take a TLBS or Mod exception. The virtual address is used to index the cache for an Index CacheOp, but need not match the cache physical tag; unmapped addresses may be used to avoid TLB exceptions.

The R10000 processor does not support the *CE* bit, and programmers must supply correct parity bits or ECC for some CacheOps.

The R10000 processor supports the *CH* bit for secondary CacheOps, Hit Invalidate, and Hit WriteBack Invalidate. As in the R4400, a hit sets the *CH* bit of the *Status* register, and a miss resets it. This bit is readable and writable by software.

For a detailed description of the individual CacheOps, see Chapter 10.

Operation:

```
32, 64    T:    vAddr ← ((offset15)^48 || offset15...0) + GPR[base]
                (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
                CacheOp (op, vAddr, pAddr)
```

Exceptions:

Coprocessor unusable exception

DMFC0 Instruction

DMFC0 Doubleword Move From System Control Coprocessor DMFC0

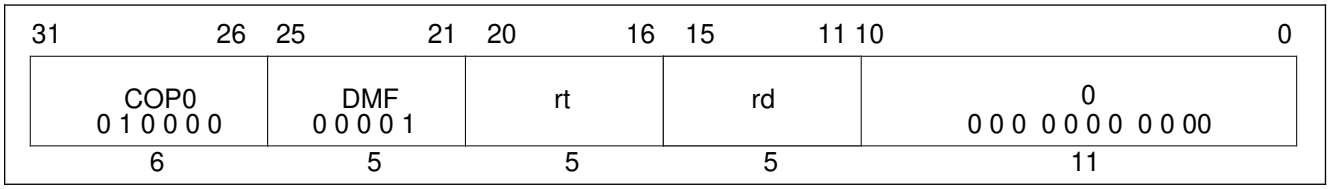

Format: DMFC0 rt, rd

Description:

The contents of coprocessor register *rd* of the CP0 are loaded into general register *rt*.

This operation is defined for the R10000 operating in 64-bit mode and in 32-bit kernel mode. Execution of this instruction in 32-bit user or supervisor mode causes a reserved instruction exception. All 64-bits of the general register destination are written from the coprocessor register source. The operation of DMFC0 on a 32-bit coprocessor 0 register is undefined.

Operation:

```
64      T:    data ← CPR[0,rd]
        T+1:  GPR[rt] ← data
```

Exceptions:

Coprocessor unusable exception

Reserved instruction exception (R10000 in 32-bit user mode
R10000 in 32-bit supervisor mode)

DMTC0 Instruction

DMTC0 — Doubleword Move To System Control Coprocessor — DMTC0

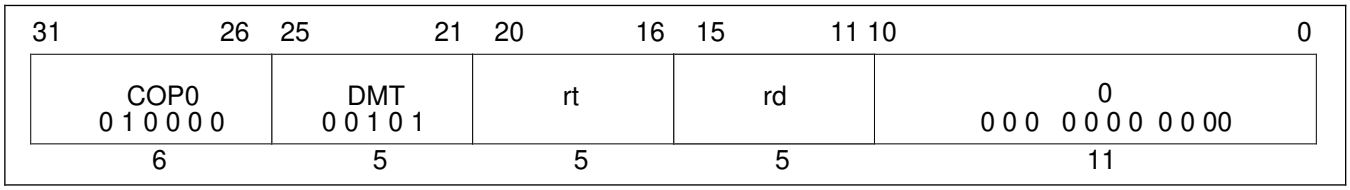

Format: DMTC0 rt, rd

Description:

The contents of general register *rt* are loaded into coprocessor register *rd* of the CP0.

This operation is defined for the R10000 operating in 64-bit mode or in 32-bit kernel mode. Execution of this instruction in 32-bit user or supervisor mode causes a reserved instruction exception.

All 64-bits of the coprocessor 0 register are written from the general register source. The operation of DMTC0 on a 32-bit coprocessor 0 register is undefined.

Because the state of the virtual address translation system may be altered by this instruction, the operation of load instructions, store instructions, and TLB operations immediately prior to and after this instruction are undefined.

Operation:

```
64      T:    data ← GPR[rt]
        T+1:  CPR[0,rd] ← data
```

Exceptions:

Coprocessor unusable exception (R10000 in 32-bit user mode
R10000 in 32-bit supervisor mode)

ERET Instruction

ERET Exception Return ERET

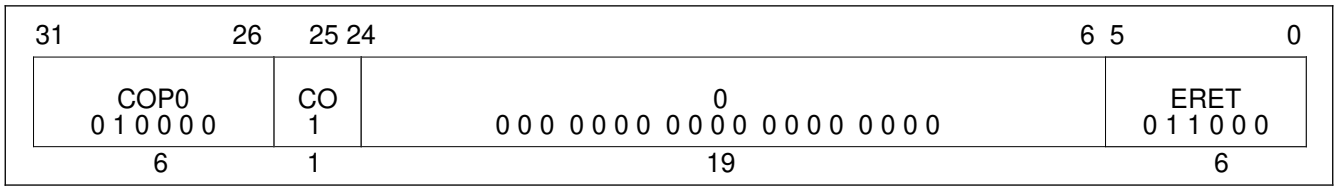

Format: ERET

Description:

ERET is the R10000 instruction for returning from an interrupt, exception, or error trap. Unlike a branch or jump instruction, ERET does not execute the next instruction.

ERET must not itself be placed in a branch delay slot.

If the processor is servicing an error trap ($SR_2 = 1$), then load the PC from the *ErrorEPC* and clear the *ERL* bit of the *Status* register (SR_2). Otherwise ($SR_2 = 0$), load the PC from the *EPC*, and clear the *EXL* bit of the *Status* register (SR_1).

An ERET executed between a LL and SC also causes the SC to fail.

If there is no exception (*EXL*=0 and *ERL*=0 in the *Status* register), execution of an ERET instruction is meaningless.

Execution of an ERET when *ERL*=0, regardless of the state of *EXL*, sets *EXL* to 0 and a jump is taken to the address presently held in the *EPC* register, even when there is no exception.

Operation:

```
32, 64      T:  if SR2 = 1 then
                    PC ← ErrorEPC
                    SR ← SR31...3 || 0 || SR1...0
                else
                    PC ← EPC
                    SR ← SR31...2 || 0 || SR0
                endif
                LLbit ← 0
```

Exceptions:

Coprocessor unusable exception

MFC0 Instruction

MFC0 — Move From System Control Coprocessor — MFC0

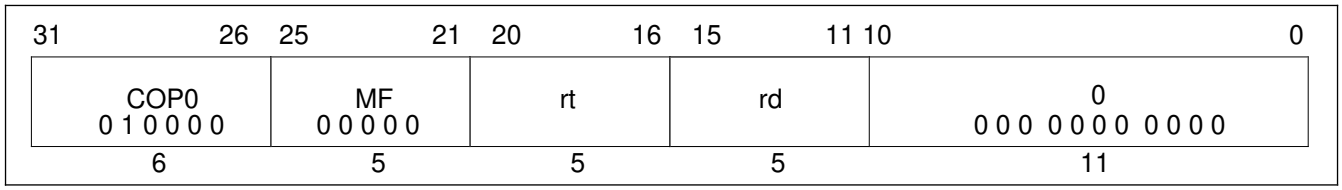

Format: MFC0 rt, rd

Description:

The contents of coprocessor register *rd* of the CP0 are loaded into general register *rt*.

Operation:

```
32    T:    data ← CPR[0,rd]
      T+1:  GPR[rt] ← data

64    T:    data ← CPR[0,rd]
      T+1:  GPR[rt] ← (data31)^32 || data31...0
```

Exceptions:

Coprocessor unusable exception

Move To/From the Performance Counter

The R10000 processor defines two performance counters, and their associated event specifier registers, which are mapped into the CP0 register 25. The following instructions are used to perform an MTC0 to or an MFC0 from a performance counter or an event specifier register.

MFPC **Move from Performance Counter** **MFPC**

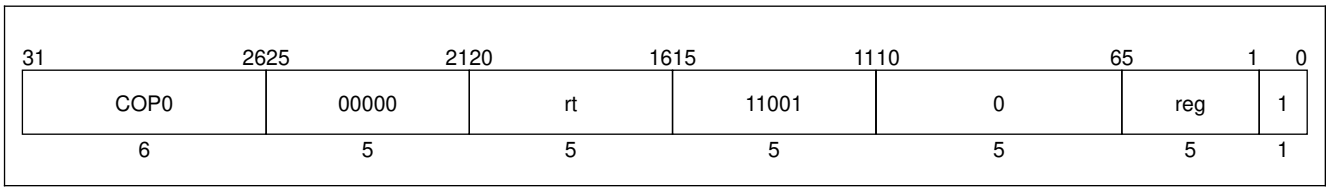

MTPC **Move to Performance Counter** **MTPC**

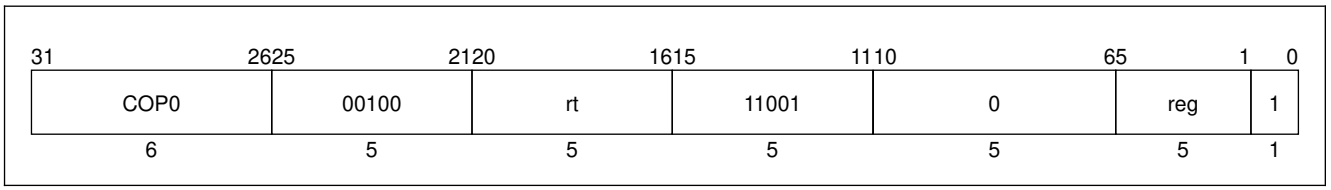

MFPS **Move from Performance Event Specifier** **MFPS**

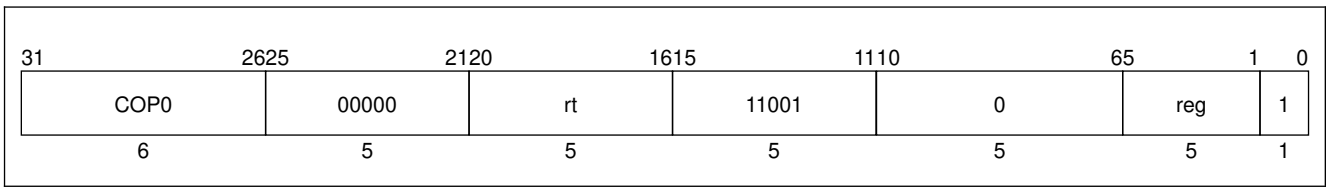

MTPS **Move to Performance Event Specifier** **MTPS**

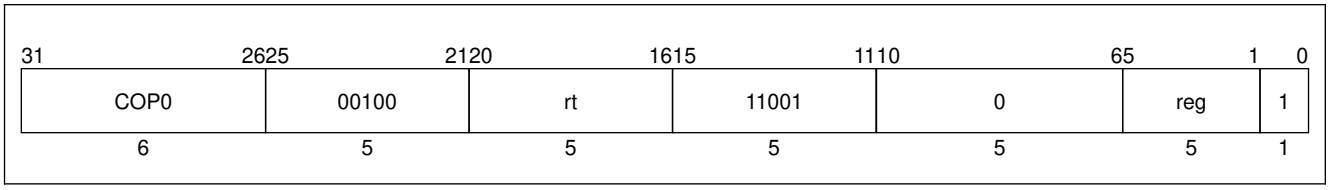

Format:

MFPC	rt, reg	— Move from performance counter
MTPC	rt, reg	— Move to performance counter
MFPS	rt, reg	— Move from performance event specifier
MTPS	rt, reg	— Move to performance event specifier

reg can be either a performance counter or an event specifier; only register 0 and 1 are valid in the R10000 implementation.

MTC0 Instruction

MTC0 — Move To System Control Coprocessor — MTC0

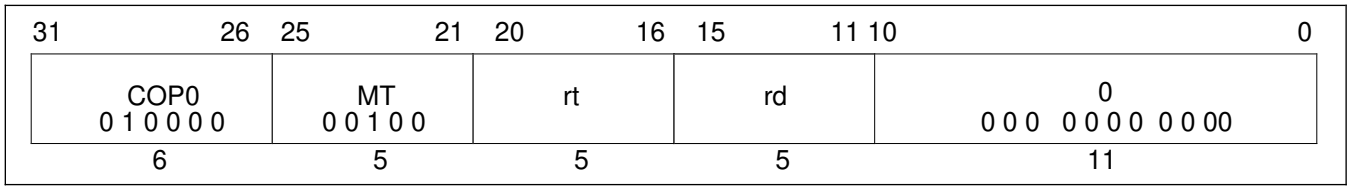

Format: MTC0 rt, rd

Description:

The contents of general register *rt* are loaded into coprocessor register *rd* of CP0.

Operation:

```
32, 64   T:      data ← GPR[rt]
         T+1:    CPR[0,rd] ← data
```

Exceptions:

Coprocessor unusable exception

TLBP Instruction

TLBP — Probe TLB For Matching Entry — TLBP

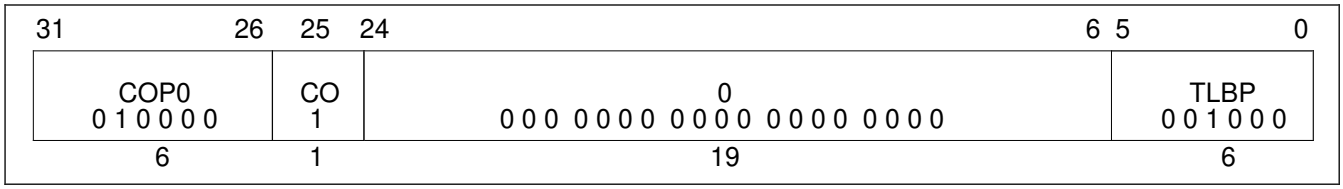

Format:

TLBP

Description:

The *Index* register is loaded with the address of the TLB entry whose contents match the contents of the *EntryHi* register. If no TLB entry matches, the high-order bit of the *Index* register is set to 0x80000000, as it is in the R4400 processor.

The architecture does not specify the operation of memory references associated with the instruction immediately after a TLBP instruction, nor is the operation specified if more than one TLB entry matches.

Operation:

```
32    T:   Index← 1 || 0^25 || undefined^6
           for i in 0...TLBEntries–1
               if (TLB[i]95...77 = EntryHi31...12) and (TLB[i]76 or
               (TLB[i]71...64 = EntryHi7...0)) then
                   Index ← 0^26 || i 5...0
               endif
           endfor
64    T:   Index← 1 || 0 ^25 || undefined^6
           for i in 0...TLBEntries–1
               if (TLB[i]171...141 and not (0^15 || TLB[i]216...205))
                = EntryHi43...13) and not (0^15 || TLB[i]216...205)) and
               (TLB[i]140 or (TLB[i]135...128 = EntryHi7...0)) then
                   Index ← 0^26 || i 5...0
               endif
           endfor
```

Exceptions:

Coprocessor unusable exception

TLBR Instruction

TLBR — Read Indexed TLB Entry — TLBR

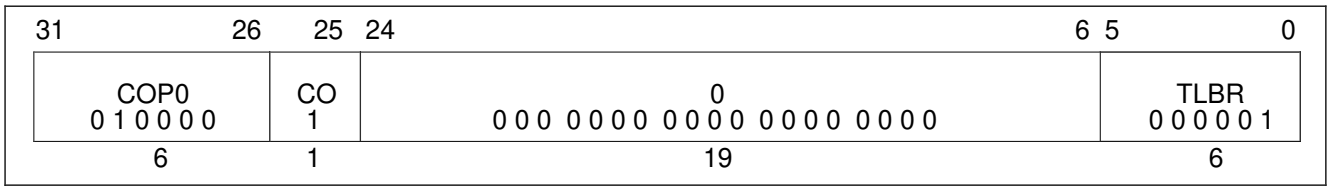

Format:

TLBR

Description:

The *G* bit (which controls ASID matching) read from the TLB is written into both of the *EntryLo0* and *EntryLo1* registers.

The *EntryHi* and *EntryLo* registers are loaded with the contents of the TLB entry pointed at by the contents of the TLB *Index* register.

In the R4400, this instruction had to be executed in unmapped spaces, and in the R10000 processor it can be executed in unmapped spaces without any hazard. In addition, TLBR can be executed in mapped spaces.

Operation:

```
32      T:  PageMask ← TLB[Index5...0]127...96
            EntryHi ← TLB[Index5...0]95...64 and not TLB[Index5...0]127...96
            EntryLo1 ←TLB[Index5...0]63...32
            EntryLo0 ← TLB[Index5...0]31...0
64      T:  PageMask ← TLB[Index5...0]255...192
            EntryHi ← TLB[Index5...0]191...128 and not TLB[Index5...0]255...192
            EntryLo1 ←TLB[Index5...0]127...65 || TLB[Index5...0]140
            EntryLo0 ← TLB[Index5...0]63...1 || TLB[Index5...0]140
```

Exceptions:

Coprocessor unusable exception

TLBWI Instruction

TLBWI Write Indexed TLB Entry TLBWI

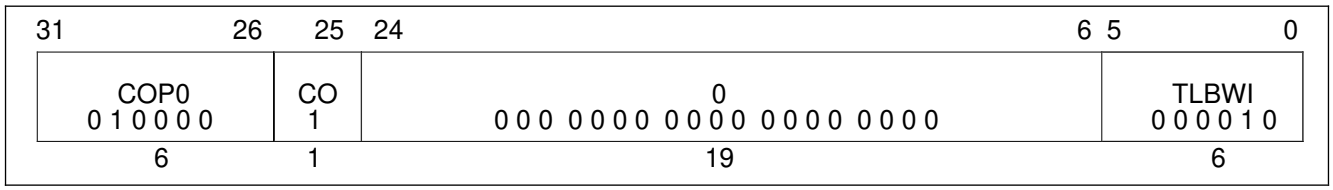

Format: TLBWI

Description:

The *G* bit of the TLB is written with the logical AND of the *G* bits in the *EntryLo0* and *EntryLo1* registers.

The TLB entry pointed at by the contents of the TLB *Index* register is loaded with the contents of the *EntryHi* and *EntryLo* registers.

The operation is invalid (and the results are unspecified) if the contents of the TLB *Index* register are greater than the number of TLB entries in the processor.

In the R4400, this instruction had to be executed in unmapped spaces, and in the R10000 processor it can be executed in unmapped spaces without any hazard.

There is no hazard to executing a TLB write in mapped space unless the write affects those instructions that have been fetched and buffered by the processor. If necessary, a flush to the instruction-fetch pipeline, such as execution of a jump register instruction, after a TLB write can avoid this hazard.

In the R4400 processor, a TLB write instruction is used to write the whole page frame number from the *EntryLo* registers to the TLB entry. Depending on the page size specified in the corresponding *PageMask* register, the lower bits of PFN may not be used for address translation. In the R10000 processor, the lower bits not used for address translation are forced to zeroes during a TLB write. This does not affect TLB address translation, however a TLB read may not retrieve what was originally written.

Operation:

```
32, 64 T:    TLB[Index5...0]  ←
                    PageMask || (EntryHi and not PageMask) || EntryLo1 || EntryLo0
```

Exceptions:

Coprocessor unusable exception

TLBWR Instruction

TLBWR Write Random TLB Entry TLBWR

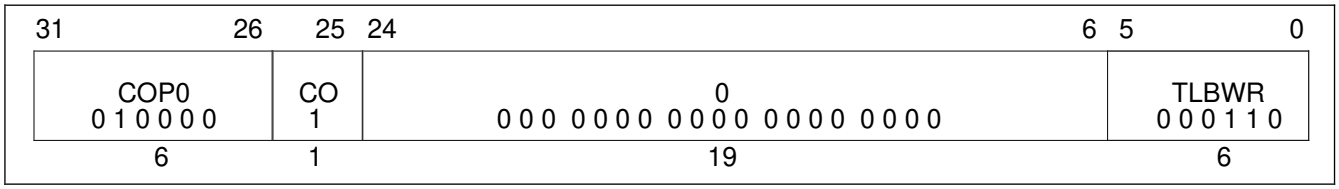

Format:

TLBWR

Description:

The G bit of the TLB is written with the logical AND of the *G* bits in the *EntryLo0* and *EntryLo1* registers.

The TLB entry pointed at by the contents of the TLB *Random* register is loaded with the contents of the *EntryHi* and *EntryLo* registers.

In the R4400, this instruction had to be executed in unmapped spaces, and in the R10000 processor it can be executed in unmapped spaces without any hazard.

There is no hazard to executing a TLB write in mapped space unless the write affects those instructions that have been fetched and buffered by the processor. If necessary, a flush to the instruction-fetch pipeline, such as execution of a jump register instruction, after a TLB write can avoid this hazard.

In the R4400 processor, a TLB write instruction is used to write the whole page frame number from the *EntryLo* registers to the TLB entry. Depending on the page size specified in the corresponding *PageMask* register, the lower bits of PFN may not be used for address translation. In the R10000 processor, the lower bits not used for address translation are forced to zeroes during a TLB write. This does not affect TLB address translation, however a TLB read may not retrieve what was originally written.

Operation:

```
32, 64T:    TLB[Random5...0]  ←
                PageMask || (EntryHi and not PageMask) || EntryLo1 || EntryLo0
```

Exceptions:

Coprocessor unusable exception

Floating-Point Unit

15

This section describes the operation of the FPU, including the register definitions.

The Floating-Point unit consists of the following functional units:

- add unit
- multiply unit
- divide unit
- square-root unit

The **add unit** performs floating-point add and subtract, compare, and conversion operations. Except for Convert Integer To Single-Precision (float), all operations have a 2-cycle latency and a 1-cycle repeat rate.

The **multiply unit** performs single-precision or double-precision multiplication with a 2-cycle latency and a 1-cycle repeat rate.

The **divide and square-root units** do single- or double-precision operations. They have long latencies and low repeat rates (20 to 40 cycles).

15.1 Floating Point Unit Operations

The floating-point add, multiply, divide, and square-root units read their operands and store their results in the floating-point register file. Values are loaded to or stored from the register file by the load/store and move units.

A logic diagram of floating-point operations is shown in Figure 15-1, in which data and instructions are read from the secondary cache into the primary caches, and then into the processor. There they are decoded and appended to the floating-point queue, passed into the FP register file where each is dynamically issued to the appropriate functional unit. After execution in the functional unit, results are stored, through the register file, in the primary data cache.

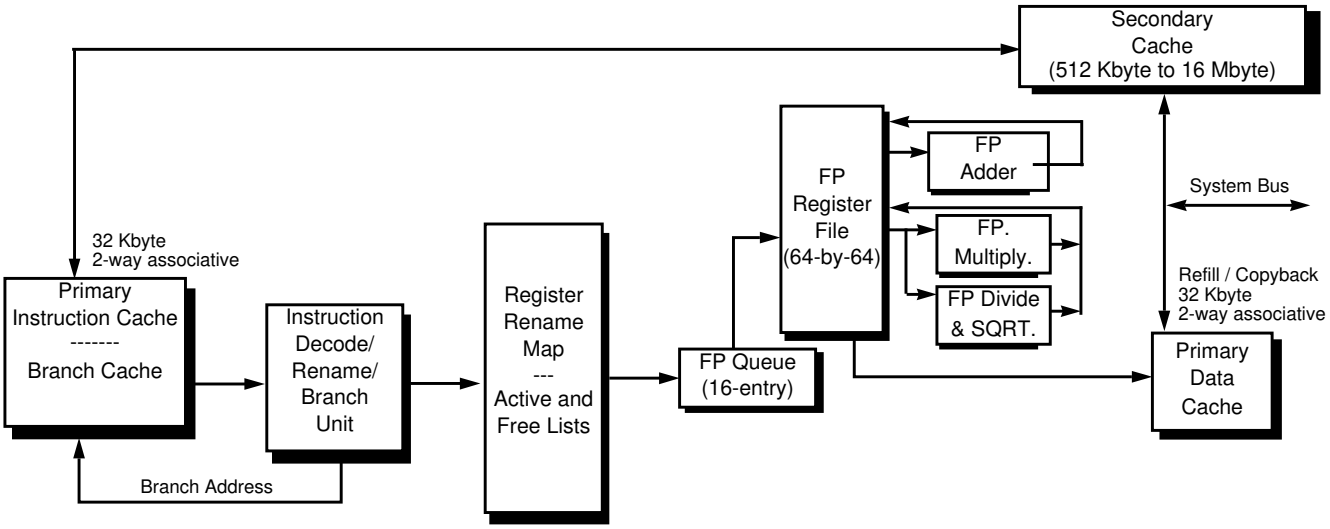

Figure 15-1 Logical Diagram of FP Operations

The floating-point queue can issue one instruction to the adder unit and one instruction to the multiplier unit. The adder and multiplier each have two dedicated read ports and a dedicated write port in the floating-point register file.

Because of their low repeat rates, the divide and square-root units do not have their own issue port. Instead, they decode instructions issued to the multiplier unit, using its operand registers and bypass logic. They appropriate a second cycle later for storing their result.

When an instruction is issued, up to two operands are read from dedicated read ports in the floating-point register file. After the operation has been completed, the result can be written back into the register file using a dedicated write port. For the add and multiply units, this write occurs four cycles after its operands were read.

15.2 Floating-Point Unit Control

The control of floating-point execution is shared by the following units:

- The floating-point queue determines operand dependencies and dynamically issues instructions to the execution units. It also controls the destination registers and register bypass.
- The execution units control the arithmetic operations and generate status.
- The graduate unit saves the status until the instructions graduate, and then it updates the *Floating-Point Status* register.

15.3 Floating-Point General Registers (FGRs)

The Floating-Point Unit is the hardware implementation of Coprocessor 1 in the MIPS IV Instruction Set Architecture. The MIPS IV ISA defines 32 logical floating-point general registers (FGRs), as shown in Figure 15-2. Each FGR is 64 bits wide and can hold either 32-bit single-precision or 64-bit double-precision values. The hardware actually contains 64 physical 64-bit registers in the Floating-Point Register File, from which the 32 logical registers are taken.

FP instructions use a 5-bit logical number to select an individual FGR. These logical numbers are mapped to physical registers by the rename unit (in pipeline stage 2), before the Floating-Point Unit executes them. Physical registers are selected using 6-bit addresses.

32- and 64-Bit Operations

The *FR* bit (26) in the *Status* register determines the number of logical floating-point registers available to the program, and it alters the operation of single-precision load/store instructions, as shown in Figure 15-2.

- *FR* is reset to 0 for compatibility with earlier MIPS I and MIPS II ISAs, and instructions use only the 16 physical even-numbered floating-point registers (32 logical registers). Each logical register is 32 bits wide.
- *FR* is set to 1 for normal MIPS III and MIPS IV operations, and all 32 of the 64-bit logical registers are available.

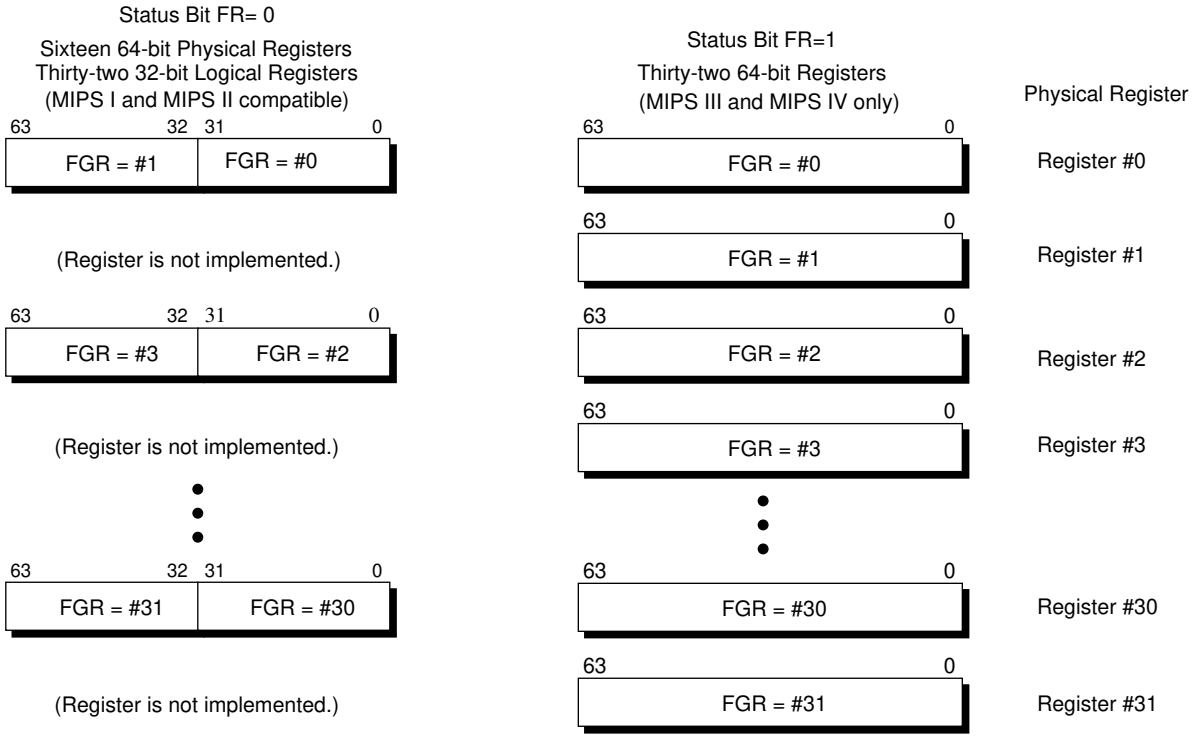

Figure 15-2 Floating-Point Registers

Load and Store Operations

When FR = 0, floating-point load and stores operate as follows:

- A doubleword load or store is handled the same as if the FR bit was set to 1, as long as the register selected is even (0, 2, 4, etc.).
- If the register selected is odd, the load/store is invalid.

These operations are shown in Figure 15-3. Singleword loads/stores to even and odd registers are also shown.

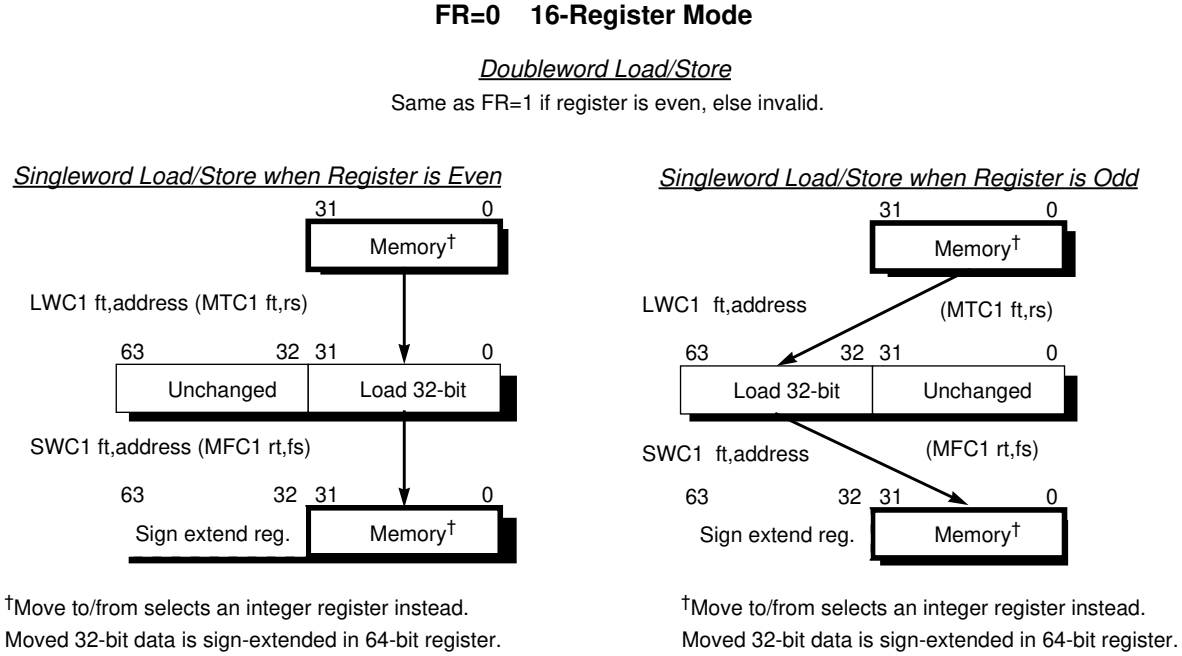

Figure 15-3 Loading and Storing Floating-Point Registers in 16-Register Mode

NOTE: Move (MOV) and conditional move (MOVC, MOVN, MOVZ are included in these arithmetic operations, although no arithmetic is actually performed.

When FR = 1, floating-point load and stores operate as follows:

- Single-precision operands are read from the low half of a register, leaving the upper half ignored. Single-precision results are written into the low half of the register. The high half of the result register is architecturally undefined; in the R10000 implementation, it is set to zero.
- Double-precision arithmetic operations use the entire 64-bit contents of each operand or result register.

Because of register renaming, every new result is written into a temporary register, and conditional move instructions select between a new operand and the previous old value. The high half of the destination register of a single-precision conditional move instruction is undefined (shown in Figure 15-5), even if no move occurs.

Singleword and doubleword loads and stores with the FPU in 32-register mode (FR=1) are shown in Figure 15-4.

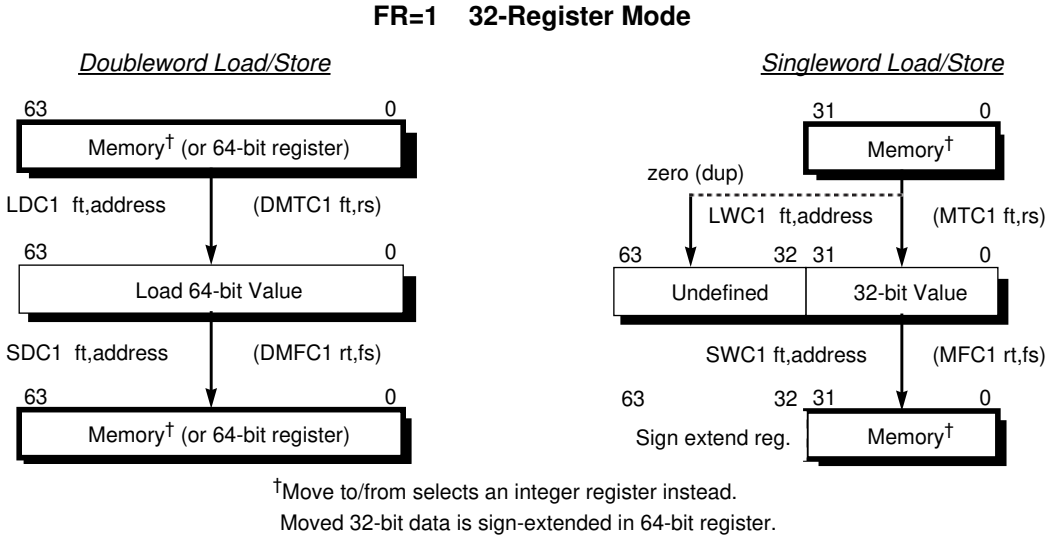

Figure 15-4 Loading and Storing Floating-Point Registers in 32-Register Mode

Doubleword load, store and move to/from instructions load or store an entire 64-bit floating-point register, as shown in Figure 15-5.

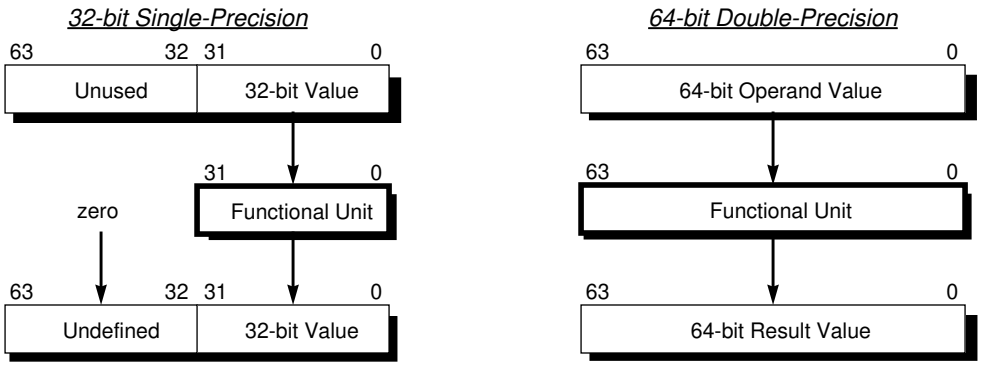

Figure 15-5 Operators on Floating-Point Registers

In MIPS I and MIPS II ISAs, all arithmetic instructions, whether single- or double-precision, are limited to using even register numbers. Load, store and move instructions transfer only a single word. Even and odd register numbers are used to access the low and high halves, respectively, of double-precision registers. When storing a floating-point register (SWC1 or MFC1), the processor reads the entire register but writes only the selected half to memory or to an integer register.

Because the register renaming scheme creates a new physical register for every destination, it is not sufficient just to enable writing half of the Floating-Point register file when loading (LWC1 or MTC1); the unchanged half must also be copied into the destination. This old value is read using the shared read port, it is then merged with the new word, and the merged doubleword value is written. (A write to the register file writes all 64 bits in parallel.)

When instructions are renamed in MIPS I or II, the low bit of any FGR field is forced to zero. Thus, each even/odd logical register number pair is treated as an even-numbered double-precision register. Odd numbered logical registers are not used in the mapping tables and dependency logic, but they remain mapped to their latest physical registers.

15.4 Floating-Point Control Registers

The MIPS IV ISA permits up to 32 control registers to be defined for each coprocessor, but the Floating-Point Unit uses only two:

- Control register 0, the FP *Implementation and Revision* register
- Control register 31, the *Floating-Point Status* register (*FSR*)

Floating-Point *Implementation and Revision* Register

The following fields are defined for control register 0 in Coprocessor 1, the FP *Implementation and Revision* register, as shown in Figure 15-6:

- The *Implementation* field holds an 8-bit number, 0x09, which identifies the R10000 implementation of the floating point coprocessor.
- The *Revision* field is an 8-bit number that defines a particular revision of the floating point coprocessor. Since it can be arbitrarily changed, it is not defined here.

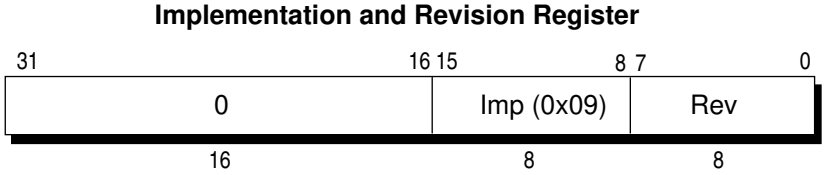

Figure 15-6 FP Implementation and Revision Register Format

Floating-Point Status Register (FSR)

Figure 15-7 shows the Floating-Point *Status* register (*FSR*), control register 31 in Coprocessor 1. It is implemented in the graduation unit rather than the Floating-Point Unit, because it is closely tied to the active list.

Bits 22:18 are unimplemented and must be set to zero. All other bits may be read or written using Control Move instructions from or to Coprocessor 1 (subfunctions CFC1 or CTC1). These move instructions are fully interlocked; they are delayed in the decode stage until all previous instructions have been graduated, and no subsequent instruction is decoded until they have been completed.

FP Status Register

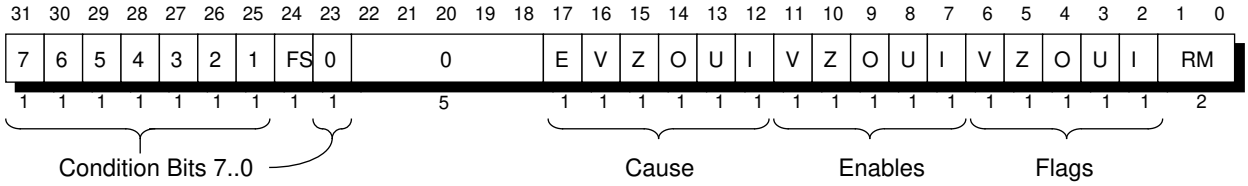

Condition bits are True/False values set by floating-point compare instructions.
Flush (FS) bit: 0: A denormalized result causes an Unimplemented Operation exception.
1: A denormalized result is replaced with zero. No exception is flagged.
Cause bits indicate the status of each floating-point arithmetic instruction. (Not by load, store, or move.)
Enable bits enable an exception if the corresponding *Cause* bit is set.
Flag bits are set whenever the corresponding *Cause* bit is a 1. These bits are cumulative. Once a bit is set, it remains set until the FSR is written by a CTC1 instruction.

E	Unimplemented operation. This exception is always enabled.
IEEE 754 Exception bits:	The following bits may be individually enabled:
V	Invalid operation.
Z	Division by zero. (Divide unit only.)
O	Overflow.
U	Underflow.
I	Inexact operation. (Result can not be stored precisely.)

Round Mode (RM): (IEEE specification)
0: *RN*, Round to nearest representable value. If two values are equally near, set the lowest bit to zero.
1: *RZ*, Round toward Zero. Round to the closest value whose magnitude is not greater than the result.
2: *RP*, Round to Plus Infinity. Round to the closest value whose magnitude is not less than the result.
3: *RM*, Round to Minus Infinity. Round to the closest value whose magnitude is not greater.

Figure 15-7 Floating-Point Status Register (FSR)

Bit Descriptions of the FSR

Description of the bits in the FSR are as follows:

Condition Bits [31:25,23]: The *Condition* bits indicate the result of floating-point compare instructions. The active list keeps track of these bits.

Cause Bits [17:12]: Each functional unit can detect exceptional cases in their function codes, operands, or results. These cases are indicated by setting one of six specific *Cause* bits. The *Cause* bits indicate the status of the floating-point arithmetic instruction which graduated most recently or caused an exception to be taken. The *FSR* is not modified by load, store, or move instructions. All cause bits, except *E*, have corresponding *Enable* and *Flag* bits in the FSR.

- E *Unimplemented operation*: the execution unit does not perform the specified operation. This exception is always enabled.
- V *Invalid operation*: this operation is not valid for the given operands.
- Z *Division by zero*: (divide unit only) the result of division by zero is not defined.
- O *Overflow*: the result is too large in magnitude to be correctly represented in the result format.
- U *Underflow*: the result is too small in magnitude to be correctly represented in the result format.
- I *Inexact Result*: the result cannot be represented exactly.

NOTE: The *FSR* is modified only for instructions issued by the floating-point queue. Move From (MFC or DMFC) instructions never set the *Cause* field; status bits from the functional unit (multiplier) must be ignored. Move or Move Conditional instructions can set the Unimplemented Operation exception only in the *Cause* field. Load and store instructions are issued by the address queue.)

The functional units generate the *Cause* bits and send them to the graduation unit when the operation is completed.

Enable Bits [11:7]: The five *Enable* bits individually enable (when set to a 1) or disable (when set to a 0) exceptions when the corresponding *Cause* bit is set.

Flag Bits [6:2]: One of the five *Flag* bits is set when a floating-point arithmetic instruction graduates, if the corresponding *Cause* bit is set. The *Flag* bits are sticky and remain set until the *FSR* is written. Thus, the *Flag* bits indicate the status of all floating-point instructions graduated since the *FSR* was last written. The *Flag* bits are not modified for any instructions which cause an exception to be taken.

Round Mode [1:0]: RM bits select one of the four IEEE rounding modes. Most floating-point results cannot be precisely represented by the 32-bit or 64-bit register formats, and must be truncated and rounded to a representable value. The modes selected by the *RM* bit values are:

0: *RN*, round to nearest representable value. If two values are equally near, set the lowest bit to zero.

1: *RZ*, round toward zero. Round to the closest value whose magnitude is not greater than the result.

2: *RP*, round to plus infinity. Round to the closest value whose magnitude is not less than the result.

3: *RM*, round to minus infinity. Round to the closest value whose magnitude is not greater.

The *Round* and *Enable* bits only change when the *FSR* is written by a CTC1 (Move To Coprocessor 1 Control Register) instruction. Each CTC1 instruction is executed sequentially, after all previous floating-point instructions have been completed, so these *FSR* bits do not change while any floating-point instruction is active. These bits are broadcast from the graduation unit to all the floating-point functional units.

When a *Cause* bit is set and its corresponding *Enable* bit is also set, an exception is taken on the instruction. The result of the instruction is not stored, and the *Flag* bits are not changed. If no exception is taken, the corresponding *Flag* bits are set.

The *Cause* and *Flag* bits may be read or written. If a CTC1 instruction sets both a *Cause* bit and its *Enable* bit, an exception is taken immediately. The *FSR* is written, but the exception is reported on the move instruction.

Loading the FSR

The *FSR* may be loaded from an integer register by a CTC1 instruction which selects control register 31. This instruction is executed serially; that is, it is delayed during decode until the entire pipeline has emptied, and it is completed before the next instruction is decoded. This instruction writes all *FSR* bits.

If any *Cause* bit and its corresponding *Enable* bit are both set, an exception is taken after FSR has been modified. The CTC1 instruction is aborted; it does not graduate, even though it has changed the processor state.

15.5 FPU Instructions

This section describes the R10000 processor-specific implementations of the following FPU instructions:

- CVT.L.fmt
- moves and conditional moves
- CFC1/CTC1CVT.L.fmt

CVT.L.fmt

The CVT.L.fmt instruction has a slight change in the R10000 processor implementation. The R4400 processor allows conversion from a single or a double up to a 53-bit long integer. If the result is greater than 53 bits after the conversion, an Unimplemented Operation exception is taken. A back- conversion from a 53-bit long integer to single/double also takes an Unimplemented Operation exception.

In the R10000, the conversion allows only up to 52 bits; otherwise an Unimplemented Operation exception is taken. The back-conversion from a 52-bit long integer to single/double no longer takes an Unimplemented Operation exception.

Moves and Conditional Moves

The only legal formats for the move and conditional move instructions are single and double precision. The move instructions do not trap if their operands are either denormalized or NaNs, which is consistent with the R4400 implementation. Execution of floating-point move and conditional move instructions do not affect the *Cause* field of the floating-point *Status* register unless they take an Unimplemented Operation exception because an illegal formats was used.†

The upper 32 bits of the destination registers are undefined in architecture for all the floating-point arithmetic operations in single-precision or 32-bit fixed format (S or W). In the R10000 processor, the implementation clears the upper 32 bits, including MOV.S, whereas R4400 and R4200 processors preserve the upper 32 bits during the move.

For the floating-point conditional move instructions, MOVT.S, MOVF.S, MOVZ.S, and MOVN.S, the R10000 processor always clears the upper 32 bits of the destination register even though the condition is false.

In *32 floating-point register* mode (FR=1), the upper 32 bits of the destination register for the MTC1 and LWC1 instructions are architecturally undefined. The R10000 processor implementation clears the upper 32 bits.

CFC1/CTC1

There are only two valid Floating-Point Control registers: 0 and 31. Access to other registers is undefined.

† The *Cause* field is set to 100000 (*E* bit is 1).

Memory Management

16

This section describes the R10000 processor memory management, including:

- processor modes and exceptions
- virtual address space
- virtual address translation

16.1 Processor Modes

The R10000 has three operating modes and two addressing modes. All are described in this section.

Processor Operating Modes

The three operating modes are listed in order of decreasing system privilege:

- **Kernel mode** (highest system privilege): can access and change any register. The innermost core of the operating system runs in kernel mode.
- **Supervisor mode**: has fewer privileges and is used for less critical sections of the operating system.
- **User mode** (lowest system privilege): prevents users from interfering with one another.

Selection between the three modes can be made by the operating system (when in Kernal mode) by writing into *Status* register's *KSU* field. The processor is forced into Kernel mode when the processor is handling an error (the *ERL* bit is set) or an exception (the *EXL* bit is set). Table 16-1 shows the selection of operating modes with respect to the *KSU*, *EXL* and *ERL* bits.

Table 16-1 also shows how different instruction sets and addressing modes are enabled by the *Status* register's *XX*, *UX*, *SX* and *KX* bits. A dash ("-") in this table indicates a "don't care." For detailed information on the address spaces available in each mode, refer to section titled, Virtual Address Space, in this chapter.

The R10000 processor was designed for use with the MIPS IV ISA; however, for compatibility with earlier machines, the useable ISAs can be limited to either MIPS III or MIPSI/II.

Table 16-1 Processor Modes

XX 31	KX 7	SX 6	UX 5	KSU 4:3	ERL 2	EXL 1	Description	ISA* III	ISA* IV	Addressing Mode 32-Bit/64-Bit
0	-†	-	0	10	0	0	User mode.	No	No	32
1	-	-	0	10	0	0		No	Yes	32
0	-	-	1	10	0	0		Yes	No	64
1	-	-	1	10	0	0		Yes	Yes	64
-	-	0	-	01	0	0	Supervisor mode.	No	Yes	32
-	-	1	-	01	0	0		Yes	Yes	64
-	0	-	-	00	0	0	Kernel mode.	Yes	Yes	32
-	1	-	-	00	0	0		Yes	Yes	64
-	0	-	-	-	0	1	Exception Level	Yes	Yes	32
-	1	-	-	-	0	1		Yes	Yes	64
-	0	-	-	-	1	X	Error Level.	Yes	Yes	32
-	1	-	-	-	1	X		Yes	Yes	64

* No means the ISA is disabled; Yes means the ISA is enabled.

† Dashes (-) are "don't care."

Addressing Modes

The processor's *addressing mode* determines whether it generates 32-bit or 64-bit memory addresses.

Refer to Table 16-1 for the following addressing mode encodings:

- In Kernel mode the *KX* bit allows 64-bit addressing; all instructions are always valid.
- In Supervisor mode, the *SX* bit allows 64-bit addressing and the MIPS III instructions. MIPS IV ISA is enabled all the time in Supervisor mode.
- In User mode, the *UX* bit allows 64-bit addressing and the MIPS III instructions; the *XX* bit allows the new MIPS IV instructions.

16.2 Virtual Address Space

The processor uses either 32-bit or 64-bit address spaces, depending on the operating and addressing modes set by the *Status* register. Table 16-1 lists the decoding of these modes.

The processor uses the following addresses:

- virtual address **VA[43:0]**
- region bits **VA[63:59]**

If a region is **mapped**, virtual addresses are translated in the TLB. Bits **VA[58:44]** are not translated in the TLB and are sign extensions of bit **VA[43]**.

In both 32-bit and 64-bit address mode, the memory address space is divided into many regions, as shown in Figure 16-3. Each region has specific characteristics and uses. The user can access only the *useg* region in 32-bit mode, or *xuseg* in 64-bit mode, as shown in Figure 16-1. The supervisor can access user regions as well as *sseg* (in 32-bit mode) or *xsseg* and *csseg* (in 64-bit mode), shown in Figure 16-2. The kernel can access all regions except those restricted because bits **VA[58:44]** are not implemented in the TLB, as shown in Figure 16-3.

The R10000 processor follows the R4400 implementation for *data* references only, ensuring compatibility with the NT kernel. If any of the upper 33 bits are nonzero for an instruction fetch, an Address Error is generated. Refer to Table 16-2 for delineation of the address spaces.

User Mode Operations

In User mode, a single, uniform virtual address space—labelled User segment—is available; its size is:

- 2 Gbytes (2^{31} bytes) in 32-bit mode (*useg*)
- 16 Tbytes (2^{44} bytes) in 64-bit mode (*xuseg*)

Figure 16-1 shows User mode virtual address space.

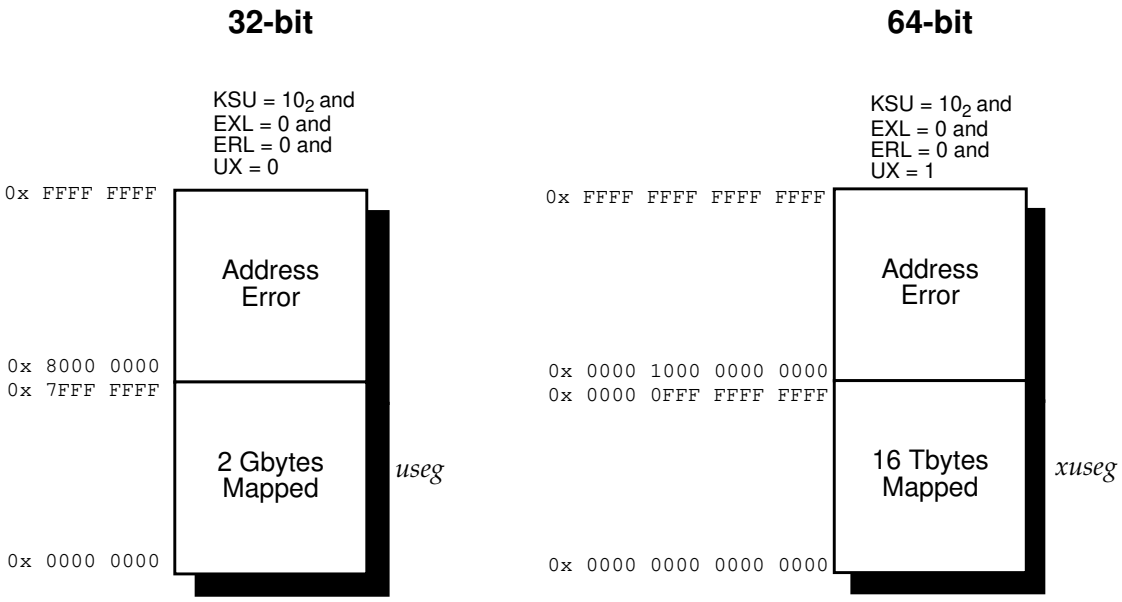

Figure 16-1 User Mode Virtual Address Space

The User segment starts at address 0 and the current active user process resides in either *useg* (in 32-bit mode) or *xuseg* (in 64-bit mode). The TLB identically maps all references to *useg*/*xuseg* from all modes, and controls cache accessibility.

32-bit User Mode (*useg*)

In User mode, when *UX* = 0 in the *Status* register, User mode addressing is compatible with the 32-bit addressing model shown in Figure 16-1, and a 2-Gbyte user address space is available, labelled *useg*.

All valid User mode virtual addresses have their most-significant bit cleared to 0; any attempt to reference an address with the most-significant bit set while in User mode causes an Address Error exception.

The system maps all references to *useg* through the TLB, and bit settings within the TLB entry for the page determine the cacheability of a reference.

64-bit User Mode (*xuseg*)

In User mode, when *UX* =1 in the *Status* register, User mode addressing is extended to the 64-bit model shown in Figure 16-1. In 64-bit User mode, the processor provides a single, uniform virtual address space of 2^{44} bytes, labelled *xuseg*.

All valid User mode virtual addresses have bits 63:44 equal to 0; an attempt to reference an address with bits 63:44 not equal to 0 causes an Address Error exception.

Although the system may be in 32-bit mode, address logic still generates 64-bit values. In this case the high 32 bits must equal the sign bit (31), or an Address Error exception is taken.

Supervisor Mode Operations

Supervisor mode is designed for layered operating systems in which a true kernel runs in processor Kernel mode, and the rest of the operating system runs in Supervisor mode.

The processor operates in Supervisor mode when the *Status* register contains the Supervisor-mode bit-values shown in Table 16-1.

Figure 16-2 shows Supervisor mode address mapping.

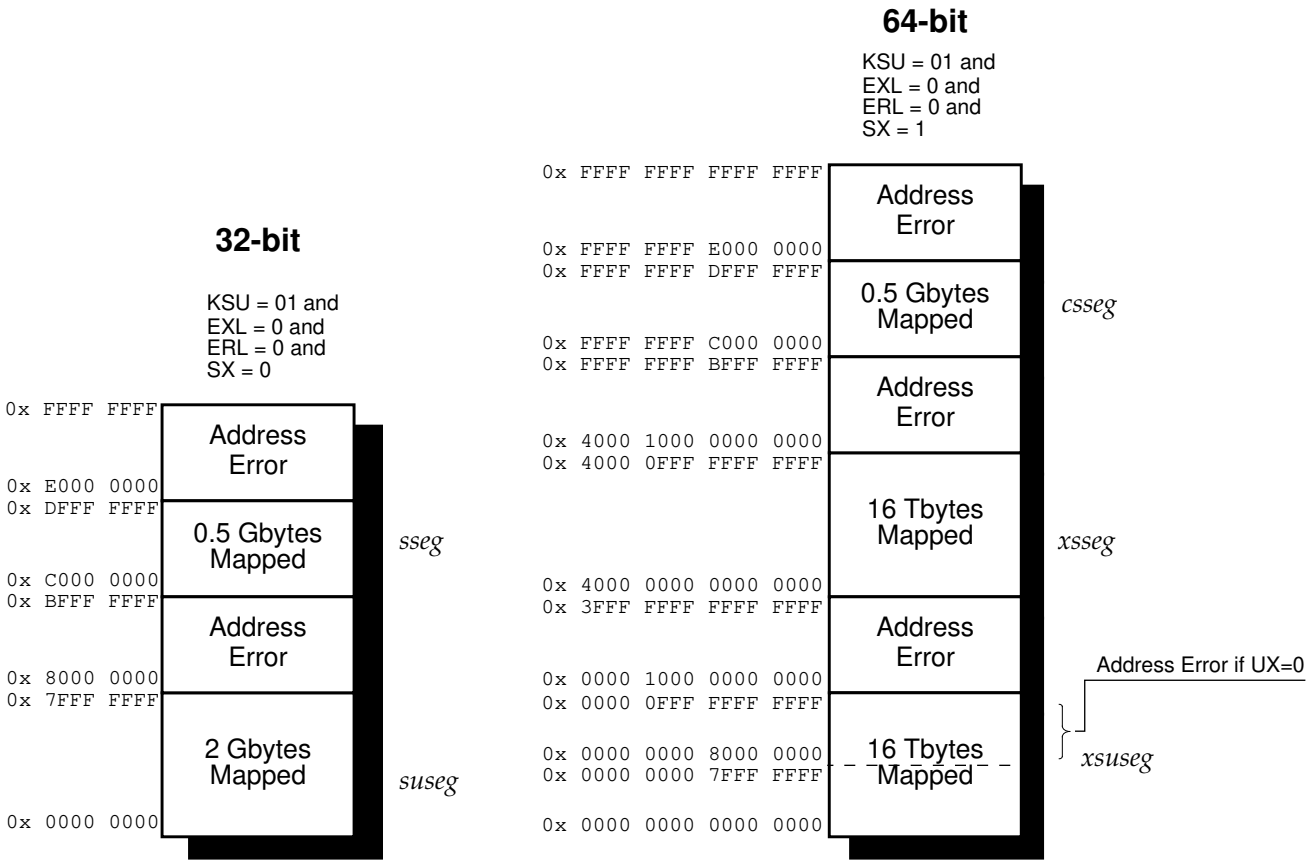

Figure 16-2 Supervisor Mode Address Space

32-bit Supervisor Mode, User Space (*suseg*)

In Supervisor mode, when *SX* = 0 in the *Status* register and the most-significant bit of the 32-bit virtual address is set to 0, the *suseg* virtual address space is selected; it covers the full 2^{31} bytes (2 Gbytes) of the current user address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

This mapped space starts at virtual address 0x0000 0000 and runs through 0x7FFF FFFF.

32-bit Supervisor Mode, Supervisor Space (*sseg*)

In Supervisor mode, when *SX* = 0 in the *Status* register and the three most-significant bits of the 32-bit virtual address are 110_2, the *sseg* virtual address space is selected; it covers 2^{29}-bytes (512 Mbytes) of the current supervisor address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

This mapped space begins at virtual address 0xC000 0000 and runs through 0xDFFF FFFF.

64-bit Supervisor Mode, User Space (*xsuseg*)

In Supervisor mode, when *SX* = 1 in the *Status* register and bits 63:62 of the virtual address are set to 00_2, selection of the *xsuseg* virtual address space is dependent upon the *UX* bit.

- if *UX* = 1, the entire space from 0x0000 0000 0000 0000 through 0000 0FFF FFFF FFFF (16 Tbytes) is selected.
- If *UX* = 0, the address space 0x0000 0000 0000 0000 through 0000 0000 7FFF FFFF (2 Gbytes) is selected. Addressing the space ranging from 0000 0000 8000 0000 through 0000 0FFF FFFF FFFF will cause an address error.

The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

64-bit Supervisor Mode, Current Supervisor Space (*xsseg*)

In Supervisor mode, when *SX* = 1 in the *Status* register and bits 63:62 of the virtual address are set to 01_2, the *xsseg* current supervisor virtual address space is selected. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

This mapped space begins at virtual address 0x4000 0000 0000 0000 and runs through 0x4000 0FFF FFFF FFFF.

64-bit Supervisor Mode, Separate Supervisor Space (*csseg*)

In Supervisor mode, when *SX* = 1 in the *Status* register and bits 63:62 of the virtual address are set to 11_2, the *csseg* separate supervisor virtual address space is selected. Addressing of the *csseg* is compatible with addressing *sseg* in 32-bit mode. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

This mapped space begins at virtual address 0xFFFF FFFF C000 0000 and runs through 0xFFFF FFFF DFFF FFFF.

Kernel Mode Operations

The processor operates in Kernel mode when the *Status* register contains the Kernel-mode bit-values shown in Table 16-1.

Kernel mode virtual address space is divided into regions differentiated by the high-order bits of the virtual address, as shown in Figure 16-3.

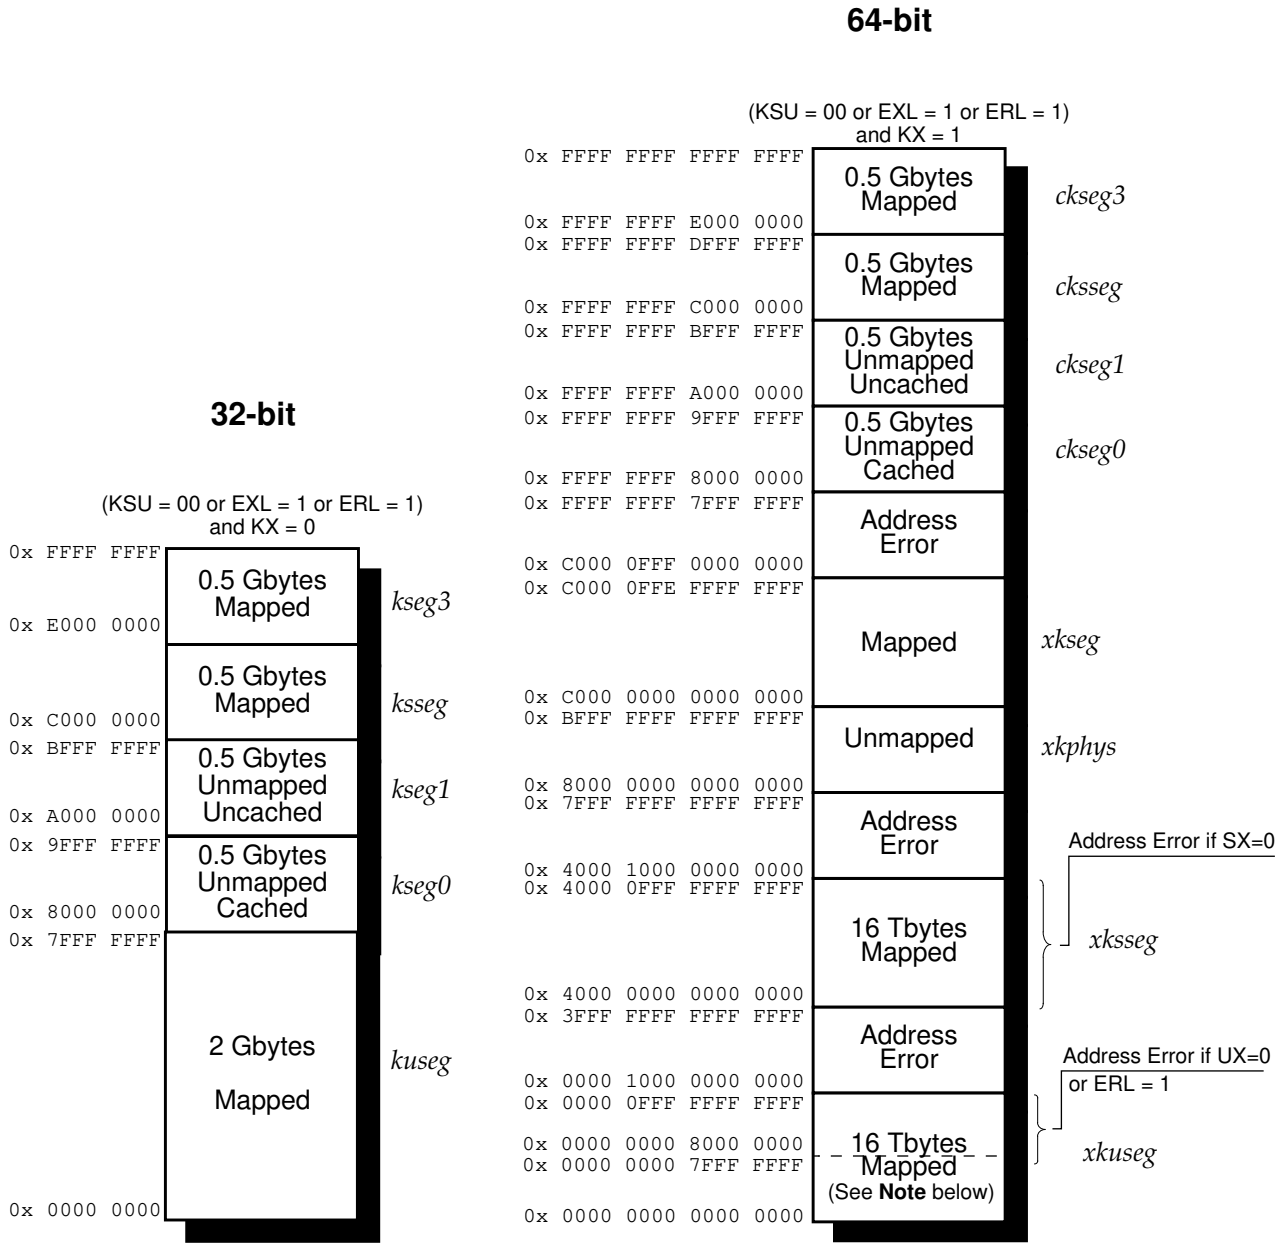

Figure 16-3 Kernel Mode Address Space

NOTE: If *ERL* = 1, the selected 2 Gbyte space becomes uncached and unmapped.

32-bit Kernel Mode, User Space (*kuseg*)

In Kernel mode, when *KX* = 0 in the *Status* register, and the most-significant bit of the virtual address, A31, is cleared, the 32-bit *kuseg* virtual address space is selected; it covers the full 2^{31} bytes (2 Gbytes) of the current user address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

32-bit Kernel Mode, Kernel Space 0 (*kseg0*)

In Kernel mode, when *KX* = 0 in the *Status* register and the most-significant three bits of the virtual address are 100_2, 32-bit *kseg0* virtual address space is selected; it is the 2^{29}-byte (512-Mbyte) kernel physical space. References to *kseg0* are not mapped through the TLB; the physical address is selected by subtracting 0x8000 0000 from the virtual address. The *K0* field of the *Config* register determines cacheability and coherency.

32-bit Kernel Mode, Kernel Space 1 (*kseg1*)

In Kernel mode, when *KX* = 0 in the *Status* register and the most-significant three bits of the 32-bit virtual address are 101_2, 32-bit *kseg1* virtual address space is selected; it is the 2^{29}-byte (512-Mbyte) kernel physical space.

References to *kseg1* are not mapped through the TLB; the physical address is selected by subtracting 0xA000 0000 from the virtual address.

Caches are disabled for accesses to these addresses, and physical memory (or memory-mapped I/O device registers) are accessed directly.

32-bit Kernel Mode, Supervisor Space (*ksseg*)

In Kernel mode, when *KX* = 0 in the *Status* register and the most-significant three bits of the 32-bit virtual address are 110_2, the *ksseg* virtual address space is selected; it is the current 2^{29}-byte (512-Mbyte) supervisor virtual space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

References to *ksseg* are mapped through the TLB.

32-bit Kernel Mode, Kernel Space 3 (*kseg3*)

In Kernel mode, when *KX* = 0 in the *Status* register and the most-significant three bits of the 32-bit virtual address are 111_2, the *kseg3* virtual address space is selected; it is the current 2^{29}-byte (512-Mbyte) kernel virtual space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

References to *kseg3* are mapped through the TLB.

64-bit Kernel Mode, User Space (*xkuseg*)

In Kernel mode, when *KX* = 1 in the *Status* register and bits 63:62 of the 64-bit virtual address are 00_2, selection of the *xkuseg* virtual address space is dependent upon the *UX* and *ERL* bits.

- if *UX* = 1 and *ERL* = 0, the entire space from 0x0000 0000 0000 0000 through 0000 0FFF FFFF FFFF (16 Tbytes) is selected.
- If *UX* = 0 or *ERL* = 1, the address space 0x0000 0000 0000 0000 through 0000 0000 7FFF FFFF (2 Gbytes) is selected. Addressing the space ranging from 0000 0000 8000 0000 through 0000 0FFF FFFF FFFF will cause an address error. Moreover, if *ERL*=1, the selected 2-Gbyte address space becomes unmapped and uncached.

The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

64-bit Kernel Mode, Current Supervisor Space (*xksseg*)

In Kernel mode, when *KX* = 1 in the *Status* register and bits 63:62 of the 64-bit virtual address are 01_2, selection of the *xksseg* virtual address space is dependent upon the *SX* bit.

- if *SX* = 1, the entire space from 0x4000 0000 0000 0000 through 4000 0FFF FFFF FFFF (16 Tbytes) is selected.
- If *SX* = 0, access to any address in the space ranging from 0x4000 0000 0000 0000 through 4000 0FFF FFFF FFFF causes an address error.

The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

64-bit Kernel Mode, Physical Spaces (*xkphys*)

In Kernel mode, when *KX* = 1 in the *Status* register and bits 63:62 of the 64-bit virtual address are 10_2, the *xkphys* virtual address space is selected; it is a set of eight kernel physical spaces. Each kernel physical space contains either one or four 2^{40}-byte physical pages.

References to this space are not mapped; the physical address selected is taken directly from bits 39:0 of the virtual address. Bits 61:59 of the virtual address specify the *cache algorithm*, described in Chapter 4. If the cache algorithm is either uncached or uncached accelerated (values of 2 or 7) the space contains four physical pages; access to addresses whose bits 56:40 are not equal to 0 cause an Address Error exception. Address bits 58:57 carry the *uncached attribute* (described in Chapter 6), and are not checked for address errors.

If the cache algorithm is neither uncached nor uncached accelerated, the space contains a single physical page, as on the R4400 processor. In this case, access to addresses whose bits 58:40 are not equal to a zero cause an Address Error exception, as shown in Figure 16-4.

0X BFFFFFFF FFFFFFFF 0X BE000100 00000000	Address Error
0X BE0000FF FFFFFFFF 0X BE000000 00000000	Uncached Accelerated
0X BDFFFFFF FFFFFFFF 0X BC000100 00000000	Address Error
0X BC0000FF FFFFFFFF 0X BC000000 00000000	Uncached Accelerated
0X BBFFFFFF FFFFFFFF 0X BA000100 00000000	Address Error
0X BA0000FF FFFFFFFF 0X BA000000 00000000	Uncached Accelerated
0X B9FFFFFF FFFFFFFF 0X B8000100 00000000	Address Error
0X B80000FF FFFFFFFF 0X B8000000 00000000	Uncached Accelerated
0X B7FFFFFF FFFFFFFF 0X B0000100 00000000	Address Error
0X B00000FF FFFFFFFF 0X B0000000 00000000	Reserved*
0X AFFFFFFF FFFFFFFF 0X A8000100 00000000	Address Error
0X A80000FF FFFFFFFF 0X A8000000 00000000	Cacheable Exclusive Write
0X A7FFFFFF FFFFFFFF 0X A0000100 00000000	Address Error
0X A00000FF FFFFFFFF 0X A0000000 00000000	Cacheable Exclusive

0X 9FFFFFFF FFFFFFFF 0X 98000100 00000000	Address Error
0X 980000FF FFFFFFFF 0X 98000000 00000000	Cacheable Noncoherent
0X 97FFFFFF FFFFFFFF 0X 96000100 00000000	Address Error
0X 960000FF FFFFFFFF 0X 96000000 00000000	Uncached
0X 95FFFFFF FFFFFFFF 0X 94000100 00000000	Address Error
0X 940000FF FFFFFFFF 0X 94000000 00000000	Uncached
0X 93FFFFFF FFFFFFFF 0X 92000100 00000000	Address Error
0X 920000FF FFFFFFFF 0X 92000000 00000000	Uncached
0X 91FFFFFF FFFFFFFF 0X 90000100 00000000	Address Error
0X 900000FF FFFFFFFF 0X 90000000 00000000	Uncached
0X 8FFFFFFF FFFFFFFF 0X 88000100 00000000	Address Error
0X 880000FF FFFFFFFF 0X 88000000 00000000	Reserved*
0X 87FFFFFF FFFFFFFF 0X 80000100 00000000	Address Error
0X 800000FF FFFFFFFF 0X 80000000 00000000	Reserved*

* Accessing a reserved space results in undefined behavior.

Figure 16-4 xkphys Virtual Address Space

64-bit Kernel Mode, Kernel Space (*xkseg*)

In Kernel mode, when *KX* = 1 in the *Status* register and bits 63:62 of the 64-bit virtual address are 11_2, the address space selected is one of the following:

- kernel virtual space, *xkseg*, the current kernel virtual space; the virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address
- one of the four 32-bit kernel mode compatibility spaces (described below).

64-bit Kernel Mode, Compatibility Spaces (*ckseg1:0, cksseg, ckseg3*)

In Kernel mode, when *KX* = 1 in the *Status* register, bits 63:62 of the 64-bit virtual address are 11_2, and bits 61:31 of the virtual address equal –1, the lower two bytes of address, as shown in Figure 16-3, select one of the following 512-Mbyte compatibility spaces.

- *ckseg0*. This 64-bit virtual address space is an unmapped region, compatible with the 32-bit address model *kseg0*. The *K0* field of the *Config* register controls cacheability and coherency.
- *ckseg1*. This 64-bit virtual address space is an unmapped and uncached region, compatible with the 32-bit address model *kseg1*.
- *cksseg*. This 64-bit virtual address space is the current supervisor virtual space, compatible with the 32-bit address model *kssе*g.
- *ckseg3*. This 64-bit virtual address space is kernel virtual space, compatible with the 32-bit address model *kseg3*.

Address Space Access Privilege Differences Between the R4400 and R10000

In the R4400, the 64-bit Supervisor mode can access the entire *xsuseg* space, and the 64-bit Kernel mode can access the entire *xksseg* and *xkuseg* spaces. Access privileges in the R10000 are also dependent on the *UX* and *SX* bits:

- Access to the 64-bit user space in 64-bit Supervisor or Kernel mode (*xsuseg* or *xkuseg*) is controlled by the *UX* bit. If *UX*=0, the 64-bit Supervisor and Kernel modes can only access the 32-bit user space (*suseg* or *kuseg*).
- Access to the 64-bit supervisor space in Kernel mode (*xksseg*) is controlled by the *SX* bit. If *SX*=0, the 64-bit Kernel mode can only access the 32-bit supervisor space (*ksseg*).

An Address Error exception is taken on an illegal access.

The R10000 processor implements the same access privileges for 32-bit processor modes as in the R4400. The Table 16-2 summarizes the access privileges for all processor modes in the R10000 processor.

Table 16-2 Access Privileges for User, Supervisor and Kernel Mode Operations

64-bit Virtual Address	32-bit Mode			64-bit Mode			
	User*	Supervisor	Kernel	User	Supervisor	Kernel & ERL=0	Kernel & ERL=1
FFFFFFFF E0000000 TO FFFFFFFF FFFFFFFF		*AddrErr*	OK		*AddrErr*	OK	OK
FFFFFFFF C0000000 TO FFFFFFFF DFFFFFFF		OK			OK		
FFFFFFFF A0000000 TO FFFFFFFF BFFFFFFF							
FFFFFFFF 80000000 TO FFFFFFFF 9FFFFFFF							
C0000FFF 00000000 TO FFFFFFFF 7FFFFFFF						*AddrErr*	*AddrErr*
C0000000 00000000 TO C0000FFE FFFFFFFF	*AddrErr*			*AddrErr*	*AddrErr*	OK	OK
80000000 00000000 TO BFFFFFFF FFFFFFFF		*AddrErr*				OK	OK
40001000 00000000 TO 7FFFFFFF FFFFFFFF			*AddrErr*			*AddrErr*	*AddrErr*
40000000 00000000 TO 40000FFF FFFFFFFF					OK	*AddrErr if SX=0*	*AddrErr if SX=0*
00001000 00000000 TO 3FFFFFFF FFFFFFFF					*AddrErr*	*AddrErr*	*AddrErr*
00000000 80000000 TO 00000FFF FFFFFFFF				OK	*AddrErr if UX=0*	*AddrErr if UX=0*	*AddrErr*
00000000 00000000 TO 00000000 7FFFFFFF	OK	OK	OK		OK	OK	OK

* For data references, the upper 32 bits of the virtual addresses are cleared before checking access privilege and TLB translation.

16.3 Virtual Address Translation

Programs can operate using either **physical** or **virtual** memory addresses:

- physical addresses correspond to hardware locations in main memory
- virtual addresses are logical values only, and do not correspond to fixed hardware locations

Virtual addresses must first be **translated** (finding the physical address at which the virtual address points) before main memory can be accessed. This translation is essential for multitasking computer systems, because it allows the operating system to load programs anywhere in main memory independent of the logical addresses used by the programs.

This translation also implements a memory protection scheme, which limits the amount of memory each program may access. The scheme prevents programs from interfering with the memory used by other programs or the operating system.

Cache Pages

Translated virtual addresses retrieve data in blocks, which are called **pages**. In the R10000 processor, the size of each page may be selected from a range that runs from 4 Kbytes to 16 Mbytes inclusive, in increments of 4 (that is, 4 Kbytes, 16 Kbytes, 64 Kbytes, etc.).

The virtual address bits which select a page (and thus are translated) are called the *page address*. The lower bits which select a byte within the selected page are called the *offset* and are not translated. The number of offset bits varies from 12 to 24 bits, depending on the page size.

Cache Page Size Encodings

Page size is defined in each TLB entry's *PageMask* field. This field is loaded or read using the *PageMask* register, as described in Chapter 14.

Each entry translates a pair of physical pages. The low bit of the virtual address page is not compared, because it is used to select between these two physical pages.

Using the TLB

Translations are maintained by the operating system, using page tables in memory. A subset of these translations are loaded into a hardware buffer called the **translation-lookaside buffer** or TLB. The contents of this buffer are maintained by the operating system; if an instruction needs a translation which is not already in the buffer, an exception is taken so the operating system can compute and load the needed translation. If all the necessary translations are present, the program is executed without any delays.

The TLB contains 64 entries, each of which maps a pair of virtual pages. Formats of TLB entries are shown in Figure 16-5.

Cache Algorithm Field

The *Cache Algorithm* fields of the TLB, *EntryLo0*, *EntryLo1*, and *Config* registers indicate how data is cached. Cache algorithms are described in Chapter 4.

Format of a TLB Entry

Figure 16-5 shows the TLB entry formats for both 32- and 64-bit modes. Each field of an entry has a corresponding field in the *EntryHi*, *EntryLo0*, *EntryLo1*, or *PageMask* registers, as shown in Chapter 14; for example the *PFN* and uncached attribute (*UC*) fields of the TLB entry are also held in the *EntryLo* registers.

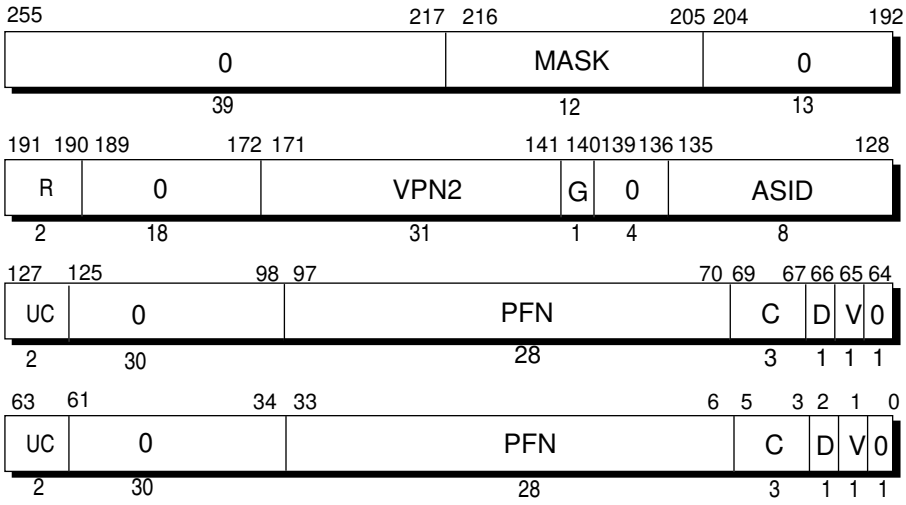

Figure 16-5 Format of a TLB Entry

Address Translation

Because a 64-bit address is unnecessarily large, only the low 44 address bits are translated. The high two virtual address bits (bits 63:62) select between user, supervisor, and kernel address spaces. The intermediate address bits (61:44) must either be all zeros or all ones, depending on the address region. The TLB does not include virtual address bits 61:59, because these are decoded only in the *xkphys* region, which is unmapped.

For data cache accesses, the **joint TLB** (JTLB) translates addresses from the address calculate unit. For instruction accesses, the JTLB translates the PC address if it misses in the instruction TLB (ITLB). That entry is copied into the ITLB for subsequent accesses. The ITLB is transparent to system software.

Address Space Identification (*ASID*)

Each independent task, or *process*, has a separate address space, assigned a unique 8-bit Address Space Identifier (**ASID**). This identifier is stored with each TLB entry to distinguish between entries loaded for different processes. The ASID allows the processor to move from one process to another (called a **context switch**) without having to invalidate TLB entries.

The processor's current ASID is stored in the low 8 bits of the *EntryHi* register. These bits are also used to load the *ASID* field of an entry during TLB refill.

The *ASID* field of each TLB entry is compared to the *EntryHi* register; if the ASIDs are equal or if the entry is global (see below), this TLB entry may be used to translate virtual addresses. The ASID comparison is performed only when a new value is loaded into the *EntryHi* register; the one-bit result of the match is stored in a static Enable latch. (This bit is set whenever a new entry is loaded.)

Global Processes (*G*)

A translation may be defined as *global* so that it can be shared by all processes. This *G* bit is set in the TLB entry and enables the entry independent of its ASID value.

Avoiding TLB Conflict

Setting the *TS* bit in the *Status* register indicates an entry being presented to the TLB matches more than one virtual page entry in the TLB. Any TLB entries that allow multiple matches, even in the *Wired* area, are invalidated before the new entry can be written into the TLB. This prevents multiple matches during address translation.

CPU Exceptions

17

This chapter describes the processor exceptions—a general view of the cause and return of an exception, exception vector locations, and the types of exceptions that are supported, including the cause, processing, and servicing of each exception.

17.1 Causing and Returning from an Exception

When the processor takes an exception, the *EXL* bit in the *Status* register is set to 1, which means the system is in Kernel mode. After saving the appropriate state, the exception handler typically changes the *KSU* bits in the *Status* register to Kernel mode and resets the *EXL* bit back to 0. When restoring the state and restarting, the handler restores the previous value of the *KSU* field and sets the *EXL* bit back to 1.

Returning from an exception also resets the *EXL* bit to 0 (see the ERET instruction in Appendix A).

17.2 Exception Vector Locations

The Cold Reset, Soft Reset, and NMI exceptions are always vectored to the dedicated Cold Reset exception vector at an uncached and unmapped address. Addresses for all other exceptions are a combination of a *vector offset* and a *base address*.

The boot-time vectors (when *BEV* = 1 in the *Status* register) are at uncached and unmapped addresses. During normal operation (when *BEV* = 0) the regular exceptions have vectors in cached address spaces; Cache Error is always at an uncached address so that cache error handling can bypass a suspect cache.

The exception vector assignments for the R10000 processor shown in Table 17-1; the addresses are the same as for the R4400.

Table 17-1 Exception Vector Addresses

BEV	Exception Type	Exception Vector Address	
		32-bit	64-bit
	Cold Reset/Soft Reset/ NMI	0xBFC00000	0xFFFFFFFF BFC00000
BEV=0	TLB Refill (EXL=0)	0x80000000	0xFFFFFFFF 80000000
	XTLB Refill (EXL=0)	0x80000080	0xFFFFFFFF 80000080
	Cache Error	0xA0000100	0xFFFFFFFF A0000100
	Others	0x80000180	0xFFFFFFFF 80000180
BEV=1	TLB Refill (EXL=0)	0xBFC00200	0xFFFFFFFF BFC00200
	XTLB Refill (EXL=0)	0xBFC00280	0xFFFFFFFF BFC00280
	Cache Error	0xBFC00300	0xFFFFFFFF BFC00300
	Others	0xBFC00380	0xFFFFFFFF BFC00380

17.3 TLB Refill Vector Selection

In all present implementations of the MIPS III ISA, there are two TLB refill exception vectors:

- one for references to 32-bit address space (TLB Refill)
- one for references to 64-bit address space (XTLB Refill)

Table 17-2 lists the exception vector addresses.

The TLB refill vector selection is based on the address space of the address (*user*, *supervisor*, or *kernel*) that caused the TLB miss, and the value of the corresponding extended addressing bit in the *Status* register (*UX*, *SX*, or *KX*). The current operating mode of the processor is not important except that it plays a part in specifying in which address space an address resides. The *Context* and *XContext* registers are entirely separate page-table-pointer registers that point to and refill from two separate page tables, however these two registers share *BadVPN2* fields (see Chapter 14 for more information). For all TLB exceptions (Refill, Invalid, TLBL or TLBS), the *BadVPN2* fields of both registers are loaded as they were in the R4400.

In contrast to the R10000, the R4400 processor selects the vector based on the current operating mode of the processor (*user*, *supervisor*, or *kernel*) and the value of the corresponding extended addressing bit in the *Status* register (*UX*, *SX* or *KX*). In addition, the *Context* and *XContext* registers are not implemented as entirely separate registers; the *PTEbase* fields are shared. A miss to a particular address goes through either TLB Refill or XTLB Refill, depending on the source of the reference. There can be only be a single page table unless the refill handlers execute address-deciphering and page table selection in software.

NOTE: Refills for the 0.5 Gbyte supervisor mapped region, *sseg/ksseg*, are controlled by the value of *KX* rather than *SX*. This simplifies control of the processor when supervisor mode is not being used.

Table 17-2 lists the TLB refill vector locations, based on the address that caused the TLB miss and its corresponding mode bit.

Table 17-2 TLB Refill Vectors

Space	**Address Range**	**Regions**	Exception Vector
Kernel	0xFFFF FFFF E000 0000 to 0xFFFF FFFF FFFF FFFF	*kseg3*	Refill (KX=0) or XRefill (KX=1)
Supervisor	0xFFFF FFFF C000 0000 to 0xFFFF FFFF DFFF FFFF	*sseg*, *ksseg*	Refill (KX=0) or XRefill (KX=1)
Kernel	0xC000 0000 0000 0000 to 0xC000 0FFE FFFF FFFF	*xkseg*	XRefill(KX=1)
Supervisor	0x4000 0000 0000 0000 to 0x4000 0FFF FFFF FFFF	*xsseg*, *xksseg*	XRefill (SX=1)
User	0x0000 0000 8000 0000 to 0x0000 0FFF FFFF FFFF	*xsuseg*, *xuseg*, *xkuseg*	XRefill (UX=1)
User	0x0000 0000 0000 0000 to 0x0000 0000 7FFF FFFF	*useg*, *xuseg*, *suseg*, *xsuseg*, *kuseg*, *xkuseg*	Refill (UX=0) or XRefill (UX=1)

Priority of Exceptions

The remainder of this chapter describes exceptions in the order of their priority shown in Table 17-3 (with certain of the exceptions, such as the TLB exceptions and Instruction/Data exceptions, grouped together for convenience). While more than one exception can occur for a single instruction, only the exception with the highest priority is reported. Some exceptions are not caused by the instruction executed at the time, and some exceptions may be deferred. See the individual description of each exception in this chapter for more detail.

Table 17-3 Exception Priority Order

Cold Reset *(highest priority)*
Soft Reset
Nonmaskable Interrupt (NMI)*
Cache error — Instruction cache*
Cache error — Data cache*
Cache error — Secondary cache*
Cache error — System interface*
Address error — Instruction fetch
TLB refill — Instruction fetch
TLB invalid — Instruction fetch
Bus error — Instruction fetch
Integer overflow, Trap, System Call, Breakpoint, Reserved Instruction, Coprocessor Unusable, or Floating-Point Exception
Address error — Data access
TLB refill — Data access
TLB invalid — Data access
TLB modified — Data write
Watch*
Bus error — Data access
Interrupt *(lowest priority)**

* These exceptions are interrupt types, and may be imprecise. Priority may not be followed when considering a specific instruction.

Generally speaking, the exceptions described in the following sections are handled ("processed") by hardware; these exceptions are then serviced by software.

Cold Reset Exception

Cause

The Cold Reset exception is taken for a power-on or "cold" reset; it occurs when the **SysGnt*** signal is asserted while the **SysReset*** signal is also asserted.† This exception is not maskable.

Processing

The CPU provides a special interrupt vector for this exception:

- location 0xBFC0 0000 in 32-bit mode
- location 0xFFFF FFFF BFC0 0000 in 64-bit mode

The Cold Reset vector resides in unmapped and uncached CPU address space, so the hardware need not initialize the TLB or the cache to process this exception. It also means the processor can fetch and execute instructions while the caches and virtual memory are in an undefined state.

The contents of all registers in the CPU are undefined when this exception occurs, except for the following register fields:

- In the *Status* register, *SR* and *TS* are cleared to 0, and *ERL* and *BEV* are set to 1. All other bits are undefined.
- *Config* register is initialized with the boot mode bits read from the serial input.
- The *Random* register is initialized to the value of its upper bound.
- The *Wired* register is initialized to 0.
- The *EW* bit in the *CacheErr* register is cleared.
- The *ErrorEPC* register gets the PC.
- The *FrameMask* register is set to 0.
- Branch prediction bits are set to 0.
- *Performance Counter* register *Event* field is set to 0.
- All pending cache errors, delayed watch exceptions, and external interrupts are cleared.

Servicing

The Cold Reset exception is serviced by:

- initializing all processor registers, coprocessor registers, caches, and the memory system
- performing diagnostic tests
- bootstrapping the operating system

† If **SysGnt*** remains deasserted (high) while **SysReset*** is asserted, the processor interprets this as a Soft Reset exception.

Soft† Reset Exception

Cause

The Soft Reset exception occurs in response to a Soft Reset (see Chapter 8).

A Soft Reset exception is not maskable.

The processor differentiates between a Cold Reset and a Soft Reset as follows:

- A Cold Reset occurs when the **SysGnt*** signal is asserted while the **SysReset*** signal is also asserted.
- A Soft Reset occurs if the **SysGnt*** signal remains negated when a **SysReset*** signal is asserted.

In R4400 processor, there is no way for software to differentiate between a Soft Reset exception and an NMI exception. In the R10000 processor, a bit labelled *NMI* has been added to the *Status* register to distinguish between these two exceptions. Both Soft Reset and NMI exceptions set the *SR* bit and use the same exception vector. During an NMI exception, the *NMI* bit is set to 1; during a Soft Reset, the *NMI* bit is set to 0.

Processing

When a Soft Reset exception occurs, the *SR* bit of the *Status* register is set, distinguishing this exception from a Cold Reset exception.

When a Soft Reset is detected, the processor initializes minimum processor state. This allows the processor to fetch and execute the instructions of the exception handler, which in turn dumps the current architectural state to external logic. Hardware state that loses architectural state is not initialized unless it is necessary to execute instructions from unmapped uncached space that reads the registers, TLB, and cache contents.

The Soft Reset can begin on an arbitrary cycle boundary and can abort multicycle operations in progress, so it may alter machine state. Hence, caches, memory, or other processor states can be inconsistent: data cache blocks may stay at the refill state and any cached loads/stores to these blocks will hang the processor. Therefore, CacheOps should be used to dump the cache contents.

After the processor state is read out, the processor should be reset with a Cold Reset sequence.

† Soft Reset is also known colloquially as *Warm* Reset.

A Soft Reset exception preserves the contents of all registers, except for:

- *ErrorEPC* register, which contains the PC
- *ERL* bit of the *Status* register, which is set to 1
- *SR* bit of the *Status* register, which is set to 1 on Soft Reset or an NMI; 0 for a Cold Reset
- *BEV* bit of the *Status* register, which is set to 1
- *TS* bit of the *Status* register, which is set to 0
- PC is set to the reset vector 0xFFFF FFFF BFC0 0000
- clears any pending Cache Error exceptions

Servicing

A Soft Reset exception is intended to quickly reinitialize a previously operating processor after a fatal error.

It is not normally possible to continue program execution after returning from this exception, since a **SysReset*** signal can be accepted anytime.

NMI Exception

Cause

The NMI exception is caused by assertion of the **SysNMI*** signal.

An NMI exception is not maskable.

In R4400 processor, there is no way for software to differentiate between a Soft Reset exception and an NMI exception. In the R10000 processor, a bit labelled *NMI* has been added to the *Status* register to distinguish between these two exceptions. Both Soft Reset and NMI exceptions set the *SR* bit and use the same exception vector. During an NMI exception, the *NMI* bit is set to 1; during a Soft Reset, the *NMI* bit is set to 0.

Processing

When an NMI exception occurs, the *SR* bit of the *Status* register is set, distinguishing this exception from a Cold Reset exception.

An exception caused by an NMI is taken at the instruction boundary. It does not abort any state machines, preserving the state of the processor for diagnosis. The *Cause* register remains unchanged and the system jumps to the NMI exception handler (see Table 17-1).

An NMI exception preserves the contents of all registers, except for:

- *ErrorEPC* register, which contains the PC
- *ERL* bit of the *Status* register, which is set to 1
- *SR* bit of the *Status* register, which is set to 1 on Soft Reset or an NMI; 0 for a Cold Reset
- *BEV* bit of the *Status* register, which is set to 1
- *TS* bit of the *Status* register, which is set to 0
- PC is set to the reset vector 0xFFFF FFFF BFC0 0000
- clears any pending Cache Error exceptions

Servicing

The NMI can be used for purposes other than resetting the processor while preserving cache and memory contents. For example, the system might use an NMI to cause an immediate, controlled shutdown when it detects an impending power failure.

It is not normally possible to continue program execution after returning from this exception, since an NMI can occur during another error exception.

Address Error Exception

Cause

The Address Error exception occurs when an attempt is made to execute one of the following:

- reference to an illegal address space
- reference the supervisor address space from User mode
- reference the kernel address space from User or Supervisor mode
- load or store a doubleword that is not aligned on a doubleword boundary
- load, fetch, or store a word that is not aligned on a word boundary
- load or store a halfword that is not aligned on a halfword boundary

This exception is not maskable.

Processing

The common exception vector is used for this exception. The *AdEL* or *AdES* code in the *Cause* register is set, indicating whether the instruction caused the exception with an instruction reference, load operation, or store operation shown by the *EPC* register and *BD* bit in the *Cause* register.

When this exception occurs, the *BadVAddr* register retains the virtual address that was not properly aligned or that referenced protected address space. The contents of the *VPN* field of the *Context*, *XContext*, and *EntryHi* registers are undefined, as are the contents of the *EntryLo* register.

The *EPC* register contains the address of the instruction that caused the exception, unless this instruction is in a branch delay slot. If it is in a branch delay slot, the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set as indication.

Servicing

The process executing at the time is handed a UNIX SIGSEGV (segmentation violation) signal. This error is usually fatal to the process incurring the exception.

TLB Exceptions

Three types of TLB exceptions can occur:

- TLB Refill occurs when there is no TLB entry that matches an attempted reference to a mapped address space.
- TLB Invalid occurs when a virtual address reference matches a TLB entry that is marked invalid.
- TLB Modified occurs when a store operation virtual address reference to memory matches a TLB entry which is marked valid but is not dirty (the entry is not writable).

The following three sections describe these TLB exceptions.

NOTE: TLB Refill vector selection is also described earlier in this chapter, in the section titled, TLB Refill Vector Selection.

TLB Refill Exception

Cause

The TLB refill exception occurs when there is no TLB entry to match a reference to a mapped address space. This exception is not maskable.

Processing

There are two special exception vectors for this exception; one for references to 32-bit address spaces, and one for references to 64-bit address spaces. The *UX, SX,* and *KX* bits of the *Status* register determine whether the user, supervisor or kernel address spaces referenced are 32-bit or 64-bit spaces; the TLB refill vector is selected based upon the address space of the address causing the TLB miss (user, supervisor, or kernel mode address space), together with the value of the corresponding extended addressing bit in the *Status* register (*UX*, *SX*, or *KX*). The current operating mode of the processor is not important except that it plays a part in specifying in which space an address resides. An address is in *user* space if it is in *useg*, *suseg*, *kuseg*, *xuseg*, *xsuseg*, or *xkuseg* (see the description of virtual address spaces in Chapter 16). An address is in *supervisor* space if it is in *sseg*, *ksseg*, *xsseg* or *xksseg*, and an address is in *kernel* space if it is in either *kseg3* or *xkseg*. *Kseg0*, *kseg1*, and kernel physical spaces (*xkphys*) are kernel spaces but are not mapped.

All references use these vectors when the *EXL* bit is set to 0 in the *Status* register. This exception sets the *TLBL* or *TLBS* code in the *ExcCode* field of the *Cause* register. This code indicates whether the instruction, as shown by the *EPC* register and the *BD* bit in the *Cause* register, caused the miss by an instruction reference, load operation, or store operation.

When this exception occurs, the *BadVAddr*, *Context*, *XContext* and *EntryHi* registers hold the virtual address that failed address translation. The *EntryHi* register also contains the ASID from which the translation fault occurred. The *Random* register normally contains a valid location in which to place the replacement TLB entry. The contents of the *EntryLo* register are undefined. The *EPC* register contains the address of the instruction that caused the exception, unless this instruction is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set.

Servicing

To service this exception, the contents of the *Context* or *XContext* register are used as a virtual address to fetch memory locations containing the physical page frame and access control bits for a pair of TLB entries. The two entries are placed into the E*ntryLo0/EntryLo1* register; the *EntryHi* and *EntryLo* registers are written into the TLB.

It is possible that the virtual address used to obtain the physical address and access control information is on a page that is not resident in the TLB. This condition is processed by allowing a TLB refill exception in the TLB refill handler. This second exception goes to the common exception vector because the *EXL* bit of the *Status* register is set.

TLB Invalid Exception

Cause

The TLB invalid exception occurs when a virtual address reference matches a TLB entry that is marked invalid (TLB valid bit cleared). This exception is not maskable.

Processing

The common exception vector is used for this exception. The *TLBL* or *TLBS* code in the *ExcCode* field of the *Cause* register is set. This indicates whether the instruction, as shown by the *EPC* register and *BD* bit in the *Cause* register, caused the miss by an instruction reference, load operation, or store operation.

When this exception occurs, the *BadVAddr*, *Context*, *XContext* and *EntryHi* registers contain the virtual address that failed address translation. The *EntryHi* register also contains the ASID from which the translation fault occurred. The *Random* register normally contains a valid location in which to put the replacement TLB entry. The contents of the *EntryLo* registers are undefined.

The *EPC* register contains the address of the instruction that caused the exception unless this instruction is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set.

Servicing

A TLB entry is typically marked invalid when one of the following is true:

- a virtual address does not exist
- the virtual address exists, but is not in main memory (a page fault)
- a trap is desired on any reference to the page (for example, to maintain a reference bit)

After servicing the cause of a TLB Invalid exception, the TLB entry is located with TLBP (TLB Probe), and replaced by an entry with that entry's *Valid* bit set.

TLB Modified Exception

Cause

The TLB modified exception occurs when a store operation virtual address reference to memory matches a TLB entry that is marked valid but is not dirty and therefore is not writable. This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *Mod* code in the *Cause* register is set.

When this exception occurs, the *BadVAddr*, *Context*, *XContext* and *EntryHi* registers contain the virtual address that failed address translation. The *EntryHi* register also contains the ASID from which the translation fault occurred. The contents of the *EntryLo* register are undefined.

The *EPC* register contains the address of the instruction that caused the exception unless that instruction is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set.

Servicing

The kernel uses the failed virtual address or virtual page number to identify the corresponding access control information. The page identified may or may not permit write accesses; if writes are not permitted, a write protection violation occurs.

If write accesses are permitted, the page frame is marked dirty/writable by the kernel in its own data structures. The TLBP instruction places the index of the TLB entry that must be altered into the *Index* register. The *EntryLo* register is loaded with a word containing the physical page frame and access control bits (with the *D* bit set), and the *EntryHi* and *EntryLo* registers are written into the TLB.

Cache Error Exception

The Cache Error exception is described in Chapter 9.

Virtual Coherency Exception

The Virtual Coherency exception is not implemented in the R10000 processor, since the virtual coherency condition is handled in hardware.

In the R4400 processor, a Virtual Coherency exception occurs when a primary cache miss hits in the secondary cache but **VA[14:12]** are not the same as the *PIdx* field of the secondary cache tag, and the cache algorithm specifies that the page is cached. When such a situation is detected in the R10000 processor, the primary cache lines at the old virtual index are invalidated and the *PIdx* field of the secondary cache is written with the new virtual index.

Bus Error Exception

Cause

A Bus Error exception occurs when a processor block read, upgrade, or double/single/partial-word read request receives an external ERR completion response, or a processor double/single/partial-word read request receives an external ACK completion response where the associated external double/single/partial-word data response contains an uncorrectable error. This exception is not maskable.

Processing

The common interrupt vector is used for a Bus Error exception. The *IBE* or *DBE* code in the *ExcCode* field of the *Cause* register is set, signifying whether the instruction (as indicated by the *EPC* register and *BD* bit in the *Cause* register) caused the exception by an instruction reference, load operation, or store operation.

The *EPC* register contains the address of the instruction that caused the exception, unless it is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set.

Servicing

The physical address at which the fault occurred can be computed from information available in the CP0 registers.

- If the *IBE* code in the *Cause* register is set (indicating an instruction fetch reference), the instruction that caused the exception is located at the virtual address contained in the *EPC* register (or 4+ the contents of the *EPC* register if the *BD* bit of the *Cause* register is set).
- If the *DBE* code is set (indicating a load or store reference), the instruction that caused the exception is located at the virtual address contained in the *EPC* register (or 4+ the contents of the *EPC* register if the *BD* bit of the *Cause* register is set).

The virtual address of the load and store reference can then be obtained by interpreting the instruction. The physical address can be obtained by using the TLBP instruction and reading the *EntryLo* registers to compute the physical page number. The process executing at the time of this exception is handed a UNIX SIGBUS (bus error) signal, which is usually fatal.

Integer Overflow Exception

Cause

An Integer Overflow exception occurs when an ADD, ADDI, SUB, DADD, DADDI or DSUB† instruction results in a 2's complement overflow. This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *OV* code in the *Cause* register is set.

The *EPC* register contains the address of the instruction that caused the exception unless the instruction is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set.

Servicing

The process executing at the time of the exception is handed a UNIX SIGFPE/FPE_INTOVF_TRAP (floating-point exception/integer overflow) signal. This error is usually fatal to the current process.

† See Appendix A for a description of these instructions.

Trap Exception

Cause

The Trap exception occurs when a TGE, TGEU, TLT, TLTU, TEQ, TNE, TGEI, TGEUI, TLTI, TLTUI, TEQI, or TNEI instruction results in a TRUE condition. This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *Tr* code in the *Cause* register is set.

The *EPC* register contains the address of the instruction causing the exception unless the instruction is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction and the *BD* bit of the *Cause* register is set.

Servicing

The process executing at the time of a Trap exception is handed a UNIX SIGFPE/FPE_INTOVF_TRAP (floating-point exception/integer overflow) signal. This error is usually fatal.

System Call Exception

Cause

A System Call exception occurs during an attempt to execute the SYSCALL instruction. This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *Sys* code in the *Cause* register is set.

The *EPC* register contains the address of the SYSCALL instruction unless it is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction.

If the SYSCALL instruction is in a branch delay slot, the *BD* bit of the *Status* register is set; otherwise this bit is cleared.

Servicing

When the System Call exception occurs, control is transferred to the applicable system routine. Additional distinctions can be made by analyzing the *Code* field of the SYSCALL instruction (bits 25:6), and loading the contents of the instruction whose address the EPC register contains.

To resume execution, the *EPC* register must be altered so that the SYSCALL instruction does not re-execute; this is accomplished by adding a value of 4 to the *EPC* register (*EPC* register + 4) before returning.

If a SYSCALL instruction is in a branch delay slot, a more complicated algorithm, beyond the scope of this description, may be required.

Breakpoint Exception

Cause

A Breakpoint exception occurs when an attempt is made to execute the BREAK instruction. This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *BP* code in the *Cause* register is set.

The *EPC* register contains the address of the BREAK instruction unless it is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction.

If the BREAK instruction is in a branch delay slot, the *BD* bit of the *Status* register is set, otherwise the bit is cleared.

Servicing

When the Breakpoint exception occurs, control is transferred to the applicable system routine. Additional distinctions can be made by analyzing the *Code* field of the BREAK instruction (bits 25:6), and loading the contents of the instruction whose address the EPC register contains. A value of 4 must be added to the contents of the EPC register (EPC register + 4) to locate the instruction if it resides in a branch delay slot.

To resume execution, the EPC register must be altered so that the BREAK instruction does not re-execute; this is accomplished by adding a value of 4 to the EPC register (EPC register + 4) before returning.

If a BREAK instruction is in a branch delay slot, interpretation of the branch instruction is required to resume execution.

Reserved Instruction Exception

Cause

The Reserved Instruction exception occurs when one of the following conditions occurs:

- an attempt is made to execute an instruction with an undefined major opcode (bits 31:26)
- an attempt is made to execute a SPECIAL instruction with an undefined minor opcode (bits 5:0)
- an attempt is made to execute a REGIMM instruction with an undefined minor opcode (bits 20:16)
- an attempt is made to execute 64-bit operations in 32-bit mode when in User or Supervisor modes
- an attempt is made to execute a COP1X when the MIPS IV ISA is not enabled

64-bit operations are always valid in Kernel mode regardless of the value of the *KX* bit in the *Status* register.

This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *RI* code in the *Cause* register is set.

The *EPC* register contains the address of the reserved instruction unless it is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction.

Servicing

No instructions in the MIPS ISA are currently interpreted. The process executing at the time of this exception is handed a UNIX SIGILL/ILL_RESOP_FAULT (illegal instruction/reserved operand fault) signal. This error is usually fatal.

Coprocessor Unusable Exception

Cause

The Coprocessor Unusable exception occurs when an attempt is made to execute a coprocessor instruction for either:

- a corresponding coprocessor unit (CP1 or CP2) that has not been marked usable, or
- CP0 instructions, when the unit has not been marked usable and the process executes in either User or Supervisor mode.

This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *CpU* code in the *Cause* register is set. The contents of the *Coprocessor Usage Error* field of the coprocessor *Control* register indicate which of the four coprocessors was referenced. The *EPC* register contains the address of the unusable coprocessor instruction unless it is in a branch delay slot, in which case the *EPC* register contains the address of the preceding branch instruction.

Servicing

The coprocessor unit to which an attempted reference was made is identified by the Coprocessor Usage Error field, which results in one of the following situations:

- If the process is entitled access to the coprocessor, the coprocessor is marked usable and the corresponding user state is restored to the coprocessor.
- If the process is entitled access to the coprocessor, but the coprocessor does not exist or has failed, interpretation of the coprocessor instruction is possible.
- If the *BD* bit is set in the *Cause* register, the branch instruction must be interpreted; then the coprocessor instruction can be emulated and execution resumed with the *EPC* register advanced past the coprocessor instruction.
- If the process is not entitled access to the coprocessor, the process executing at the time is handed a UNIX SIGILL/ILL_PRIVIN_FAULT (illegal instruction/privileged instruction fault) signal. This error is usually fatal.

Floating-Point Exception

Cause

The Floating-Point exception is used by the floating-point coprocessor. This exception is not maskable.

Processing

The common exception vector is used for this exception, and the *FPE* code in the *Cause* register is set.

The contents of the *Floating-Point Control/Status* register indicate the cause of this exception.

Servicing

This exception is cleared by clearing the appropriate bit in the *Floating-Point Control/Status* register.

Watch Exception

Cause

A Watch exception occurs when a load or store instruction references the physical address specified in the *WatchLo/WatchHi* System Control Coprocessor (CP0) registers. The *WatchLo* register specifies whether a load or store initiated this exception.

The CACHE instruction never causes a Watch exception.

The Watch exception is postponed if either the *EXL* or *ERL* bit is set in the *Status* register. If either bit is set, the instruction referencing the *WatchLo/WatchHi* address is executed and the exception is delayed until the delay condition is cleared; that is, until *ERL* and *EXL* both are cleared (set to 0). The *EPC* contains the address of the next unexecuted instruction.

A delayed Watch exception is cleared by system reset or by writing a value to the *WatchLo* register.†

Watch is maskable by setting the *EXL* or *ERL* bits in the *Status* register.

Processing

The common exception vector is used for this exception, and the *Watch* code in the *Cause* register is set.

Servicing

The Watch exception is a debugging aid; typically the exception handler transfers control to a debugger, allowing the user to examine the situation.

To continue program execution, the Watch exception must be disabled to execute the faulting instruction. The Watch exception must then be reenabled. The faulting instruction can be executed either by interpretation or by setting breakpoints.

† An MTC0 to the *WatchLo* register clears a delayed Watch exception.

Interrupt Exception

Cause

The Interrupt exception occurs when one of the eight interrupt conditions is asserted. The significance of these interrupts is dependent upon the specific system implementation.

Each of the eight interrupts can be masked by clearing the corresponding bit in the *Interrupt-Mask (IM)* field of the *Status* register, and all of the eight interrupts can be masked at once by clearing the *IE* bit of the *Status* register.

Processing

The common exception vector is used for this exception, and the *Int* code in the *Cause* register is set.

The *IP* field of the *Cause* register indicates current interrupt requests. It is possible that more than one of the bits can be simultaneously set (or even *no* bits may be set) if the interrupt is asserted and then deasserted before this register is read.

On Cold Reset, an R4400 processor can be configured with *IP[7]* either as a sixth external interrupt, or as an internal interrupt set when the *Count* register equals the *Compare* register. There is no such option on the R10000 processor; *IP[7]* is always an internal interrupt that is set when one of the following occurs:

- the *Count* register is equal to the *Compare* register
- either one of the two performance counters overflows

Software needs to poll each source to determine the cause of the interrupt (which could come from more than one source at a time). For instance, writing a value to the *Compare* register clears the timer interrupt but it may not clear *IP[7]* if one of the performance counters is simultaneously overflowing. Performance counter interrupts can be disabled individually without affecting the timer interrupt, but there is no way to disable the timer interrupt without disabling the performance counter interrupt.

Servicing

If the interrupt is caused by one of the two software-generated exceptions (described in Chapter 6, Software Interrupts), the interrupt condition is cleared by setting the corresponding *Cause* register bit, *IP[1:0],* to 0. Software interrupts are imprecise. Once the software interrupt is enabled, program execution may continue for several instructions before the exception is taken.

Timer interrupts are cleared by writing to the *Compare* register. The Performance Counter interrupt is cleared by writing a 0 to bit 31, the overflow bit, of the counter.

Cold Reset and Soft Reset exceptions clear all the outstanding external interrupt requests**,** *IP[2]* to *IP[6].*

If the interrupt is hardware-generated, the interrupt condition is cleared by correcting the condition causing the interrupt pin to be asserted.

17.4 MIPSIV Instructions

The system must either be in Kernel or Supervisor mode, or have set the *XX* bit of the *Status* register to a 1 in order to use the MIPS IV instruction set. In User mode, if *XX* is a 0 and an attempt is made to execute MIPS IV instructions, an exception will be taken. The type of exception that will be taken depends upon the type of instruction whose execution was attempted; a list is given in Table 17-4. Note that operating with MIPS IV instructions does not require that MIPS III instruction set or 64-bit addressing is enabled.

MIPS IV instructions that use or modify the floating-point registers (CP1 state) are also affected by the *CU1* bit of the CP0 Status register. If *CU1* is not set, a Coprocessor Unusable exception may be signaled.

The Reserved Instruction (RI), Coprocessor Unusable (CU), and Unimplemented Operation (UO) exceptions for MIPS IV instructions are listed in the Table 17-4 below.

Table 17-4 MIPS IV Instruction Exceptions

Exceptions	Instructions	CU1	MIPS4
RI	CPU (undefined)	-	-
RI	MOVN,Z	-	0
RI	MOVT,F	-	0
CU		0	1
RI	PREF	-	0
CU	COP1 (all instructions)	0	-
UO	(undefined)	1	-
RI	BC (cc>0)	1	0
UO	C (cc>0)	1	0
UO	MOVN,Z,T,F	1	0
UO	RECIP, RSQRT	1	0
RI	COP1X (all instructions)	-	0
CU	(all instructions)	0	1
RI	(undefined)	1	1

17.5 COP0 Instructions

Execution of an RFE instruction causes a Reserved Instruction exception in the R10000 processor.

The execution of undefined COP0 functions is undefined in the R10000 processor.

17.6 COP1 Instructions

The R10000 and R4400 processors do not generate the same exceptions for undefined COP1 instructions. In the R4400 processor, undefined opcodes or formats in the *sub* field take an Unimplemented Operation exceptions. In the R10000 processor, undefined opcodes (bits 25:24 are 0 or 1) take Reserved Instruction exceptions and undefined formats (bits 25:24 are 2 or 3) take Unimplemented Operation exceptions.

In MIPS II on an R4400 processor, the execution of DMTC1, DMFC1, and L format take Unimplemented Operation exceptions. In MIPS II on the R10000 processor, the execution of DMTC1 and DMFC1 take Reserved Instruction exceptions

The attempted execution of the L format takes an Unimplemented Operation exception when the MIPS III mode is not enabled.

A CTC1 instruction that sets both *Cause* and *Enable* bits also forces an immediate floating-point exception; the *EPC* register points to the offending CTC1 instruction.

17.7 COP2 Instructions

If the *CU2* bit of the CP0 *Status* register is not set during an attempted execution of such Coprocessor 2 instructions as COP2, LWC2, SWC2, LDC2, and SDC2, the system takes a Coprocessor Unusable exception.

In the R4400 processor, if the *CU2* bit is set, COP2 instructions are handled as NOPs; the operations of Coprocessor 2 load/store instructions are undefined. In the R10000 processor, an execution of a Coprocessor 2 instruction takes a Reserved Instruction exception when *CU2* bit is set.

Cache Test Mode

18

The R10000 processor provides a cache test mode that may be used during manufacturing test and system debug to access the following internal RAM arrays:

- data cache data array
- data cache tag array
- instruction cache data array
- instruction cache tag array
- secondary cache way predication table

18.1 Interface Signals

Cache test mode is accessed by using a subset of the system interface signals. By not requiring the use of any secondary cache interface signals, the internal RAM arrays may be accessed for single-chip LGA as well as R10000/secondary cache module configurations.

The following system interface signals are used during cache test mode:

- **SysAD(57:0)**
- **SysVal***

Any input signals not listed above are ignored by the processor when it is operating in cache test mode, and any output signals not listed above are undefined during cache test mode.

18.2 System Interface Clock Divisor

Cache test mode is supported for all system interface clock speeds. However, since cache test mode repeat rates and latencies are expressed in terms of **PClk** cycles, the external agent must take care when operating at any system interface clock divisor other than Divide-by-1.

18.3 Entering Cache Test Mode

In order for the processor to enter cache test mode, the external agent must begin a Power-on or Cold Reset sequence.

Rather than negating **SysReset*** at the end of the reset sequence, the external agent loads the mode bits into the processor by driving the mode bits (with the **CTM** signal asserted) on **SysAD(63:0)**, waits at least two **SysClk** cycles, and then asserts **SysGnt*** for at least one **SysClk** cycle.

After waiting at least another 100 ms, the external agent may issue the first cache test mode command.

Figure 18-1 shows the cacheCache Test Mode Entry Sequenc test mode entry sequence.

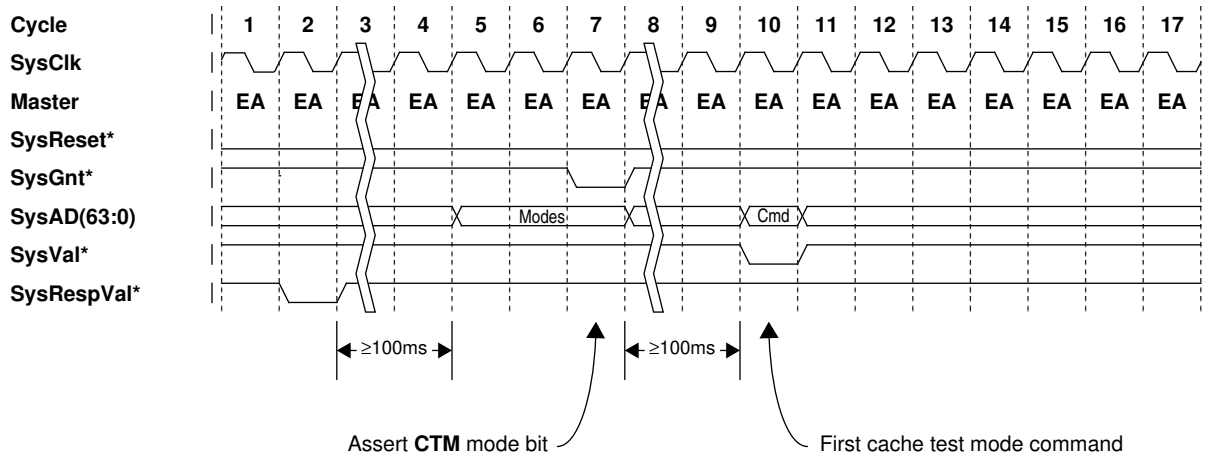

Figure 18-1 Cache Test Mode Entry Sequence

18.4 Exit Sequence

To leave cache test mode, the external agent does the following:

- loads the mode bits into the processor by driving the mode bits (with the **CTM** mode bit negated) on **SysAD(63:0)**
- waits at least two **SysClk** cycles
- asserts **SysGnt*** for at least one **SysClk** cycle

After at least one **SysClk** cycle, the external agent may negate **SysReset*** to end the reset sequence.

Figure 18-2 shows the cache test mode exit sequence.

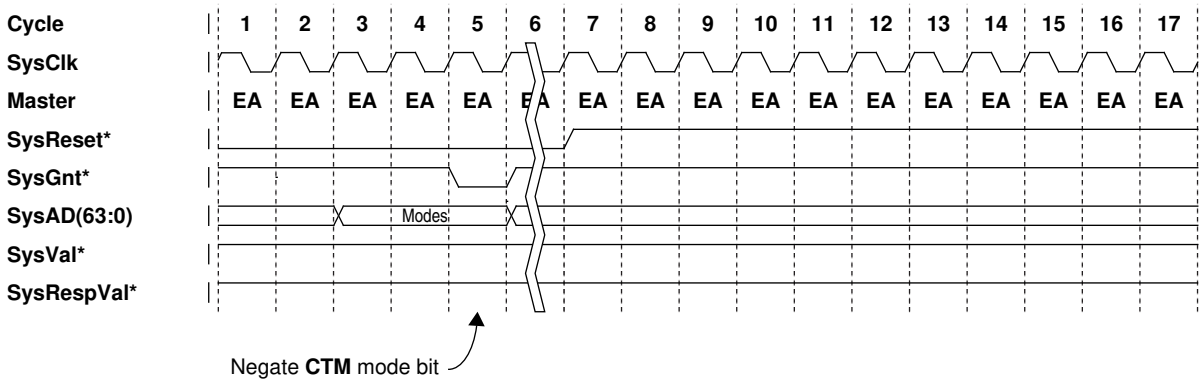

Figure 18-2 Cache Test Mode Exit Sequence

18.5 SysAD(63:0) Encoding

Encoding of the **SysAD(63:0)** bus during cache test mode is shown in Table 18-1. "Unused" fields are read as "undefined," and must be written as zeroes.

Table 18-1 Cache Test Mode **SysAD(63:0)** *Encoding*

<table>
<tr><th>SysAD Bit</th><th>Data Cache Data Array</th><th>Data Cache Tag Array</th><th>Instruction Cache Data Array</th><th>Instruction Cache Tag Array</th><th>Secondary Cache Way Predication Array</th></tr>
<tr><td>0</td><td rowspan="10">Data</td><td>Tag parity</td><td rowspan="11">Data</td><td>Tag parity</td><td>MRU</td></tr>
<tr><td>1</td><td>SCWay</td><td>Unused</td><td rowspan="13">Unused</td></tr>
<tr><td>2</td><td>State parity</td><td>State parity</td></tr>
<tr><td>3</td><td>LRU</td><td>LRU</td></tr>
<tr><td>4</td><td rowspan="2">Unused</td><td rowspan="2">Unused</td></tr>
<tr><td>5</td></tr>
<tr><td>6</td><td rowspan="2">State</td><td>State</td></tr>
<tr><td>7</td><td>Unused</td></tr>
<tr><td>31:8</td><td rowspan="2">Tag</td><td rowspan="2">Tag</td></tr>
<tr><td>35:32</td><td>Data parity</td></tr>
<tr><td>36</td><td rowspan="3">Unused</td><td rowspan="2">StateMod</td><td>Data parity</td><td rowspan="3">Unused</td></tr>
<tr><td>38:37</td><td rowspan="2">Unused</td></tr>
<tr><td>39</td><td>Unused</td></tr>
<tr><td rowspan="2">42:40</td><td>0</td><td>1</td><td>2</td><td>3</td><td>4</td></tr>
<tr><td colspan="5">Array select</td></tr>
<tr><td>43</td><td colspan="5">Write/Read select</td></tr>
<tr><td>44</td><td colspan="5">Auto-increment select</td></tr>
<tr><td>45</td><td colspan="5">Way</td></tr>
<tr><td>57:46</td><td colspan="5">Address</td></tr>
<tr><td>63:58</td><td colspan="5">Unused</td></tr>
</table>

18.6 Cache Test Mode Protocol

This section describes the cache test mode protocol in detail, including:

- normal write protocol
- auto-increment protocol
- normal read protocol
- auto-increment read protocol

Normal Write Protocol

A cache test mode **normal write** operation writes a selected RAM array. The write address, way, array, and data are specified in the write command.

The external agent issues a normal write command by:

- driving the address on **SysAD(57:46)**
- driving the way on **SysAD(45)**
- negating the auto-increment select on **SysAD(44)**
- asserting the Write/Read select on **SysAD(43)**
- driving the array select on **SysAD(42:40)**
- driving the write data on **SysAD(39:0)**
- asserting **SysVal*** for one **SysClk** cycle

Normal writes have a repeat rate of 8 **PClk** cycles.

Figure 18-3 depicts two cache test mode normal writes.

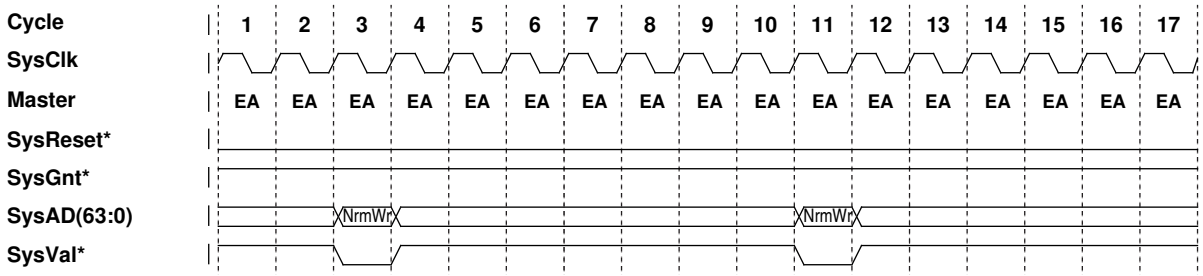

Figure 18-3 Cache Test Mode Normal Write Protocol

Auto-Increment Write Protocol

A cache test mode **auto-increment write** operation writes a selected RAM array. The write address is obtained by incrementing the previous write address, and the write way is obtained from the previous write way.

If an overflow occurs when incrementing the previous write address, the address wraps to 0, and the way is toggled.

The write data is identical to the previous write data.

For proper results, an auto-increment write must always be proceeded by a normal or auto-increment write.

The external agent issues an auto-increment write command by:

- asserting the auto-increment select on **SysAD(44)**
- asserting the Write/Read select on **SysAD(43)**
- driving the array select on **SysAD(42:40)**
- asserting **SysVal*** for one **SysClk** cycle

Auto-increment writes have a repeat rate of one **PClk** cycle.

Figure 18-4 depicts three cache test mode auto-increment writes.

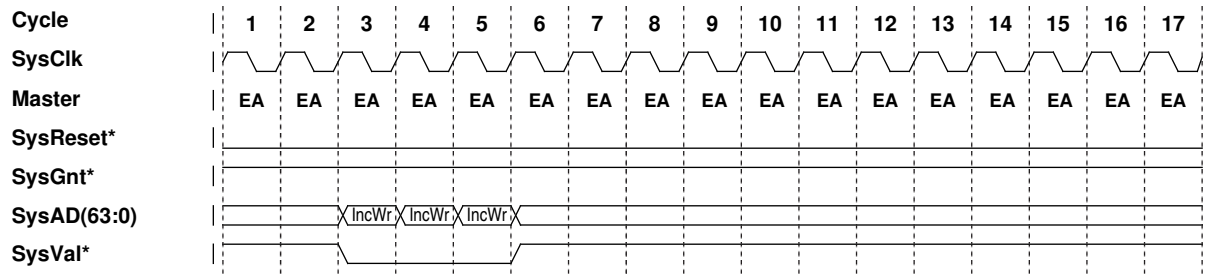

Figure 18-4 Cache Test Mode Auto-Increment Write Protocol

Normal Read Protocol

A cache test mode **normal read** operation reads a selected RAM array. The read address, way, and array are specified by the read command.

The external agent issues a normal read command by:

- driving the address on **SysAD(57:46)**
- driving the way on **SysAD(45)**
- negating the auto-increment select on **SysAD(44)**
- negating the Write/Read select on **SysAD(43)**
- driving the array select on **SysAD(42:40)**
- asserting **SysVal*** for one **SysClk** cycle.

After a read latency of 15 **PClk** cycles, the processor provides the read response by:

- entering *Master* state
- driving the read data on **SysAD(39:0)**
- asserting **SysVal*** for one **SysClk** cycle.

In the following **SysClk** cycle, the processor reverts to *Slave* state.

Normal reads have a repeat rate of 17 **PClk** cycles.

Figure 18-5 depicts two cache test mode normal reads.

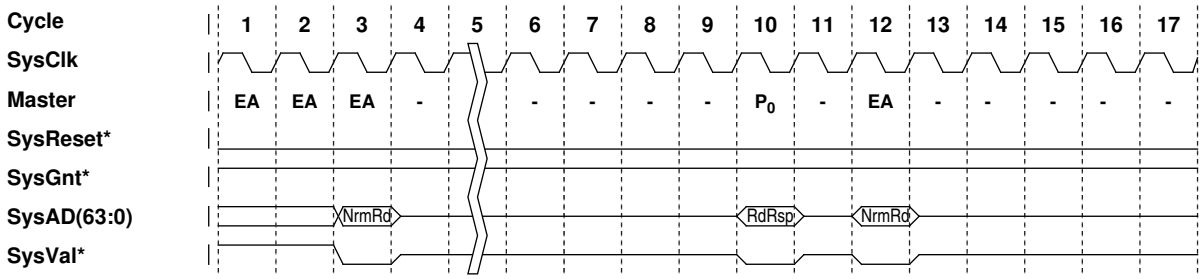

Figure 18-5 Cache Test Mode Normal Read Protocol

Auto-Increment Read Protocol

A cache test mode **auto-increment read** operation reads a selected RAM array. The read address is obtained by incrementing the previous access address, and the read way is obtained from the previous access way.

If an overflow occurs when incrementing the previous access address, the address wraps to 0, and the way is toggled.

The external agent issues an auto-increment read command by:

- asserting the auto-increment select on **SysAD(44)**
- negating the Write/Read select on **SysAD(43)**
- driving the array select on **SysAD(42:40)**
- asserting **SysVal*** for one **SysClk** cycle.

After a read latency of 15 **PClk** cycles, the processor provides the read response by:

- entering *Master* state
- driving the read data on **SysAD(39:0)**
- asserting **SysVal*** for one **SysClk** cycle.

In the following **SysClk** cycle, the processor reverts to *Slave* state.

Auto-increment reads have a repeat rate of 17 **PClk** cycles.

Figure 18-6 depicts two cache test mode auto-increment reads.

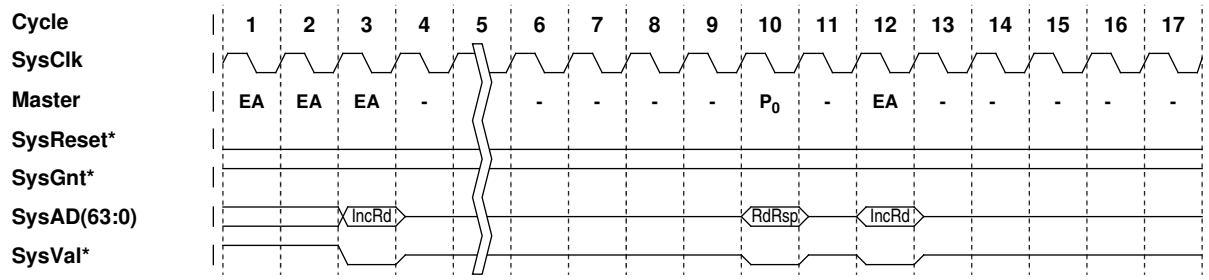

Figure 18-6 Cache Test Mode Auto-Increment Read Protocol

Glossary

A

The following terms are defined in this Glossary:

- superscalar processor
- pipeline
- pipeline latency
- pipeline repeat rate
- out-of-order execution
- dynamic scheduling
- instruction fetch, decode, issue, execution, completion, and graduation
- active list
- free list and busy registers
- register renaming and unnaming
- nonblocking loads and stores
- speculative branching
- logical and physical registers
- register files
- ANDES architecture

A.1 Superscalar Processor

A superscalar processor is one that can fetch, execute and complete more than one instruction in parallel. By implication, a superscalar processor has more than one pipeline (see below).

A.2 Pipeline

In the processor **pipeline**, the execution of each instruction is divided into a sequence of simpler suboperations. Each suboperation is performed by a separate hardware section called a **stage**, and each stage passes its result to a succeeding stage.

Normally, each instruction only remains in each stage for a single cycle, and each stage begins executing a new instruction as previous instructions are being completed in later stages. Thus, a new instruction can often begin during every cycle.

Pipelines greatly improve the rate at which instructions can be executed, as long as there are no dependencies. The efficient use of a pipeline requires that several instructions be executed in parallel, however the result of any instruction is not available for several cycles after that instruction has entered the pipeline. Thus, new instructions must not depend on the results of instructions which are still in the pipeline.

A.3 Pipeline Latency

The **latency** of an execution pipeline is the number of cycles between the time an instruction is issued and the time a dependent instruction (which uses its result as an operand) can be issued.

In the R10000 processor, most integer instructions have a single-cycle latency, load instructions have a 2-cycle latency for cache hits, and floating-point addition and multiplication have a 2-cycle latency. Integer multiply, floating-point square-root, and all divide instructions are computed iteratively and have longer latencies.

A.4 Pipeline Repeat Rate

The **repeat rate** of the pipeline is the number of cycles that occur between the issuance of one instruction and the issuance of the next instruction to the same execution unit. In the R10000 processor, the main five pipelines all have repeat rates of one cycle, but the iterative units have longer repeat delays.

A.5 Out-of-Order Execution

The “program order” of instructions is the sequence in which they are fetched and decoded. In the R10000 processor, instructions may be issued, executed, and completed **out of program order**. They are always graduated in program order.

A.6 Dynamic Scheduling

The R10000 processor can issue instructions to functional units out of program order; this capability is known as **dynamic scheduling** or **dynamic issuing**.

The R10000 processor can dynamically issue an instruction as soon as all its operands are available and the required execution unit is not busy. Thus, an instruction is not delayed by a stalled previous instruction unless it needs the results of that previous instruction.

A.7 Instruction Fetch, Decode, Issue, Execution, Completion, and Graduation

In general, instructions are *fetched*, *decoded*, and *graduated* in their original program order, but may be *issued*, *executed*, and *completed* out of program order, as shown in Figure A-1.

- **Instruction fetching** is the process of reading instructions from the instruction cache.
- **Instruction decode** includes register renaming and initial dependency checks. For branch instructions, the branch path is predicted and the target address is computed.
- An instruction is **issued** when it is handed over to a functional unit for **execution**.
- An instruction is **complete** when its result has been computed and stored in a temporary physical register.
- An instruction **graduates** when this temporary result is committed as the new state of the processor. An instruction can graduate only after it and all previous instructions have been successfully completed.

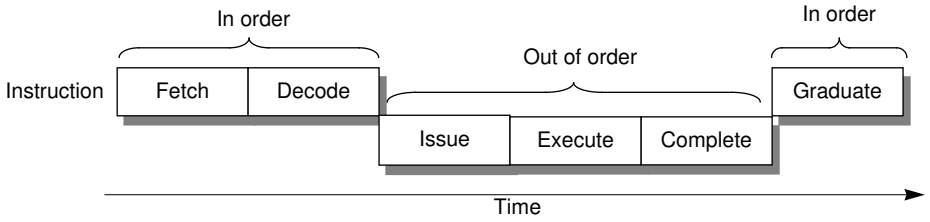

Figure A-1 Dynamic Scheduling

A.8 Active List

The R10000 processor's **active list** is a program-order list of decoded instructions. For each instruction, the active list indicates the physical register which contained the *previous* value of the destination register (if any). If this instruction graduates, that previous value is discarded and the physical register is returned to the free list. The active list records status, such as those instructions that have completed, or those instructions that have detected exceptions. Instructions are appended to the bottom of the list as they are decoded and instructions are removed from the top as they graduate.

A.9 Free List and Busy Registers

A **busy-bit table** indicates whether or not a result has been written into each of the physical registers. Each register is initially defined to be busy when it is moved from the **free list** to the active list; the register becomes available ("not busy") when its instruction completes and its result is stored in the register file.

The busy-bit table is read for each operand while an instruction is decoded, and these bits are written into the queue with the instruction. If an operand is busy, the instruction must wait in the queue until the operand is "not busy." The queues determine when an operand is ready by comparing the register number of the result coming out of each execution unit with the register number of each operand of the instructions waiting in the queue.

With a few exceptions, the integer and address queues have integer operand registers, and the floating-point queue has floating-point operand registers.

A.10 Register Renaming

As it executes instructions, the processor generates a myriad of *temporary* register results. These temporary values are stored in register files together with *permanent* values. The temporary values become new permanent values when their corresponding instructions graduate.

Register renaming is used to resolve data dependencies during the dynamic execution of instructions.

To ensure each instruction is given correct operand values, the logical register numbers (**names**) used in the instruction are mapped to physical registers. Each time a new value is put in a logical register, it is assigned to a new physical register. Thus, each physical register has only a single value. Dependencies are determined using these physical register numbers.

An example of register renaming is shown below. The following Doubleword Shift Left Logical instruction,

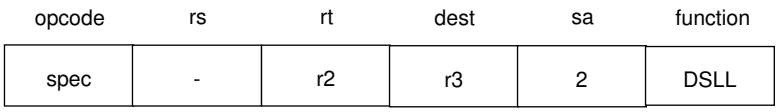


```
DSLL                r3,r2,2
```

has one register operand (*r2*) plus a 5-bit shift count of value two stored in the *sa* field; the value in *r2* is shifted left by two and this value is stored in *r3*.

The physical execution of the instruction above, with register renaming, is given below:

```
            Physical execution          Rename operation

            p3←p2 shift left 2          r3 = p3
```

When the DSLL instruction is executed, the *l*ogical destination register *r3* is assigned a new physical register, p3, from the free list.

Register renaming also allows exceptions to be handled in a precise manner. *Out-of-order execution* means that an instruction can change its result register even before all prior instructions have been completed. However, if any of the prior instructions cause an exception, the original register value must be restored. Since each new register value is loaded into a new physical register (physical register values are not overwritten until the physical register is placed in the free list), previous values remain unchanged in the original physical registers and these previous values can be restored.†

An instruction can be aborted up until the time it graduates, and all register and memory values can be restored to a precise state following any exception. This state is restored by *unnaming* the temporary physical registers assigned to subsequent instructions.

Registers are **unnamed** by writing the old destination register into the mapping table and returning the new destination register to the free list. Unnaming is done in reverse program order, in case a logical register was used more than once. After renaming, the register files contain only the permanent values which were created by instructions prior to the exception.

Once an instruction has graduated, all previous values are lost.

A.11 Nonblocking Loads and Stores

Loads and stores are **nonblocking**; that is, cache misses do not stall the processor. All other parts of the processor may continue to work on non-dependent instructions while as many as four cache misses are being processed.

† This same technique is used to reverse mispredicted speculative branches.

A.12 Speculative Branching

Normally, about one of every six instructions is a branch. Since four instructions are fetched each cycle, the R10000 processor encounters, on average, a branch instruction every other cycle, as shown in Figure A-2.

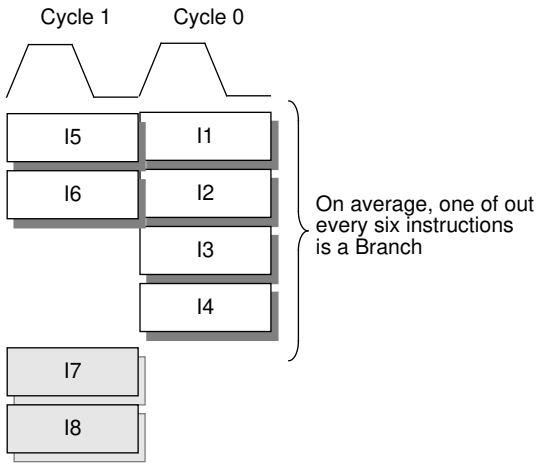

Figure A-2 Speculative Branching

When a branch instruction was encountered in previous processors, the instruction fetch and instruction issue halted until it was determined whether or not to take the branch. For instance, a branch delay slot was designed into the MIPS architecture to handle the intrinsic delay of a branch and to keep the pipeline filled.

Since the processor fetches up to four instructions each clock cycle, there is not enough time to resolve branches without stalling the fetch/decode circuitry. The processor therefore **predicts** the outcome of every branch and **speculatively executes** the branch based on this branch prediction.

The branch prediction circuit consists of a 512-entry RAM, using a 2-bit prediction scheme: two bits are assigned to a branch instruction, and indicate whether or not the branch was taken the last time it occurred. The four possible prediction states are: strongly taken, weakly taken, weakly not taken, strongly not taken. If the branch was taken the last two times, there is a good probability it will be taken this time too — or the inverse.†

The R10000 processor can speculate up to four branches deep. Shadow copies of the mapping tables are kept every time a prediction is made, allowing the R10000 processor to recover from a mispredicted branch in a single cycle.

† Simulations have shown the R10000 branch prediction algorithm to be over 90% accurate.

A.13 Logical and Physical Registers

Register renaming (described above) distinguishes between **logical registers**, which are referenced within instruction fields, and **physical registers**, which are actually located in the hardware register file. The programmer is only aware of logical registers; the implementation of physical registers is entirely transparent.

Logical register numbers are dynamically mapped onto physical register numbers. This mapping uses *mapping tables* which are updated after each instruction is decoded; each new result is written into a new physical register. This value is temporary and the previous contents of each logical register can be restored if its instruction must be aborted following an exception or a mispredicted branch.

Register renaming simplifies dependency checks. Logical register numbers can be ambiguous when instructions are executed out of order, since a succession of different values may be assigned to the same register. But physical register numbers uniquely identify each result, making dependency checking unambiguous.

The queues and execution units use physical register numbers. Integer and floating-point registers are implemented with separate renaming hardware and multi-port register files.

A.14 Register Files

The R10000 processor has two 64-bit-wide register files to store integer and floating-point values. Each file contains 64 registers. The integer register file has seven read and three write ports; the floating-point register file has five read and three write ports.

The integer and floating-point pipelines each use two dedicated operand ports and one dedicated result port in the appropriate register file. The Load/Store unit uses two dedicated integer operand ports for address calculation. It must also load or store either integer or floating-point values, sharing a result port and a read port in both register files.

These shared ports are also used to move data between the integer and floating-point register files, to store branch and link return addresses, and to read the target address for branch register instructions.

A.15 ANDES Architecture

The R10000 processor uses the MIPS ANDES architecture, or *Architecture with Non-sequential Dynamic Execution Scheduling*.

Index

Numerics

16-word, cache refill
read sequence 63
write sequence 67
32-bit
address space 293
mode, TLB entry format 305
32-word, cache refill
read sequence 63
write sequence 67
4-word, cache refill
read sequence 61
write sequence 65
599CLGA, *see* CLGA
64-bit
address space 293
mode, TLB entry format 305
8-word, cache refill
read sequence 62
write sequence 66

A

AC electrical specifications 200
asynchronous inputs 201
delay time 201
hold time 201
maximum operating conditions 200
setup time 201
test specification 200
timing
secondary cache 200
System interface 200
access privileges, address space 302
active list, definition of A-3
add unit, FPU 277
address
calculation of 19
encodings, mode 293
Kernel mode 298
mapping
Kernel mode 298
Supervisor mode 296
User mode 294
mode 293
page 304
queue 6, 12
instruction graduation 12
issue ports 12
number of entries 12
number of instructions written per cycle 12
organized as FIFO 12
sequencing 12
space
access privileges 302
kernel 293
supervisor 293
user 293
virtual 293
Supervisor mode 296
translation 306
User mode 294
Address Error exception 316
Address Space Identifier, see ASID
address/data bus signals 36
algorithms
cache, five types of 47
aliasing, virtual 59
ALU (arithmetic logic unit)
No. 1 18
No. 2 18
ALU1 9, 11
ALU2 9, 11
ANDES, Architecture with Non-sequential Dynamic Execution Scheduling 4, A-7
arbitration protocol, System interface 98
arbitration rules, System interface 99

arbitration signals 36
arbitration, cluster bus 72
Architecture with Non-sequential Dynamic Execution Scheduling, see ANDES
arithmetic instructions, FPU 286
arithmetic logic unit, see ALU
array, page table entry (PTE) 224
ASID (Address Space Identifier)
  context switch 306
  relationship to Global (G) bit in TLB entry 306
ASID (Address Space Indentifier)
  stored in EntryHi register 306
asynchronous inputs, AC electrical specification 201

B

Bad Virtual Address register (BadVAddr) 227
block
  instruction cache 9
  primary data cache 9
  secondary cache 10
  size
    primary data cache 42
    primary instruction cache 40
    secondary cache 45
block data transfers 84
  external block data responses 84
  processor block write requests 84
  processor coherency data responses 84
boundary scan register, JTAG 192
branch
  determining next address 17
  instruction, limits on execution 17
  mispredicted 14
  prediction 14, 27, A-6
  prediction rates, improving 21
  speculative A-6
  unit 10, 17
Branch on Coprocessor 0 instructions 262
Breakpoint exception 326
buffer
  cached request 79, 80
  cluster request 79, 82
  incoming 79
  outgoing 79, 80
  uncached 79, 81
bus
  SysAD 92
  SysCmd 85
  SysResp 95
  SysState 94
Bus Error exception 322
busy-bit table A-4
bypass register, JTAG 191

C

cache 4
  algorithms 47
    and processor requests 50
    cacheable coherent exclusive on write, description of 48
    cacheable coherent exclusive, description of 48
    cacheable noncoherent, description of 48
    encoding 47
    for kseg0 address space 47
    for mapped address space 47
    for xkphys address space 47
    uncached accelerated, description of 48
    uncached, description of 48
    where specified 47
  associativity 39
  block ownership 51
  misses 21
  nonblocking 21
  ordering constraints 15
  pages 304
  primary 4
  primary data 9
    block size 42
    changing states 43
    description of 42
    error handling 164
    index and tag 43
    interleaving 28
    refill 27
    state diagram 44
    states 43
    subset of secondary cache 43
    write back protocol 42
  primary instruction 9
    block size 40
    changing states 41
    description of 40
    error handling 163

primary instruction (cont.)
error protection 163
index and tag 40
refill 27
state diagram 41
states 40
rules, ownership of a cache block 51
secondary 4
associativity 10, 45
block size 45
block state 60
blocks 10
changing states 46
clock domain 144
data array 54
data array width 55
description of 45
ECC 10
error handling 165
index and tag 45
indexing the data array 55
indexing the tag array 56
interface frequencies 54
sizes 10
specifying block size 54
specifying cache size 54
state diagram 46
states 45
tag 59
tag and data array ECC 54
tag array 54
way prediction 57
write back protocol 45
strong ordering
example of 16
structure, two-level 39
Cache Barrier CACHE instruction 184
Cache Error exception 160, 321
precision 160
prioritization 160
CACHE instructions 161, 173, 264
and a hit in the cache 175
and Address Error exception 175
and CE bit 176
and CH bit 176
and CP0 174
and invalidation 176
and TLB Invalid exception 175
and TLB Refill exception 175
and Watch exception 175
and write back 175
Cache Barrier 184
Hit Writeback Invalidate 185
Index Hit Invalidate 183
Index Invalidate 178
Index Load Data 187
Index Load Tag 180, 181, 183, 184, 185, 187, 188
Index Store Data 188
Index Store Tag 181
Index Writeback Invalidate 178
op field encoding 177
serial operations 176
unsupported instructions 176
using the physical address 174
using the virtual address 174
cache miss stalls 21
Cache Operation, *see also* CACHE instructions 262
cached request buffer 79, 80
CacheErr register 161, 252
CacheOp, *see also* CACHE instructions 173
CacheOp, see CACHE instructions
capacitors, decoupling 203
cause bits, FPU 286
Cause register 236
CH, bit 262
ckseg0 space 302
ckseg1 space 302
ckseg3 space 302
cksseg space 302
CLGA (ceramic land grid array) 206
electrical characteristics 207
layout 206
mechanical characteristics 206
package 206
pinout 210
thermal characteristics 208
clock
domain
in secondary cache 144
internal processor clock domain 143
secondary cache clock domain 143
System interface clock domain 143

clock (cont.)
signal
HSTL 144
PClk 144
SCClk 144
SysClk 143
SysClkRET 144
signals, overview of 36
cluster bus 32, 72
operation 136
cluster coordinator 72
cluster request buffer 79, 82
coherency conflicts 131
coherency protocol, directory-based 140
coherency schemes 32
coherency, System interface
external intervention exclusive request 129
external intervention shared request 129
external invalidate request 129
CohPrcReqTar, mode bit 137, 139, 153
cold reset 147
sequence 150
Compare register 228
completing, an instruction A-3
completion, definition of A-5
condition bit dependencies 14
Config register 240
conflict, TLB, avoiding 306
conflicts
coherency 131
external 132
internal 131
Context register 224
context switch 306
control registers, FPU 284
controller, TAP 190
COP0 instructions (see also CP0) 333
COP1 instructions 333
COP2 instructions 333
Coprocessor 0, see CP0
Coprocessor Unusable exception 328
correctable error 158
Count register 228
CP0
branch on Coprocessor 0 instructions 262
CP0 (Coprocessor 0) 219
CP0 (coprocessor 0)
hazards 262
instructions 262
load hazards 262
move instructions 263
registers, list of 220
csseg space 297
CTM, mode bit 155

D

data cache
see also cache, primary data 42
data dependencies 20
data path, secondary cache 10
DC electrical specifications 196
input and output 199
input level sensing 198
maximum operating conditions 197
mode definitions 198
power supply levels 196
unused inputs 198
Vref, voltage reference 198
DC power supply levels 196
DC voltage, reference 198
DCOk, signal 34
DE, bit in Status register 161
decoding, an instruction A-3
decoupling capacitance 203
delay times, AC electrical 201
dependencies
condition bit 14
exception 14
instruction 13
memory 14
register 14, A-7
dependency 13
DevNum, mode bits 153
Diagnostic register 245
directory-based coherency protocol 140
divide unit, FPU 277
DMFC0, instruction 263, 267
DMTC0, instruction 263, 268

Done, bit 11
done, see also completion A-5
Doubleword Move From CP0, instruction 262
Doubleword Move To CP0, instruction 262
duplicate tags, external 30
dynamic issue 13, A-3
dynamic scheduling A-3

E

ECC (error correcting code)
  matrix for secondary cache data array 165
  matrix for secondary cache tag array 167
  matrix for System interface 170
  register 251
  secondary cache 10
efficiency, program, suggestions for increasing 21
electrical specifications
  AC 200
  DC 196
EntryHi register 229, 305
  ASID field in 306
EntryLo0 register 223, 305
EntryLo1 register 223, 305
ERET, instruction 269
ERL bit, in Status register 292
error
  correctable 158
  uncorrectable 158
    handling an 160
    limiting the propagation of 159
  units that detect and report uncorrectable errors 160
error correcting code see also ECC 162
Error Exception Program Counter (ErrorEPC) register 261
error handling 157
  protocol 171
error levels, in the Status register 292
error protection 157
  schemes used in R10000 162
error protection schemes, used in R10000
  ECC 162
  parity 162
  sparse encoding 162
EW, bit in CacheErr register 161
exception levels, in the Status register 292
exception processing, CPU
  exception types
    Address Error 316
    Breakpoint 326
    Bus Error 322
    Cache Error 160, 321
    Coprocessor Unusable 328
    Floating-Point 329
    Integer Overflow 323
    Interrupt 331
    Reserved Instruction 327
    Soft Reset 313
    System Call 325
    TLB 317
    TLB Invalid 317, 319
    TLB Modified 317, 320
    TLB Refill 317, 318
    Trap 324
    Virtual Coherency 321
    Watch 330
  exception vector location
    Reset 308
    TLB Refill 308
  exception vector selection 309
  precise handling 14
  priority of 309, 311
  TLB refill vector locations 310
Exception Program Counter (EPC) register 238
Exception Return, instruction 262
executing, an instruction A-3
execution order 13
execution pipelines 6
execution units, iterative A-7
execution, speculative 20, A-6
EXL bit, in Status register 292
external agent 30, 31, 69
  connecting to 71
external allocate request number request protocol 123
external block data response 84
  protocol 116
external coherency conflicts 132
external coherency request latency 133
external coherency requests, action taken 130
external completion response protocol 119
external double/single/partial-word data response protocol 118

external duplicate tags, support for 139
external interface 10
memory accesses 28
priority operations 28
external interrupt request 95
protocol 125
external intervention exclusive request 129
external intervention request protocol 121
external intervention shared request 129
external invalidate request 129
protocol 124
external request 70, 77
protocol 121
external response 70, 77
protocol 116

F

fetch pipeline 6, 17
fetching, an instruction A-3
FGR (Floating-Point General register)
32-bit operations 280
5-bit select 279
64-bit operations 280
load operations 281
operations 280
Status register FR bit 280
store operations 281
floating-point
adder 18
adder pipeline 6
divide 18, 278
multiplier 18
pipeline 7
queue 6, 11
instructions written each cycle 11
number of allowable entries 11
ports 11
sequencing 11
registers 279
rounding mode 287
square root 18
Floating-Point exception 329
Floating-Point Status register see FSR 285
Floating-Point Unit, see FPU
flow control 83
external data response 83
external request 83
processor coherency data response 83
processor eliminate request 83
processor read request 83
processor upgrade request 83
processor write request 83
signals 36
format, TLB entry 305
FPU 277
Active List, control of FSR 285
add unit 277
arithmetic instructions 286
changing rounding mode using a CTC1 287
compare 286
condition bits 286
control registers 284
divide unit 277
FGRs (general registers) 279
FSR cause bits 286
FSR, (Status register in FPU) 285
graduation, control of FSR 285
instructions, processor specific 288
latency 277
logic diagram 278
move to floating-point 283
multiply unit 277
operations 278
queue
controlling units 279
move unit, FPU 278
read ports 278
register file 278
repeat rate 277
rounding modes 287
serial dependency circuit 283
square-root unit 277
FrameMask register 244
free list A-4
FSR (Floating-Point Status register)
cause bits 286
condition bits 286
division by zero 286
enable bits 286
flag bits 286
inexact result 286
invalid operation 286

FSR (Floating-Point Status register) (cont.)
load exceptions 287
loading the FSR 287
overflow 286
RM, round to minus infinity 287
RN, round to nearest representable value 287
RP, round to plus infinity 287
RZ, round toward zero 287
underflow 286
unimplemented operation 286
functional unit 9
branch 10
floating-point adder 9
floating-point multiplier 9
instruction decode and rename 10
integer ALU 9
Load/Store Unit 9

G

G, bit in TLB 306
global processes (G bit in TLB) 306
graduation
definition of A-5
of an instruction A-3
Grant parking 98

H

hardware emulation, support for 141
hardware interrupts 95
hazards, CP0 262
Hit Writeback Invalidate CACHE instruction 185
hold times, AC electrical 201

I

I/O, support for 139
incoming buffer 79
Index Hit Invalidate CACHE instruction 183
Index Invalidate CACHE instruction 178
Index Load Data CACHE instruction 187
Index Load Tag CACHE instruction 180, 181, 183, 184, 185, 187, 188
Index register 221
Index Store Data CACHE instruction 188
Index Store Tag CACHE instruction 181
Index Writeback Invalidate CACHE instruction 178
initialization 147
input voltage levels, maximum 202
instruction
CACHE, *see also* CACHE instructions 161, 173, 262, 264
CacheOp, see CACHE instructions
completion 20, A-3
COP0 (see also CP0) 333
COP1 333
COP2 333
decoding A-3
dependencies 13
DMFC0 262, 267
DMFC1 286
DMTC0 262, 268
ERET 262, 269
execution A-3
fetching A-3
FPU, processor specific 288
CFC1 289
CTC1 289
CVT.L.fmt 288
for valid FP control registers 289
moves and conditional moves 289
graduation A-3
issue 20, A-3
superscalar 20
latencies 25
load linked 23
MFC0 262, 270
MFC1 286
in MIPS II ISA 283
MFPC 271
MFPS 271
MTC0 262, 272
MTPC 271
MTPS 271
prefetch 21
processor-specific 22
queue 11, 17
repeat rates 25
store conditional 23
SWC1, in MIPS II ISA 283
SYNC 24, 135
TLBP 262, 273
TLBR 262, 274
TLBWI 262, 275
TLBWR 262, 276

instruction cache, block size 40
instruction register, JTAG 191
integer
queue 11
branch instructions 11
divide instructions 11
multiply instructions 11
ports 11
shift instructions 11
integer ALU pipeline 6
Integer Overflow exception 323
integer queue 6
interface, external 10
internal coherency conflicts 131
internal processor clock domain 144
Interrupt exception 331
Interrupt register 95
interrupt request, external 95
interrupts 95
hardware 95
nonmaskable 96
software 96
timer 96
ISA (Instruction Set Architecture)
MIPS I 2
MIPS II 2
MIPS III 2
MIPS IV 2, 279
issue, dynamic A-3
issuing, an instruction A-3
iterative 9
iterative execution units A-7
ITLB (instruction TLB) 306

J

JTAG
boundary scan register 192
bypass register 191
instruction register 191
interface 189
instruction register 191
JTCK signal 190
JTDI signal 190
JTDO signal 190
JTMS signal 190
JTAG (cont.)
interface (cont.)
Tap controller 190
test access port 190
signals 38
JTCK, signal 38, 190
JTDI, signal 38
JTDO, signal 38
JTLB (joint TLB) 306
JTMS, signal 38, 190

K

Kernel mode 292
address mapping 298
ckseg0 space 302
ckseg1 space 302
ckseg3 space 302
cksseg space 302
kseg0 space 299
kseg1 space 299
kseg3 space 299
ksseg space 299
kuseg space 299
operations 298
xkphys space 300
xkseg space 302
xksseg space 300
xkuseg space 300
kseg0 space 299
Kseg0CA, mode bits 153
kseg1 space 299
kseg3 space 299
ksseg space 299
kuseg space 299
KX bit, in Status register 292

L

latency 25
accessing secondary cache 27
definition of A-2
external coherency request 133
FPU 277
least-recently used replacement algorithm (LRU) 9
level sensing, input 198
list, free A-4

LL instruction 23
LLAddr register 241
load hazards, CP0 262
load linked 23
load operations, FPU registers 281
Load/Store Unit pipeline 6
loads
  nonblocking A-5
logic diagram, FPU 278
logical register
  initialization (necessity for) 148
logical register, *see also* physical register A-7
LRU (least-recently used) replacement algorithm 9

M

mapped, virtual address region 293
mapping table A-7
matches, multiple, in TLB 306
MemEnd, mode bits 154
memory dependencies 14
memory ordering 15
memory protection 304
MFC0, instruction 263, 270
MIPS IV, instruction set *see also* ISA 332
miscellaneous system signals 37
mispredicted branch 14, 27
mode
  addressing 293
  addressing, encodings 293
    Kernel mode 293
    Supervisor mode 293
    User mode 293
  operating 292
mode bits 153
  CohPrcReqTar 137, 139, 153
  CTM 155
  DevNum 153
  Kseg0CA 153
  MemEnd 154
  ODrainSys 155
  PrcElmReq 140, 153, 184
  PrcReqMax 114, 153
  SCBlkSize 54, 154
  SCClkDiv 54, 154
  SCClkTap 144, 155
mode bits (cont.)
  SCCorEn 154, 165
  SCSize 54, 154
  SysClkDiv 70, 154
mode definitions, DC 198
Move from CP0, instruction 262
Move from performance counter, instruction 271
Move from performance event specifier, instruction 271
move instructions, CP0 263
Move to CP0, instruction 262
Move to performance counter, instruction 271
Move to performance event specifier, instruction 271
Move To/From the Performance Counter, instructions 271
MTC0, instruction 263, 272
multiple matches, in TLB 306
multiplier pipeline 6
multiply unit, FPU 277
multiprocessor system 31
  cluster bus 31
  with external agent 31

N

nonblocking cache 21
nonblocking, loads and stores A-5
nonmaskable interrupt 96
number, request 77

O

ODrainSys, mode bit 155
offset, in page address 304
op field encoding of CACHE instructions 177
operating conditions, AC 200
operating mode
  Kernel 292, 298
  Supervisor 292, 296
  User 292, 294
operations, FPU 278
ordering, memory 15
ordering, strong 15
out of program order, execution A-2
outgoing buffer 79, 80

P

package, *see* CLGA
page
  address 304
  offset 304
  size
    code 304
    defined 304
    in TLB 19
  virtual 304
page table entry (PTE) array 224
PageMask register 225, 304, 305
parity protection 162
PClk, signal 54, 70
pending response 132
performance
  branch prediction 27
  cache 27
  R10000 24, 27
Performance Counter register 247
permanent register A-4
phase-locked loop 145
physical memory addresses 304
physical register, *see also* logical register A-7
PIdx, primary cache index 59
pipeline 17
  definition of A-2
  fetch 6, 17
  floating-point 7
  floating-point multiplier 6
  integer ALU 6
  Load/Store Unit 6
  stage 1 17, 18
  stage 2 17
  stages 4-6 18
  stalls 13
PLL 145
PLLDis, signal 38
power interface signals, *see also* individual signals 34
power supply
  levels, DC 196
  regulation 202
power-on reset 147
  sequence 148
PrcElmReq, mode bit 140, 153, 184
PrcReqMax, mode bits 114, 153
precise exceptions 14
prediction, branch A-6
prefetch instruction 21
primary data cache, *see also* cache, primary data 9
primary instruction cache, *see also* cache, primary instruction 9
Probe TLB for Matching Entry, instruction 262
processor block read request protocol 102
processor block write request 84
  protocol 106
processor coherency data response 84
  protocol 127
processor coherency state response protocol 126
processor double/single/partial-word read request protocol 104
processor double/single/partial-word write request protocol 108
processor eliminate request protocol 112
processor request 70, 76
  flow control protocol 114
  protocol 102
processor response 70, 77
  protocols 126
Processor Revision Identifier (PRId) register 239
processor upgrade request protocol 110
processor-specific instructions 22
program order 13
  dynamic execution 13
  instruction completion A-3
  instruction decoding A-3
  instruction execution A-3
  instruction fetching A-3
  instruction graduation A-3
  instruction issue A-3
protection
  ECC 162
  memory 304
  parity 162
  SECDED 162
  sparse encoding 162
protocol
  arbitration, System interface 98
  error handling 171
PTE (page table entry) 224

Q

queue
address 6
instruction 17
integer 6

R

R10000 processor
ANDES architecture 4
caches 4
execution pipelines 6
geometry 4
overview 4
pipeline stages 5
superscalar pipeline 5
R4000 superpipeline 3
Random register 222
Read Indexed TLB Entry, instruction 262
read port, FPU 278
read sequences 60
16-word 63
32-word 63
4-word 61
8-word 62
tag 64
reference voltage 202
DC 198
register
Bad Virtual Address register (BadVAddr) 227
boundary scan, JTAG 192
bypass, JTAG 191
CacheErr 161, 252
Cause 236
Compare Register 228
Config 240
Context 224
Counter Register 228
CP0 219
dependency 14, A-7
Diagnostic 245
ECC 251
EntryHi 229
EntryLo0 223
EntryLo1 223
Error Exception Program Counter (ErrorEPC) register 261
Exception Program Counter 238
register (cont.)
file
FPU 278
ports A-7
FrameMask 244
Index 221
instruction, JTAG 191
LLAddr 241
logical, *see also* physical register 17, A-7
PageMask 225, 304
Performance Counter 247
permanent A-4
physical, *see also* logical register 17, A-7
Processor Revision Identifier (PRId) 239
Random 222
renaming 14, A-4
Status 161
ERL bit 292
EXL bit 292
SX bit 302
TS bit 306
USL field 292
UX bit 302
Status, *see also* register, Status 230
TagHi 255
TagLo 255
temporary A-4
unnamed A-5
WatchHi 242
WatchLo 242
Wired 226
write before reading (necessity for) 148
XContext 243
renaming, register A-4
repeat rate 25
accessing secondary cache 27
definition of A-2
FPU 277
replacement algorithm, cache 9
request cycle 70
request number 77
Reserved Instruction exception 327
reset
cold 147, 150
power-on 147, 148
soft (warm) 152
warm 147

response bus signals 37
response cycle 70
rounding modes, in FSR 287
rules, arbitration for System interface 99

S

SC instruction 23
SC(A,B)Addr(), signals 56
SC(A,B)DWay, signals 62, 66
SCAAddr, signal 35
SCADCS, signal 35
SCADOE, signal 35
SCADWay, signal 35
SCADWr, signal 35
SCBDCS, signal 35
SCBDOE, signal 35
SCBDWay, signal 35
SCBDWr, signal 35
SCBlkSize, mode bit 54
SCBlkSize, mode bits 154
SCClk, signal 35, 54
SCClkDiv, mode bits 54, 154
SCClkTap, mode bits 144, 155
SCCorEn, mode bits 154, 165
SCData, signal 35
SCDataChk, signal 35, 165
scheduling, dynamic A-3
SCSize, mode bits 54, 154
SCTag, signals 35, 59
SCTagChk, signal 35
SCTagLSBAddr, signal 35, 56
SCTCS, signal 35
SCTOE, signal 35
SCTWay, signal 35, 56, 58, 62
SCTWr, signal 35
SECDED 162
secondary cache interface signals, *see also* individual signals 35
secondary cache, *see also* cache, secondary 45
SelDVCO, signal 38
setup times, AC electrical 201
signal integrity 202
  decoupling capacitance 203
  maximum input voltage levels 202
  power supply regulation 202
  reference voltage 202
signals
  power interface, *see also* individual signals 34
  secondary cache interface, *see also* individual signals 35
  System interface, *see also* individual signals 36
  test interface, *see also* individual signals 38
size, page in memory 304
soft (warm) reset 147, 152
Soft Reset exception 313
software interrupts 96
sparse encoding protection 162
specifications, test, AC electrical 200
speculative branching A-6
speculative execution 14, 21, A-6
square-root unit, FPU 277
sseg space 297
SSRAM
  address signals 35
  clock signals 35
  data signals 35
  tag signals 35
stage, definition of A-2
stalls, improving performance 13
state bus signals 37
Status register 230
  in FPU, *see also* FSR 280
Status register, *see also* register, Status 161
store conditional 23
store operations, FPU registers 281
stores
  and uncached buffer 49
  nonblocking A-5
strong ordering 15
  example of 16
superpipeline, architecture 3
superpipeline, R4000 3
superscalar 13
  pipeline 3
  processor
    definition of 3, A-2

Supervisor mode 292
address mapping 296
csseg space 297
operations 296
sseg space 297
suseg space 296
xsseg space 297
xsuseg space 297
suseg space 296
switch, context 306
SX bit, in Status register 292, 302
SYNC
instruction 24, 135
SysAD, bus signals 36, 85, 92
SysADChk, signal 36
SysClk, signal 24, 36, 70, 143
SysClkDiv, mode bits 154
SysClkRet, signal 36
SysCmd, bus 36, 85
SysCmdPar, signal 36
SysCorErr, signal 37, 165
SysCyc, signal 37
SysGblPerf, signal 24, 37, 135
SysGnt, signal 36
SysNMI, signal 37, 96
SysRdRdy, signal 36
SysRel, signal 36
SysReq, signal 36
SysReset, signal 37
SysResp, bus 37, 85, 95
SysRespPar, signal 37
SysRespVal, signal 37
SysState, bus 37, 85, 94
SysStatePar, signal 37
SysStateVal, signal 37
System Call exception 325
system configuration
multiprocessor 31
uniprocessor 30
System interface 10, 69
arbitration
in a cluster bus system 72, 101
in a uniprocessor system 100
protocol 98
System interface (cont.)
arbitration (cont.)
rules 99
block write request protocol 106
buffers 79
bus encoding
description of buses 85
SysAD 92
SysCmd 85
SysResp 95
SysState 94
cached request buffer 80
clock domain 144
cluster bus 72
cluster request buffer 82
coherency 129
coherency conflicts, action taken 131
connecting to an external agent 71
connections to various system configurations 73
directory-based coherency protocol 140
error handling
on buses 168
on SysAD bus 169
on SysCmd bus 168
on SysResp bus 170
on SysState bus 170
schemes 168
error protection
for buses 168
schemes 168
external agent 69
external allocate request number request protocol 123
external block data response protocol 116
external coherency conflicts 132
external coherency requests, action taken 130
external completion response protocol 119
external data response flow control 83, 84
external double/single/partial-word data response protocol 118
external duplicate tags, support for 139
external interrupt request protocol 125
external intervention exclusive request 129
external intervention request protocol 121
external intervention shared request 129
external invalidate request 129
protocol 124

System interface (cont.)
external request 70, 77
flow control 83
protocol 121
external response 70, 77
protocol 116
flow control 83
frequencies 70
grant parking 98
hardware emulation, support for 141
I/O 139
incoming buffer 79
internal coherency conflicts 131
interrupts 95
master state 71
multiprocessor connections
with cluster bus 75
with dedicated external agents 74
outgoing buffer 80
outstanding processor requests 77
outstanding requests on the System interface 77
port 4
processor block read request protocol 102
processor coherency data response protocol 127
processor coherency state response protocol 126
processor double/single/partial-word read request protocol 104
processor double/single/partial-word write request protocol 108
processor eliminate request protocol 112
processor request 70, 76
flow control protocol 114
protocol 102
processor response 70, 77
protocols 126
processor upgrade request protocol 110
register-to-register operation 70
request 76
cycle 70
number field 77
protocol 102
response 76
cycle 70
protocol 102
signals 36, 71
slave state 71
split transaction 77
support for I/O 139
uncached attribute 140
System interface (cont.)
uncached buffer 81
uniprocessor connections 73
SysUncErr, signal 37
SysVal, signal 36
SysWrRdy, signal 36

T

table
busy-bit A-4
mapping A-7
tag bus, secondary cache, SCTag 59
tag read sequence 64
tag write sequence 68
TagHi register 255
TagLo register 255
tags, external, duplicate 139
TAP controller 190
TCA, signal 38
TCB, signal 38
temporary register A-4
test access port (TAP) 190
test interface signals, *see also* individual signals 38
test signals, miscellaneous 38
timer interrupt 96
TLB 19, 305
32-bit-mode entry format 305
64-bit-mode entry format 305
address
calculation 19
translation, avoiding multiple matches 306
ASID field 306
avoiding conflict 306
Cache Algorithm fields 305
entry formats 305
exceptions 317
Global (G) bit 306
ITLB 306
misses 224
multiple matches, avoiding 306
number of entries 19, 305
page size code 304
physical address 19
used with Context register 224

TLB (Translation Lookaside Buffer) 7
JTLB 306
TLB Invalid exception 317, 319
TLB Modified exception 317, 320
TLB Refill exception 317, 318
TLBP, instruction 273
TLBR, instruction 274
TLBWI, instruction 275
TLBWR, instruction 276
Translation Look-Aside Buffer, see TLB
translation, virtual address 304, 306
Trap exception 324
TriState, signal 38
TS, bit, in Status register 306
two-level cache structure 39

U

uncached
accelerated
blocks, completely gathered 49
blocks, incompletely gathered 49
stores 49
attribute, support for 140
buffer 79, 81
uncorrectable error 158
detection, suppressed 161
uniprocessor system 30
unnaming, register A-5
useg space 294, 295
User mode 292
address mapping 294
operations 294
useg space 295
xuseg space 295
UX bit, in Status register 292, 302

V

Vcc, signal 34
VccPa, signal 34
VccPd, signal 34
VccQSC, signal 34
VccQSys, signal 34
virtual address
space 293
translation 304
virtual aliasing 59
Virtual Coherency exception 321
virtual memory addresses 304
voltage
input, maximum 202
reference 202
VrefByp, signal 34
VrefSC, signal 34
VrefSys, signal 34
Vss, signal 34
VssPa, signal 34
VssPd, signal 34

W

Watch exception 330
WatchHi register 242
WatchLo register 242
Wired register 222, 226
write back protocol 39
primary data cache 42
Write Indexed TLB Entry, instruction 262
Write Random TLB Entry, instruction 262
write sequences 65
16-word 67
32-word 67
4-word 65
8-word 66
tag 68

X

XContext register 243
xkphys
decoding virtual address bits VA(61:59) 306
xkphys space 300
xkseg space 302
xksseg space 300
xkuseg space 300
xsseg space 297
xsuseg space 297
xuseg space 294, 295
XX bit, in Status register 292