

G2 PowerPC™ Core Reference Manual



Home Page:

www.freescale.com

email:

support@freescale.com

USA/Europe or Locations Not Listed:

Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
(800) 521-6274
480-768-2130

support@freescale.com

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064, Japan
0120 191014
+81 2666 8080
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
(800) 441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document. Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.





Overview	1
Register Model	2
Instruction Set Model	3
Instruction and Data Cache Operation	4
Exceptions	5
Memory Management	6
Instruction Timing	7
Signal Descriptions	8
Core Interface Operation	9
Power Management	10
Debug Features	11
PowerPC Instruction Set Listings	A
Revision History	B
Glossary of Terms and Abbreviations	GLO
Index	IND



1 Overview

2 Register Model

3 Instruction Set Model

4 Instruction and Data Cache Operation

5 Exceptions

6 Memory Management

7 Instruction Timing

8 Signal Descriptions

9 Core Interface Operation

10 Power Management

11 Debug Features

A PowerPC Instruction Set Listings

B Revision History

GLO Glossary of Terms and Abbreviations

IND Index

Contents

Paragraph Number	Title	Page Number
About This Book		
	Audience	xxxiii
	Organization.....	xxxiii
	Suggested Reading.....	xxxiv
	General Information.....	xxxiv
	Related Documentation.....	xxxv
	Conventions	xxxv
	Acronyms and Abbreviations	xxxvi
	Terminology Conventions.....	xli
 Chapter 1 Overview		
1.1	Overview.....	1-1
1.1.1	Features	1-3
1.1.2	G2_LE-Specific Features.....	1-6
1.1.2.1	True Little-Endian Mode	1-7
1.1.2.2	Critical Interrupt	1-7
1.1.2.3	Other New Signals.....	1-7
1.1.2.4	Additional Supervisor-Level SPRs.....	1-8
1.1.3	Instruction Unit.....	1-8
1.1.3.1	Instruction Queue and Dispatch Unit	1-8
1.1.3.2	Branch Processing Unit (BPU).....	1-9
1.1.4	Independent Execution Units.....	1-9
1.1.4.1	Integer Unit (IU).....	1-9
1.1.4.2	Floating-Point Unit (FPU)	1-9
1.1.4.3	Load/Store Unit (LSU)	1-10
1.1.4.4	System Register Unit (SRU).....	1-10
1.1.5	Completion Unit	1-10
1.1.6	Memory Subsystem Support.....	1-11
1.1.6.1	Memory Management Units (MMUs).....	1-11
1.1.6.2	Cache Units.....	1-12
1.1.7	Core Interface	1-13
1.1.8	System Support Functions	1-13
1.1.8.1	Power Management	1-14



Contents

Paragraph Number	Title	Page Number
1.1.8.2	Time Base/Decrementer	1-14
1.1.8.3	IEEE 1149.1 (JTAG)/COP Test Interface	1-15
1.1.8.4	Clock Multiplier.....	1-15
1.2	PowerPC Architecture Implementation	1-15
1.3	Implementation-Specific Information.....	1-16
1.3.1	Register Model.....	1-17
1.3.1.1	General-Purpose Registers (GPRs).....	1-17
1.3.1.2	Floating-Point Registers (FPRs).....	1-17
1.3.1.3	Condition Register (CR).....	1-19
1.3.1.4	Floating-Point Status and Control Register (FPSCR)	1-19
1.3.1.5	Machine State Register (MSR).....	1-19
1.3.1.6	Segment Registers (SRs)	1-19
1.3.1.7	Special-Purpose Registers (SPRs)	1-19
1.3.1.7.1	User-Level SPRs.....	1-20
1.3.1.7.2	Supervisor-Level SPRs.....	1-20
1.3.2	Instruction Set and Addressing Modes	1-22
1.3.2.1	PowerPC Instruction Set and Addressing Modes.....	1-22
1.3.2.2	Implementation-Specific Instruction Set	1-24
1.3.3	Cache Implementation	1-24
1.3.3.1	PowerPC Cache Characteristics	1-24
1.3.3.2	Implementation-Specific Cache Implementation	1-25
1.3.3.3	Instruction and Data Cache Way-Locking.....	1-26
1.3.4	Exception Model.....	1-26
1.3.4.1	PowerPC Exception Model.....	1-26
1.3.4.2	Implementation-Specific Exception Model.....	1-28
1.3.5	Memory Management.....	1-31
1.3.5.1	PowerPC Memory Management.....	1-31
1.3.5.2	Implementation-Specific Memory Management.....	1-31
1.3.6	Instruction Timing	1-32
1.3.7	System Interface	1-34
1.3.7.1	Memory Accesses.....	1-35
1.3.7.2	Signals.....	1-35
1.3.8	Debug Features (G2_LE Only).....	1-37
1.3.8.1	Instruction Address Breakpoint Registers (IABR and IABR2).....	1-37
1.3.8.2	Data Address Breakpoint Registers (DABR and DABR2)	1-38
1.3.8.3	Breakpoint Signaling	1-38
1.3.8.4	Other Debug Resources	1-38
1.4	Differences Between the MPC603e and the G2 and G2_LE Cores	1-39

Contents

Paragraph Number	Title	Page Number
---------------------	-------	----------------

Chapter 2 Register Model

2.1	Register Set	2-1
2.1.1	PowerPC Register Set	2-1
2.1.2	Implementation-Specific Registers	2-9
2.1.2.1	Hardware Implementation Register 0 (HID0)	2-10
2.1.2.2	Hardware Implementation Register 1 (HID1)	2-14
2.1.2.3	Hardware Implementation Register 2 (HID2)	2-14
2.1.2.4	Data and Instruction TLB Miss Address Registers (DMISS and IMISS)	2-16
2.1.2.5	Data and Instruction TLB Compare Registers (DCMP and ICMP)	2-16
2.1.2.6	Primary and Secondary Hash Address Registers (HASH1 and HASH2)	2-17
2.1.2.7	Required Physical Address Register (RPA)	2-17
2.1.2.8	BAT Registers (BAT4–BAT7)—G2_LE Only	2-18
2.1.2.9	Critical Interrupt Save/Restore Register 0 (CSRR0)—G2_LE Only	2-19
2.1.2.10	Critical Interrupt Save/Restore Register 1 (CSRR1)—G2_LE Only	2-19
2.1.2.11	SPRG4–SPRG7 (G2_LE Only)	2-20
2.1.2.12	System Version Register (SVR)—G2_LE Only	2-20
2.1.2.13	System Memory Base Address (MBAR)—G2_LE Only	2-20
2.1.2.14	Instruction Address Breakpoint Registers (IABR and IABR2)	2-21
2.1.2.14.1	Instruction Address Breakpoint Control Registers (IBCR)— G2_LE Only	2-21
2.1.2.15	Data Address Breakpoint Register (DABR and DABR2)— G2_LE Only	2-22
2.1.2.15.1	Data Address Breakpoint Control Registers (DBCR)— G2_LE-Only	2-24

Chapter 3 Instruction Set Model

3.1	Operand Conventions	3-1
3.1.1	Data Organization in Memory and Memory Operands	3-1
3.1.2	Endian Modes and Byte Ordering	3-1
3.1.3	Alignment and Misaligned Accesses	3-2
3.1.4	Floating-Point Execution Model	3-3
3.1.5	Effect of Operand Placement on Performance	3-4
3.2	Instruction Set Summary	3-5
3.2.1	Classes of Instructions	3-6
3.2.1.1	Definition of Boundedly Undefined	3-6

Contents

Paragraph Number	Title	Page Number
3.2.1.2	Defined Instruction Class	3-7
3.2.1.3	Illegal Instruction Class	3-7
3.2.1.4	Reserved Instruction Class	3-8
3.2.2	Addressing Modes	3-8
3.2.2.1	Memory Addressing	3-8
3.2.2.2	Memory Operands	3-9
3.2.2.3	Effective Address Calculation	3-9
3.2.2.4	Synchronization	3-10
3.2.2.4.1	Context Synchronization	3-10
3.2.2.4.2	Execution Synchronization	3-10
3.2.2.4.3	Instruction-Related Exceptions	3-10
3.2.3	Instruction Set Overview	3-11
3.2.4	PowerPC UISA Instructions	3-11
3.2.4.1	Integer Instructions	3-12
3.2.4.1.1	Integer Arithmetic Instructions	3-12
3.2.4.1.2	Integer Compare Instructions	3-13
3.2.4.1.3	Integer Logical Instructions	3-13
3.2.4.1.4	Integer Rotate and Shift Instructions	3-14
3.2.4.2	Floating-Point Instructions	3-15
3.2.4.2.1	Floating-Point Arithmetic Instructions	3-16
3.2.4.2.2	Floating-Point Multiply-Add Instructions	3-16
3.2.4.2.3	Floating-Point Rounding and Conversion Instructions	3-17
3.2.4.2.4	Floating-Point Compare Instructions	3-17
3.2.4.2.5	Floating-Point Status and Control Register Instructions	3-18
3.2.4.2.6	Floating-Point Move Instructions	3-18
3.2.4.3	Load and Store Instructions	3-19
3.2.4.3.1	Self-Modifying Code	3-19
3.2.4.3.2	Integer Load and Store Address Generation	3-19
3.2.4.3.3	Register Indirect Integer Load Instructions	3-20
3.2.4.3.4	Integer Store Instructions	3-20
3.2.4.3.5	Integer Load and Store with Byte-Reverse Instructions	3-21
3.2.4.3.6	Integer Load and Store Multiple Instructions	3-22
3.2.4.3.7	Integer Load and Store String Instructions	3-23
3.2.4.3.8	Floating-Point Load and Store Address Generation	3-24
3.2.4.3.9	Floating-Point Load Instructions	3-24
3.2.4.3.10	Floating-Point Store Instructions	3-25
3.2.4.4	Branch and Flow Control Instructions	3-25
3.2.4.4.1	Branch Instruction Address Calculation	3-26
3.2.4.4.2	Branch Instructions	3-26
3.2.4.4.3	Condition Register Logical Instructions	3-27
3.2.4.5	Trap Instructions	3-27

Contents

Paragraph Number	Title	Page Number
3.2.4.6	Processor Control Instructions.....	3-28
3.2.4.6.1	Move To/From Condition Register Instructions.....	3-28
3.2.4.7	Memory Synchronization Instructions—UISA	3-28
3.2.5	PowerPC VEA Instructions	3-30
3.2.5.1	Processor Control Instructions.....	3-30
3.2.5.2	Memory Synchronization Instructions—VEA	3-30
3.2.5.3	Memory Control Instructions—VEA	3-31
3.2.5.4	External Control Instructions.....	3-32
3.2.6	PowerPC OEA Instructions	3-32
3.2.6.1	System Linkage Instructions.....	3-33
3.2.6.2	Processor Control Instructions—OEA	3-33
3.2.6.2.1	Move To/From Machine State Register Instructions.....	3-33
3.2.6.2.2	Move To/From Special-Purpose Register Instructions.....	3-33
3.2.6.3	Memory Control Instructions—OEA	3-35
3.2.6.3.1	Supervisor-Level Cache Management Instruction	3-36
3.2.6.3.2	Segment Register Manipulation Instructions	3-36
3.2.6.3.3	Translation Lookaside Buffer Management Instructions	3-36
3.2.7	Recommended Simplified Mnemonics.....	3-37
3.2.8	Implementation-Specific Instructions.....	3-37

Chapter 4 Instruction and Data Cache Operation

4.1	Overview.....	4-1
4.2	Instruction Cache Organization and Control	4-3
4.2.1	Instruction Cache Organization	4-3
4.2.2	Instruction Cache Fill Operations	4-4
4.2.3	Instruction Cache Control.....	4-4
4.2.3.1	Instruction Cache Invalidation.....	4-4
4.2.3.2	Instruction Cache Disabling	4-5
4.2.3.3	Instruction Cache Locking.....	4-5
4.3	Data Cache Organization and Control	4-5
4.3.1	Data Cache Organization	4-5
4.3.2	Data Cache Fill Operations.....	4-6
4.3.3	Data Cache Control.....	4-6
4.3.3.1	Data Cache Invalidation	4-6
4.3.3.2	Data Cache Disabling	4-7
4.3.3.3	Data Cache Locking	4-7
4.3.3.4	Data Cache Operations and Address Broadcasts.....	4-7
4.3.4	Data Cache Touch Load Support	4-8
4.4	Basic Data Cache Operations	4-8
4.4.1	Data Cache Fill	4-8

Contents

Paragraph Number	Title	Page Number
4.4.2	Data Cache Cast-Out Operation	4-9
4.4.3	Cache Block Push Operation	4-9
4.5	Data Cache Transactions on Bus	4-9
4.5.1	Single-Beat Transactions	4-9
4.5.2	Burst Transactions	4-9
4.5.3	Access to Direct-Store Segments.....	4-10
4.6	Memory Management/Cache Access Mode Bits—W, I, M, and G.....	4-10
4.6.1	Write-Through Attribute (W)	4-11
4.6.2	Caching-Inhibited Attribute (I).....	4-12
4.6.3	Memory Coherency Attribute (M).....	4-12
4.6.4	Guarded Attribute (G).....	4-13
4.6.5	W, I, and M Bit Combinations	4-13
4.6.5.1	Out-of-Order Execution and Guarded Memory	4-14
4.6.5.2	Effects of Out-of-Order Data Accesses	4-14
4.6.5.3	Effects of Out-of-Order Instruction Fetches.....	4-15
4.7	Cache Coherency—MEI Protocol	4-15
4.7.1	MEI State Definitions	4-16
4.7.2	MEI State Diagram	4-16
4.7.3	MEI Hardware Considerations	4-17
4.7.4	Coherency Precautions	4-19
4.7.4.1	Coherency in Single-Processor Systems	4-19
4.7.5	Load and Store Coherency Summary	4-19
4.7.6	Atomic Memory References.....	4-20
4.7.7	Cache Reaction to <u>Specific Bus Operations</u>	4-20
4.7.8	Operations Causing core_artry Assertion.....	4-21
4.7.9	Enveloped High-Priority Cache Block Push Operation	4-22
4.8	Cache Control Instructions	4-22
4.8.1	Data Cache Block Invalidate (dcbi) Instruction.....	4-24
4.8.2	Data Cache Block Touch (dcbt) Instruction.....	4-24
4.8.3	Data Cache Block Touch for Store (dcbstst) Instruction.....	4-24
4.8.4	Data Cache Block Clear to Zero (dcbz) Instruction.....	4-24
4.8.5	Data Cache Block Store (dcbst) Instruction.....	4-25
4.8.6	Data Cache Block Flush (dcbf) Instruction.....	4-25
4.8.7	Enforce In-Order Execution of I/O (eieio) Instruction	4-26
4.8.8	Instruction Cache Block Invalidate (icbi) Instruction	4-26
4.8.9	Instruction Synchronize (isync) Instruction	4-26
4.9	System Bus Interface and Cache Instructions.....	4-26
4.10	Bus Interface	4-27
4.11	MEI State Transactions	4-29
4.12	Cache Locking	4-31
4.12.1	Cache Locking Terminology	4-32

Contents

Paragraph Number	Title	Page Number
4.12.2	Cache Locking Register Summary	4-32
4.12.3	Performing Cache Locking.....	4-33
4.12.3.1	Data Cache Locking	4-34
4.12.3.1.1	Enabling the Data Cache	4-34
4.12.3.1.2	Address Translation for Data Cache Locking	4-34
4.12.3.1.3	Disabling Exceptions for Data Cache Locking	4-35
4.12.3.1.4	Invalidating the Data Cache	4-36
4.12.3.1.5	Loading the Data Cache	4-37
4.12.3.1.6	Entire Data Cache Locking.....	4-37
4.12.3.1.7	Data Cache Way-Locking.....	4-37
4.12.3.1.8	Invalidating the Data Cache (Even if Locked).....	4-38
4.12.3.2	Instruction Cache Locking.....	4-38
4.12.3.2.1	Enabling the Instruction Cache.....	4-38
4.12.3.2.2	Address Translation for Instruction Cache Locking.....	4-39
4.12.3.2.3	Disabling Exceptions for Instruction Cache Locking.....	4-40
4.12.3.2.4	Preloading Instructions into the Instruction Cache.....	4-40
4.12.3.2.5	Entire Instruction Cache Locking.....	4-42
4.12.3.2.6	Instruction Cache Way-Locking.....	4-42
4.12.3.2.7	Invalidating the Instruction Cache (Even if Locked)	4-43

Chapter 5 Exceptions

5.1	Exception Classes	5-2
5.1.1	Exception Priorities.....	5-6
5.1.2	Summary of Front-End Exception Handling.....	5-8
5.2	Exception Processing	5-9
5.2.1	Exception Processing Registers	5-9
5.2.1.1	SRR0 and SRR1 Bit Settings.....	5-9
5.2.1.2	CSRR0 and CSRR1 Bit Settings—G2_LE Only.....	5-11
5.2.1.3	SPRG4–SPRG7 (G2_LE Only).....	5-11
5.2.1.4	MSR Bit Settings	5-12
5.2.2	Enabling and Disabling Exceptions.....	5-14
5.2.3	Steps for Exception Processing.....	5-15
5.2.4	Setting MSR[RI].....	5-16
5.2.5	Returning From an Exception Handler with rfi	5-16
5.2.6	Returning From an Interrupt with rfci	5-16
5.3	Process Switching	5-17
5.4	Exception Latencies	5-17
5.5	Exception Definitions	5-18
5.5.1	Reset Exceptions (0x00100).....	5-19
5.5.1.1	Hard Reset and Power-On Reset	5-19

Contents

Paragraph Number	Title	Page Number
5.5.1.2	Soft Reset.....	5-21
5.5.1.3	Byte Ordering Considerations for G2_LE Only.....	5-21
5.5.2	Machine Check Exception (0x00200)	5-22
5.5.2.1	Machine Check Exception Enabled (MSR[ME] = 1).....	5-23
5.5.2.2	Checkstop State (MSR[ME] = 0)	5-24
5.5.3	DSI Exception (0x00300)	5-24
5.5.4	ISI Exception (0x00400).....	5-27
5.5.5	External Interrupt (0x00500)	5-27
5.5.6	Alignment Exception (0x00600)	5-28
5.5.6.1	Integer Alignment Exceptions	5-29
5.5.6.2	Load/Store Multiple Alignment Exceptions	5-30
5.5.7	Program Exception (0x00700).....	5-31
5.5.7.1	IEEE Floating-Point Exception Program Exceptions.....	5-31
5.5.7.2	Illegal, Reserved, and Unimplemented Instructions Program Exceptions.....	5-32
5.5.8	Floating-Point Unavailable Exception (0x00800)	5-32
5.5.9	Decrementer Exception (0x00900).....	5-32
5.5.10	Critical Interrupt Exception (0x00A00)—G2_LE Only.....	5-33
5.5.11	System Call Exception (0x00C00)	5-34
5.5.12	Trace Exception (0x00D00).....	5-34
5.5.12.1	Single-Step Instruction Trace Mode	5-35
5.5.12.2	Branch Trace Mode	5-36
5.5.13	Instruction TLB Miss Exception (0x01000).....	5-36
5.5.14	Data TLB Miss on Load Exception (0x01100).....	5-36
5.5.15	Data TLB Miss on Store Exception (0x01200)	5-37
5.5.16	Instruction Address Breakpoint Exception (0x01300)	5-37
5.5.17	System Management Interrupt (0x01400).....	5-39

Chapter 6 Memory Management

6.1	MMU Features.....	6-2
6.1.1	Memory Addressing	6-3
6.1.2	MMU Organization.....	6-3
6.1.3	Address Translation Mechanisms	6-8
6.1.4	Memory Protection Facilities.....	6-10
6.1.5	Page History Information.....	6-11
6.1.6	General Flow of MMU Address Translation	6-11
6.1.6.1	Real Addressing Mode and Block Address Translation Selection	6-11
6.1.6.2	Page Address Translation Selection	6-12
6.1.7	MMU Exceptions Summary	6-14
6.1.8	MMU Instructions and Register Summary.....	6-17

Contents

Paragraph Number	Title	Page Number
6.2	Real Addressing Mode.....	6-19
6.3	Block Address Translation.....	6-20
6.4	Memory Segment Model	6-21
6.4.1	Page History Recording.....	6-21
6.4.1.1	Referenced Bit	6-22
6.4.1.2	Changed Bit	6-22
6.4.1.3	Scenarios for Referenced and Changed Bit Recording	6-23
6.4.2	Page Memory Protection	6-24
6.4.3	TLB Description	6-25
6.4.3.1	TLB Organization.....	6-25
6.4.3.2	TLB Entry Invalidation.....	6-26
6.4.4	Page Address Translation Summary.....	6-27
6.5	Page Table Search Operation	6-27
6.5.1	Page Table Search Operation—Conceptual Flow	6-27
6.5.2	Implementation-Specific Table Search Operation.....	6-31
6.5.2.1	Resources for Table Search Operations.....	6-32
6.5.2.1.1	Data and Instruction TLB Miss Address Registers (DMISS and IMISS).....	6-34
6.5.2.1.2	Data and Instruction TLB Compare Registers (DCMP and ICMP).....	6-34
6.5.2.1.3	Primary and Secondary Hash Address Registers (HASH1 and HASH2).....	6-35
6.5.2.1.4	Required Physical Address Register (RPA)	6-35
6.5.2.2	Software Table Search Operation	6-36
6.5.2.2.1	Flow for Example Exception Handlers	6-37
6.5.2.2.2	Code for Example Exception Handlers	6-42
6.5.3	Page Table Updates.....	6-48
6.5.4	Segment Register Updates	6-48

Chapter 7 Instruction Timing

7.1	Terminology and Conventions.....	7-1
7.2	Instruction Timing Overview.....	7-3
7.3	Timing Considerations.....	7-8
7.3.1	General Instruction Flow	7-8
7.3.2	Instruction Fetch Timing.....	7-9
7.3.2.1	Cache Arbitration.....	7-10
7.3.2.2	Cache Hit	7-10
7.3.2.3	Cache Miss.....	7-13
7.3.3	Instruction Dispatch and Completion Considerations	7-13
7.3.3.1	Rename Register Operation	7-15
7.3.3.2	Instruction Serialization.....	7-15

Contents

Paragraph Number	Title	Page Number
7.3.3.3	Execution Unit Considerations	7-16
7.4	Execution Unit Timings	7-16
7.4.1	Branch Processing Unit Execution Timing	7-16
7.4.1.1	Branch Folding	7-17
7.4.1.2	Static Branch Prediction	7-18
7.4.1.2.1	Predicted Branch Timing Examples	7-19
7.4.2	Integer Unit Execution Timing	7-20
7.4.3	Floating-Point Unit Execution Timing	7-21
7.4.4	Load/Store Unit Execution Timing	7-21
7.4.5	System Register Unit Execution Timing	7-21
7.5	Memory Performance Considerations	7-22
7.5.1	Copy-Back Mode	7-22
7.5.2	Write-Through Mode	7-23
7.5.3	Cache-Inhibited Accesses	7-23
7.6	Instruction Scheduling Guidelines	7-23
7.6.1	Branch, Dispatch, and Completion Unit Resource Requirements	7-24
7.6.1.1	Branch Resolution Resource Requirements	7-24
7.6.1.2	Dispatch Unit Resource Requirements	7-25
7.6.1.3	Completion Unit Resource Requirements	7-25
7.7	Instruction Latency Summary	7-26

Chapter 8 Signal Descriptions

8.1	Signal Groupings	8-1
8.2	Signal Configurations	8-3
8.2.1	Functional Groupings	8-3
8.2.2	Input/Output Enable and High-Impedance Control Signals	8-3
8.2.2.1	Unidirectional/Bidirectional Signals	8-5
8.2.2.2	Logic Gate Equivalent and Bidirectional Signals	8-5
8.2.3	Signal Summary	8-6
8.3	Signal Descriptions	8-10
8.3.1	Address Bus Arbitration Signals	8-11
8.3.1.1	Bus Request (<u>core_br</u>)—Output	8-11
8.3.1.2	Bus Grant (<u>core_bg</u>)—Input	8-11
8.3.1.3	Address Bus Busy	8-12
8.3.1.3.1	Address Bus Busy In (<u>core_abb_in</u>)	8-12
8.3.1.3.2	Address Bus Busy Out (<u>core_abb_out</u>)	8-13
8.3.1.3.3	Address Bus Busy Output Enable (<u>core_abb_oe</u>)—Output	8-13
8.3.1.3.4	Address Bus Busy High-Impedance Enable (<u>core_abb_tre</u>)—Input	8-14
8.3.2	Address Transfer Start Signals	8-14
8.3.2.1	Transfer Start	8-14

Contents

Paragraph Number	Title	Page Number
8.3.2.1.1	Transfer Start In (<u>core_ts_in</u>).....	8-14
8.3.2.1.2	Transfer Start Out (<u>core_ts_out</u>).....	8-15
8.3.3	Address Transfer Signals	8-15
8.3.3.1	Address Bus	8-15
8.3.3.1.1	Address Bus In (<u>core_a_in</u> [0:31])	8-15
8.3.3.1.2	Address Bus Out (<u>core_a_out</u> [0:31]).....	8-15
8.3.3.1.3	Address Bus Output Enable (<u>core_a_oe</u>)—Output	8-16
8.3.3.1.4	Address Bus High-Impedance Enable (<u>core_a_tre</u>)—Input.....	8-16
8.3.3.2	Address Bus Parity	8-17
8.3.3.2.1	Address Bus Parity In (<u>core_ap_in</u> [0:3]).....	8-17
8.3.3.2.2	Address Bus Parity Input Enable (<u>core_ap_ien</u>)—Output	8-17
8.3.3.2.3	Address Bus Parity Out (<u>core_ap_out</u> [0:3])	8-17
8.3.3.3	Address Parity Error (<u>core_ape</u>)—Output	8-18
8.3.3.3.1	Address Parity Error Output Enable (<u>core_ape_oe</u>)—Output.....	8-18
8.3.3.3.2	Address Parity Error High-Impedance Enable (<u>core_ape_tre</u>)— Input.....	8-19
8.3.4	Address Transfer Attribute Signals.....	8-19
8.3.4.1	Transfer Type	8-19
8.3.4.1.1	Transfer Type In (<u>core_tt_in</u> [0:4])	8-20
8.3.4.1.2	Transfer Type Out (<u>core_tt_out</u> [0:4])	8-21
8.3.4.2	Transfer Size (<u>core_tsiz</u> [0:2])—Output	8-22
8.3.4.3	Transfer Burst	8-23
8.3.4.3.1	Transfer Burst In (<u>core_tbst_in</u>)	8-23
8.3.4.3.2	Transfer Burst Out (<u>core_tbst_out</u>).....	8-23
8.3.4.4	Transfer Code (<u>core_tc</u> [0:1])—Output	8-24
8.3.4.5	Cache Inhibit (<u>core_ci</u>)—Output	8-24
8.3.4.6	Write-Through (<u>core_wt</u>)—Output.....	8-24
8.3.4.7	Global Signals.....	8-25
8.3.4.7.1	Global In (<u>core_gbl_in</u>)	8-25
8.3.4.7.2	Global Out (<u>core_gbl_out</u>).....	8-25
8.3.4.8	Cache Set Entry (<u>core_cse</u> [0:1])—Output.....	8-25
8.3.5	Address Transfer Termination Signals.....	8-26
8.3.5.1	Address Acknowledge (<u>core_aack</u>)—Input.....	8-26
8.3.5.2	Address Retry	8-26
8.3.5.2.1	Address Retry In (<u>core_artry_in</u>).....	8-26
8.3.5.2.2	Address Retry Out (<u>core_artry_out</u>)	8-27
8.3.5.2.3	Address Retry Output Enable (<u>core_artry_oe</u>)—Output	8-28
8.3.5.2.4	Address Retry High-Impedance Enable (<u>core_artry_tre</u>)—Input	8-28
8.3.6	Data Bus Arbitration Signals	8-29
8.3.6.1	Data Bus Grant (<u>core_dbg</u>)—Input	8-29
8.3.6.2	Data Bus Write Only (<u>core_dbwo</u>)—Input.....	8-29

Contents

Paragraph Number	Title	Page Number
8.3.6.3	Data Bus Busy	8-30
8.3.6.3.1	Data Bus Busy In (<u>core_dbb_in</u>)	8-30
8.3.6.3.2	Data Bus Busy Out (<u>core_dbb_out</u>).....	8-30
8.3.6.3.3	Data Bus Busy Output Enable (<u>core_dbb_oe</u>)—Output.....	8-31
8.3.6.3.4	Data Bus Busy High-Impedance Enable (<u>core_dbb_tre</u>)—Input	8-31
8.3.7	Data Transfer Signals.....	8-31
8.3.7.1	Data Bus	8-32
8.3.7.1.1	Data Bus In (<u>core_dh_in</u> [0:31], <u>core_dl_in</u> [0:31]).....	8-32
8.3.7.1.2	Data Bus Input Enable (<u>core_dh_ien</u> , <u>core_dl_ien</u>)—Output	8-32
8.3.7.1.3	Data Bus Out (<u>core_dh_out</u> [0:31], <u>core_dl_out</u> [0:31])—Output	8-33
8.3.7.1.4	Data Bus Output Enable (<u>core_d_oe</u>)—Output.....	8-33
8.3.7.1.5	Data Bus High-Impedance Enable (<u>core_d_tre</u>)—Input	8-34
8.3.7.2	Data Bus Parity (DP[0:7])	8-34
8.3.7.2.1	Data Bus Parity In (<u>core_dp_in</u> [0:7])	8-34
8.3.7.2.2	Data Bus Parity Input Enable (<u>core_dp_ien</u>)—Output.....	8-35
8.3.7.2.3	Data Bus Parity Out (<u>core_dp_out</u> [0:7])	8-35
8.3.7.3	Data Parity Error (<u>core_dpe</u>)—Output	8-35
8.3.7.3.1	Data Parity Error Output Enable (<u>core_dpe_oe</u>)—Output	8-36
8.3.7.3.2	Data Parity Error High-Impedance Enable (<u>core_dpe_tre</u>)—Input.....	8-36
8.3.7.4	Data Bus Disable (<u>core_dbdis</u>)—Input.....	8-36
8.3.8	Data Transfer Termination Signals	8-37
8.3.8.1	Transfer Acknowledge (<u>core_ta</u>)—Input.....	8-37
8.3.8.2	Data Retry (<u>core_drtry</u>)—Input	8-38
8.3.8.3	Transfer Error Acknowledge (<u>core_tea</u>)—Input.....	8-38
8.3.9	Interrupt and Checkstop Signals	8-39
8.3.9.1	External Interrupt (<u>core_int</u>)—Input.....	8-39
8.3.9.2	Critical Interrupt (<u>core_cint</u>)—Input: G2_LE Core-Only	8-39
8.3.9.3	System Management Interrupt (<u>core_smi</u>)—Input.....	8-40
8.3.9.4	Machine Check Interrupt (<u>core_mcp</u>)—Input	8-40
8.3.9.5	Checkstop Signals.....	8-41
8.3.9.5.1	Checkstop Input (<u>core_ckstp_in</u>).....	8-41
8.3.9.5.2	Checkstop Output (<u>core_ckstp_out</u>)	8-41
8.3.9.5.3	Checkstop Output Enable (<u>core_ckstp_oe</u>)—Output	8-42
8.3.9.5.4	Checkstop High-Impedance Enable (<u>core_ckstp_tre</u>)—Input.....	8-42
8.3.10	Reset Signals.....	8-42
8.3.10.1	Hard Reset (<u>core_hreset</u>)—Input.....	8-42
8.3.10.2	Soft Reset (<u>core_sreset</u>)—Input.....	8-43
8.3.10.3	Reset Configuration Signals	8-43
8.3.10.3.1	32-Bit Mode (<u>core_32bitmode</u>)—Input.....	8-43
8.3.10.3.2	Reduced Pinout Mode (<u>core_redpinmode</u>)—Input	8-44
8.3.10.3.3	MSR IP Bit Set Mode (<u>core_msrip</u>)—Input.....	8-44

Contents

Paragraph Number	Title	Page Number
8.3.10.3.4	DRTRY Mode (<code>core_drtrymode</code>)—Input	8-44
8.3.10.3.5	True Little-Endian Mode (<code>core_tle</code>)—Input	8-45
8.3.10.3.6	System Version Register (<code>core_svr[0:31]</code>)—Input	8-45
8.3.11	Processor Status Signals	8-45
8.3.11.1	Quiescent Acknowledge (<code>core_qack</code>)—Input.....	8-45
8.3.11.2	Quiescent Request (<code>core_qreq</code>)—Output.....	8-46
8.3.11.3	Reservation (<code>core_rsrv</code>)—Output	8-46
8.3.11.4	Time Base Enable (<code>core_tben</code>)—Input	8-46
8.3.11.5	TLBI Sync (<code>core_tlbisync</code>)—Input.....	8-46
8.3.11.5.1	Output Enable (<code>core_outputs_oe</code>)—Output.....	8-47
8.3.12	COP/Scan Interface.....	8-47
8.3.12.1	JTAG Test Clock (<code>core_tck</code>)—Input.....	8-48
8.3.12.2	JTAG Test Data Input (<code>core_tdi</code>)—Input	8-48
8.3.12.3	JTAG Test Data Output (<code>core_tdo</code>)—Output.....	8-49
8.3.12.3.1	JTAG Test Data Output Enable (<code>core_tdo_oe</code>)—Output.....	8-49
8.3.12.4	JTAG Test Mode Select (<code>core_tms</code>)—Input	8-49
8.3.12.5	JTAG Test Reset (<code>core_trst</code>)—Input	8-49
8.3.12.6	TLM TAP Enable (<code>core_tap_en</code>)—Input.....	8-50
8.3.12.7	Test Linking Module Select (<code>core_tlmselect</code>)—Output	8-50
8.3.13	Test Interface.....	8-50
8.3.13.1	Disable (<code>core_disable</code>)—Input.....	8-51
8.3.13.2	LSSD Test Clock (<code>core_l1_tstclk</code> , <code>core_l2_tstclk</code>)—Input.....	8-51
8.3.13.3	LSSD Test Control (<code>core_lssd_mode</code>)—Input.....	8-51
8.3.14	Debug Control Signals.....	8-51
8.3.14.1	Instruction Address Breakpoint Register Watchpoint (<code>core_iabr</code>)—Output	8-52
8.3.14.2	Instruction Address Breakpoint Register Watchpoint (<code>core_iabr2</code>)—Output	8-52
8.3.14.3	Data Address Breakpoint Register Watchpoint (<code>core_dabr</code>)—Output	8-52
8.3.14.4	Data Address Breakpoint Register Watchpoint (<code>core_dabr2</code>)—Output	8-53
8.3.15	Clock Signals	8-53
8.3.15.1	System Clock (<code>core_sysclk</code>)—Input.....	8-53
8.3.15.2	Test Clock Output (<code>core_clk_out</code>)	8-54
8.3.15.3	PLL Configuration (<code>core_pll_cfg[0:4]</code>)—Input.....	8-55

Chapter 9 Core Interface Operation

9.1	Overview.....	9-1
9.1.1	Operation of the Instruction and Data Caches	9-2
9.1.2	Operation of the System Interface	9-4
9.1.3	Optional 32-Bit Data Bus Mode	9-5

Contents

Paragraph Number	Title	Page Number
9.1.4	Direct-Store Accesses	9-5
9.2	Memory Access Protocol.....	9-5
9.2.1	Arbitration Signals.....	9-6
9.2.2	Address Pipelining and Split-Bus Transactions.....	9-7
9.2.3	Timing Diagram Conventions.....	9-8
9.3	Address Bus Tenure	9-9
9.3.1	Address Bus Arbitration	9-9
9.3.2	Address Transfer.....	9-11
9.3.2.1	Address Bus Parity	9-12
9.3.2.2	Address Transfer Attribute Signals.....	9-12
9.3.2.2.1	Transfer Type (core_tt_in[0:4], core_tt_out[0:4]) Signals	9-13
9.3.2.2.2	Transfer Size (core_tsiz[0:2]) Signals	9-13
9.3.2.3	Burst Ordering During Data Transfers	9-14
9.3.2.4	Effect of Alignment in Data Transfers (64-Bit Bus)	9-14
9.3.2.5	Effect of Alignment in Data Transfers (32-Bit Bus)	9-16
9.3.2.5.1	Alignment of External Control Instructions	9-18
9.3.2.6	Transfer Code (core_tc[0:1]) Signals	9-19
9.3.3	Address Transfer Termination	9-19
9.4	Data Bus Tenure.....	9-21
9.4.1	Data Bus Arbitration.....	9-21
9.4.1.1	Using the core_dbb_out Signal.....	9-22
9.4.2	Data Bus Write Only.....	9-23
9.4.3	Data Transfer	9-23
9.4.4	Data Transfer Termination.....	9-24
9.4.4.1	Normal Single-Beat Termination.....	9-25
9.4.4.2	Normal Burst Termination.....	9-26
9.4.4.3	Data Transfer Termination Due to a Bus Error.....	9-27
9.4.5	Memory Coherency—MEI Protocol	9-29
9.5	Timing Examples	9-31
9.6	Optional Bus Configurations	9-37
9.6.1	32-Bit Data Bus Mode	9-37
9.6.2	No-core_drtry Mode	9-39
9.6.3	Reduced-Pinout Mode	9-40
9.7	Interrupt, Checkstop, and Reset Signals	9-40
9.7.1	External Interrupts	9-41
9.7.2	Checkstops.....	9-41
9.7.3	Reset Inputs.....	9-41
9.7.4	Core Quiesce Control Signals.....	9-41
9.8	Processor State Signals	9-42
9.8.1	Support for the lwarx/stwex . Instruction Pair	9-42
9.8.2	core_tlbisync Input	9-42

Contents

Paragraph Number	Title	Page Number
9.9	IEEE 1149.1-Compliant Interface.....	9-42
9.9.1	IEEE 1149.1 Interface Description.....	9-42
9.10	Using core_dbwo (Data Bus Write Only).....	9-43

Chapter 10 Power Management

10.1	Overview.....	10-1
10.2	Dynamic Power Management.....	10-1
10.3	Programmable Power Modes.....	10-2
10.3.1	Power Management Modes	10-3
10.3.1.1	Full-Power Mode with DPM Disabled.....	10-3
10.3.1.2	Full-Power Mode with DPM Enabled.....	10-3
10.3.1.3	Doze Mode.....	10-3
10.3.1.4	Nap Mode	10-4
10.3.1.5	Sleep Mode	10-5
10.3.2	Power Management Software Considerations.....	10-6
10.4	Example Code Sequence for Entering Processor Sleep Mode	10-6

Chapter 11 Debug Features

11.1	Breakpoint Facilities	11-1
11.1.1	Instruction Address Breakpoint Registers (IABR, IABR2).....	11-1
11.1.2	Instructional Address Control Register (IBCR).....	11-2
11.1.3	Data Address Breakpoint Registers (DABR, DABR2)	11-2
11.1.4	Data Address Control Register (DBCR).....	11-3
11.1.5	Other Debug Resources	11-3
11.1.6	Software Debug Features.....	11-3
11.2	Expanded Debugging Facilities in Breakpoint Registers	11-4
11.2.1	Breakpoint Enabled.....	11-4
11.2.2	Single-Step Enabled.....	11-4
11.2.3	Branch Trace Enabled.....	11-5
11.2.4	Address Matching	11-5
11.2.5	Combinational Matching	11-5
11.3	Watchpoint Signaling.....	11-5
11.4	Exception Vectors and Priority	11-6
11.5	Instruction Address Breakpoint Examples	11-6
11.6	Synchronization Requirements.....	11-8



Contents

Paragraph Number	Title	Page Number
------------------	-------	-------------

Appendix A PowerPC Instruction Set Listings

A.1	Instructions Sorted by Mnemonic.....	A-1
A.2	Instructions Sorted by Opcode.....	A-8
A.3	Instructions Grouped by Functional Categories	A-15
A.4	Instructions Sorted by Form	A-25
A.5	Instruction Set Legend	A-36

Appendix B Revision History

B.1	Revision Changes From Revision 0 to Revision 1	B-1
-----	--	-----

Glossary of Terms and Abbreviations

Index

Figures

Figure Number	Title	Page Number
1-1	G2 Core Block Diagram	1-2
1-2	Programming Model—Registers	1-18
1-3	Data Cache Organization	1-25
1-4	System Interface	1-34
2-1	Programming Model—Registers	2-3
2-2	Hardware Implementation Register 0 (HID0)	2-10
2-3	Hardware Implementation Register 1 (HID1)	2-14
2-4	Hardware Implementation-Dependent Register 2 (HID2).....	2-15
2-5	DMISS and IMISS Registers	2-16
2-6	DCMP and ICMP Registers.....	2-16
2-7	HASH1 and HASH2 Registers	2-17
2-8	Required Physical Address Register (RPA).....	2-18
2-9	Upper BAT Register.....	2-18
2-10	Lower BAT Register	2-19
2-11	Critical Interrupt Save/Restore Register 0 (CSRR0)	2-19
2-12	Critical Interrupt Save/Restore Register 1 (CSRR1)	2-19
2-13	SPRG0–SPRG7 Registers.....	2-20
2-14	Instruction Address Breakpoint Registers (IABR and IABR2).....	2-21
2-15	Instruction Address Breakpoint Control Register (IBCR).....	2-22
2-16	Data Address Breakpoint Registers (DABR and DABR2)	2-22
2-17	Data Address Breakpoint Control Register (DBCR).....	2-24
4-1	Instruction Cache Organization	4-4
4-2	Data Cache Organization	4-6
4-3	Double-Word Address Ordering—Critical-Double-Word-First	4-10
4-4	MEI Cache Coherency Protocol—State Diagram (WIM = 001).....	4-17
4-5	Bus Interface Address Buffers	4-28
5-1	Machine Status Save/Restore Register 0 (SSR0).....	5-9
5-2	Machine Status Save/Restore Register 1 (SSR1).....	5-10
5-3	Critical Interrupt Save/Restore Register 0 (CSRR0)	5-11
5-4	Critical Interrupt Save/Restore Register 1 (CSRR1)	5-11
5-5	Special-Purpose Registers (SPRG0–SPRG7)	5-12
5-6	Machine State Register (MSR)	5-12
6-1	MMU Conceptual Block Diagram—32-Bit Implementations.....	6-5
6-2	G2 Core IMMU Block Diagram	6-6
6-3	G2 Core DMMU Block Diagram	6-7

Figures

Figure Number	Title	Page Number
6-4	Address Translation Types	6-9
6-5	General Flow of Address Translation (Real Addressing Mode and Block)	6-12
6-6	General Flow of Page and Direct-Store Interface Address Translation.....	6-13
6-7	Segment Register and TLB Organization	6-26
6-8	Page Address Translation Flow for 32-Bit Implementations—TLB Hit.....	6-28
6-9	Primary Page Table Search—Conceptual Flow.....	6-30
6-10	Secondary Page Table Search Flow—Conceptual Flow	6-31
6-11	DMISS and IMISS Registers	6-34
6-12	DCMP and ICMP Registers.....	6-34
6-13	HASH1 and HASH2 Registers	6-35
6-14	Required Physical Address (RPA) Register.....	6-36
6-15	Flow for Example Software Table Search Operation	6-38
6-16	Check and Set R and C Bit Flow	6-39
6-17	Page Fault Setup Flow	6-40
6-18	Setup for Protection Violation Exceptions.....	6-41
7-1	Pipelined Execution Unit	7-4
7-2	Instruction Flow Diagram	7-5
7-3	G2 Core Processor Pipeline Stages.....	7-7
7-4	Instruction Timing—Cache Hit	7-11
7-5	Instruction Timing—Cache Miss.....	7-14
7-6	Branch Instruction Timing	7-20
8-1	Functional Signal Groups	8-3
8-2	Logic Diagram for Bidirectional Signals.....	8-5
8-3	Detailed Signal Groups	8-10
8-4	IEEE 1149.1-Compliant Boundary Scan Interface	8-48
9-1	G2 Core Block Diagram	9-3
9-2	Overlapping Tenures on the Bus for a Single-Beat Transfer	9-5
9-3	Address Bus Arbitration	9-10
9-4	Address Bus Arbitration Showing Bus Parking	9-11
9-5	Address Bus Transfer.....	9-12
9-6	Snooped Address Cycle with <u>core_artry_out</u>	9-21
9-7	Data Bus Arbitration.....	9-22
9-8	Normal Single-Beat Read Termination.....	9-25
9-9	Normal Single-Beat Write Termination.....	9-26
9-10	Normal Burst Transaction.....	9-26
9-11	Termination with <u>DRTRY</u>	9-27
9-12	Read Burst with <u>core_ta</u> Wait States and <u>core_drtry</u>	9-28
9-13	MEI Cache Coherency Protocol—State Diagram (WIM = 001).....	9-30
9-14	Fastest Single-Beat Reads.....	9-31
9-15	Fastest Single-Beat Writes	9-32
9-16	Single-Beat Reads Showing Data-Delay Controls	9-33

Figures

Figure Number	Title	Page Number
9-17	Single-Beat Writes Showing Data-Delay Controls.....	9-34
9-18	Burst Transfers with Data-Delay Controls.....	9-35
9-19	Use of Transfer Error Acknowledge ($\overline{\text{TEA}}$)	9-36
9-20	32-Bit Data Bus Transfer (8-Beat Burst)	9-38
9-21	<u>32-Bit Data Bus Transfer (Two-Beat Burst with $\overline{\text{DRTRY}}$)</u>	9-39
9-22	<u>core_dbwo</u> Transaction	9-44



Figures

**Figure
Number**

Title

**Page
Number**



Tables

Table Number	Title	Page Number
i	Acronyms and Abbreviated Terms.....	xxxvi
ii	Terminology Conventions.....	xli
iii	Instruction Field Conventions.....	xlii
1-1	Endian Mode Indication.....	1-7
1-2	Critical Interrupt Enabling Bit.....	1-7
1-3	Exception Classifications.....	1-28
1-4	Exceptions and Conditions.....	1-29
1-5	Other Debug and Support Register Bits.....	1-39
1-6	Differences Between G2 and G2_LE Cores.....	1-39
2-1	PVR Field Descriptions.....	2-4
2-2	Architectural PVR Field Descriptions.....	2-5
2-3	Assigned PVR Values.....	2-5
2-4	MSR Bit Settings.....	2-6
2-5	HID0 Bit Functions.....	2-11
2-6	HID0[SBCLK] and HID0[ECLK] core_clk_out Configuration.....	2-14
2-7	HID1 Bit Settings.....	2-14
2-8	HID2 Bit Descriptions.....	2-15
2-9	DCMP and ICMP Bit Settings.....	2-17
2-10	HASH1 and HASH2 Bit Settings.....	2-17
2-11	RPA Bit Settings.....	2-18
2-12	System Version Register (SVR) Bit Settings.....	2-20
2-13	Instruction Address Breakpoint Register (IABR and IABR2) Bit Settings.....	2-21
2-14	Instruction Address Breakpoint Control Registers (IBCR).....	2-22
2-15	Data Address Breakpoint Registers (DABR and DABR2) Bit Settings.....	2-23
2-16	Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only.....	2-24
3-1	Endian Mode Indication.....	3-2
3-2	Memory Operands.....	3-3
3-3	Integer Arithmetic Instructions.....	3-12
3-4	Integer Compare Instructions.....	3-13
3-5	Integer Logical Instructions.....	3-14
3-6	Integer Rotate Instructions.....	3-15
3-7	Integer Shift Instructions.....	3-15
3-8	Floating-Point Arithmetic Instructions.....	3-16
3-9	Floating-Point Multiply-Add Instructions.....	3-16
3-10	Floating-Point Rounding and Conversion Instructions.....	3-17



Tables

Table Number	Title	Page Number
3-11	Floating-Point Compare Instructions	3-18
3-13	Floating-Point Move Instructions.....	3-18
3-14	Integer Load Instructions.....	3-20
3-15	Integer Store Instructions.....	3-21
3-16	Integer Load and Store with Byte-Reverse Instructions.....	3-22
3-17	Integer Load and Store Multiple Instructions.....	3-23
3-18	Integer Load and Store String Instructions.....	3-23
3-19	Floating-Point Load Instructions.....	3-24
3-20	Floating-Point Store Instructions.....	3-25
3-21	Branch Instructions.....	3-27
3-22	Condition Register Logical Instructions.....	3-27
3-23	Trap Instructions.....	3-27
3-24	Move To/From Condition Register Instructions.....	3-28
3-25	Memory Synchronization Instructions—UISA.....	3-29
3-26	Move From Time Base Instruction.....	3-30
3-27	Memory Synchronization Instructions—VEA.....	3-31
3-28	User-Level Cache Instructions	3-31
3-29	External Control Instructions	3-32
3-30	System Linkage Instructions	3-33
3-31	Move To/From Machine State Register Instructions.....	3-33
3-32	Move To/From Special-Purpose Register Instructions.....	3-34
3-33	Implementation-Specific SPR Encodings (mf spr)	3-34
3-34	Segment Register Manipulation Instructions	3-36
3-35	Translation Lookaside Buffer Management Instructions	3-37
4-1	Combinations of W, I, and M Bits.....	4-13
4-2	MEI State Definitions.....	4-16
4-3	core_cse[0:1] Signal Encoding.....	4-18
4-4	Memory Coherency Actions on Load Operations.....	4-19
4-5	Memory Coherency Actions on Store Operations.....	4-20
4-6	Response to Bus Transactions	4-20
4-7	Bus Operations Caused by Cache Control Instructions (WIM = 001).....	4-27
4-8	MEI State Transitions	4-29
4-9	Cache Organization	4-32
4-10	HID0 Bits Used to Perform Cache Locking.....	4-33
4-11	HID2 Bits Used to Perform Cache Way-Locking	4-33
4-12	MSR Bits Used to Perform Cache Locking	4-33
4-13	Example BAT Settings for Cache Locking	4-35
4-14	MSR Bits for Disabling Exceptions	4-35
4-15	G2 Core DWLCK[0–2] Encodings	4-38
4-16	Example BAT Settings for Cache Locking	4-39
4-17	MSR Bits for Disabling Exceptions	4-40

Tables

Table Number	Title	Page Number
4-18	G2 Core IWLCK[0–2] Encodings	4-42
5-1	Exception Classifications	5-3
5-2	Exceptions and Conditions	5-4
5-3	Exception Priorities	5-6
5-4	SRR1 Bit Settings for Machine Check Exceptions	5-10
5-5	SRR1 Bit Settings for Software Table Search Operations	5-10
5-6	Conventional Uses of SPRG4–SPRG7	5-12
5-7	MSR Bit Settings	5-12
5-8	IEEE Floating-Point Exception Mode Bits	5-14
5-9	MSR Setting Due to Exception	5-18
5-10	Hard Reset MSR Value and Exception Vector	5-20
5-11	Settings Caused by Hard Reset	5-20
5-12	Soft Reset Exception—Register Settings	5-21
5-13	Machine Check Exception—Register Settings	5-24
5-14	DSI Exception—Register Settings	5-25
5-15	External Interrupt—Register Settings	5-28
5-16	Alignment Interrupt—Register Settings	5-29
5-17	Access Types	5-30
5-18	Critical Interrupt—Register Settings	5-34
5-19	Trace Exception—Register Settings	5-35
5-20	Instruction and Data TLB Miss Exceptions—Register Settings	5-37
5-21	Instruction Address Breakpoint Exception—Register Settings	5-38
5-22	Breakpoint Action for Multiple Modes Enabled for the Same Address	5-39
5-23	System Management Interrupt—Register Settings	5-40
6-1	MMU Features Summary	6-2
6-2	Access Protection Options for Pages	6-10
6-3	Translation Exception Conditions	6-15
6-4	Other MMU Exception Conditions	6-16
6-5	Instruction Summary—MMU Control	6-17
6-6	MMU Registers	6-18
6-7	Table Search Operations to Update History Bits—TLB Hit Case	6-22
6-8	Model for Guaranteed R and C Bit Settings	6-24
6-9	Implementation-Specific Resources for Table Search Operations	6-32
6-10	Implementation-Specific SRR1 Bits	6-33
6-11	DCMP and ICMP Bit Settings	6-35
6-12	HASH1 and HASH2 Bit Settings	6-35
6-13	RPA Bit Settings	6-36
7-1	Branch Instructions	7-26
7-2	System Register Instructions	7-26
7-3	Condition Register Logical Instructions	7-27
7-4	Integer Instructions	7-27

Tables

Table Number	Title	Page Number
7-5	Floating-Point Instructions	7-29
7-6	Load and Store Instructions	7-30
8-1	Input/Output Enable and High-Impedance Signal Mappings.....	8-4
8-2	Conditions for Unidirectional/Bidirectional Signals	8-5
8-3	Truth Table for Bidirectional Signals	8-6
8-4	G2 Core Signal Cross Reference	8-6
8-5	G2 Core Snoop Hit Response	8-20
8-6	Transfer Type Encoding for the G2 Core as a Bus Master.....	8-21
8-7	Implementation-Specific Transfer Type Encoding.....	8-22
8-8	Data Transfer Size	8-23
8-9	Encodings for core_tc[0:1] Signals	8-24
8-10	Data Bus Lane Assignments.....	8-32
8-11	Data Bus Parity Signal Assignments	8-34
8-12	core_clk_out Signal Configuration.....	8-54
8-13	Core PLL Configuration.....	8-55
9-1	Timing Diagram Legend.....	9-8
9-2	Transfer Size Signal Encodings.....	9-13
9-3	Burst Ordering—64-Bit Bus.....	9-14
9-4	Burst Ordering—32-Bit Bus.....	9-14
9-5	Aligned Data Transfers (64-Bit Bus).....	9-15
9-6	Misaligned Data Transfers (4-Byte Examples)	9-16
9-7	Aligned Data Transfers (32-Bit Bus Mode)	9-17
9-8	Misaligned 32-Bit Data Bus Transfer (4-Byte Examples)	9-18
9-9	Transfer Code Encoding.....	9-19
9-10	core_cse[0:1] Signals.....	9-30
9-11	IEEE Interface Pin Descriptions.....	9-43
10-1	G2 core Programmable Power Modes.....	10-3
11-1	Other Debug and Support Register Bits	11-3
11-2	Related Debug Exceptions and Conditions	11-6
11-3	Single Address Matching (G2 Core Emulation)	11-7
11-4	Two Addresses OR Matching.....	11-7
11-5	Address Matching for Inside Address Range.....	11-7
11-6	Address Matching for Outside Address Range	11-8
A-1	Complete Instruction List Sorted by Mnemonic.....	A-1
A-2	Complete Instruction List Sorted by Opcode.....	A-8
A-3	Integer Arithmetic Instructions	A-15
A-4	Integer Compare Instructions.....	A-16
A-5	Integer Logical Instructions	A-16
A-6	Integer Rotate Instructions	A-17
A-7	Integer Shift Instructions.....	A-17
A-8	Floating-Point Arithmetic Instructions	A-18

Tables

Table Number	Title	Page Number
A-9	Floating-Point Multiply-Add Instructions	A-18
A-10	Floating-Point Rounding and Conversion Instructions.....	A-18
A-11	Floating-Point Compare Instructions	A-19
A-12	Floating-Point Status and Control Register Instructions.....	A-19
A-13	Integer Load Instructions	A-19
A-14	Integer Store Instructions	A-20
A-15	Integer Load and Store with Byte-Reverse Instructions	A-20
A-16	Integer Load and Store Multiple Instructions	A-20
A-17	Integer Load and Store String Instructions	A-21
A-18	Memory Synchronization Instructions.....	A-21
A-19	Floating-Point Load Instructions	A-21
A-20	Floating-Point Store Instructions	A-21
A-21	Floating-Point Move Instructions	A-22
A-22	Branch Instructions	A-22
A-23	Condition Register Logical Instructions	A-22
A-24	System Linkage Instructions	A-23
A-25	Trap Instructions	A-23
A-26	Processor Control Instructions	A-23
A-27	Cache Management Instructions.....	A-23
A-28	Segment Register Manipulation Instructions.....	A-24
A-29	Lookaside Buffer Management Instructions.....	A-24
A-30	External Control Instructions	A-24
A-31	I-Form	A-25
A-32	B-Form.....	A-25
A-33	SC-Form.....	A-25
A-34	D-Form.....	A-25
A-35	DS-Form.....	A-27
A-36	X-Form.....	A-27
A-37	XL-Form	A-31
A-38	XFX-Form.....	A-32
A-39	XFL-Form	A-32
A-40	XS-Form.....	A-32
A-41	XO-Form.....	A-33
A-42	A-Form.....	A-33
A-43	M-Form	A-34
A-44	MD-Form	A-35
A-45	MDS-Form	A-35
A-46	PowerPC Instruction Set Legend.....	A-36



Tables

**Table
Number**

Title

**Page
Number**

About This Book

The primary objective of this reference manual is to define the functionality of the G2 core, a derivative of the original MPC603e PowerPC™ microprocessor design. The G2 core is an implementation of the PowerPC microprocessor family. This reference manual also describes the G2_LE core, which is a derivative of the G2 core. It is written from the perspective of the G2 core and unless otherwise noted, the information applies to both the G2 and G2_LE core. The G2_LE core has the similar functionality to the G2 core and any differences regarding registers, signals, exception model, and debug features are summarized in Section 1.4, “Differences Between the MPC603e and the G2 and G2_LE Cores.” This book is intended as a companion to the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture* (referred to as the *Programming Environment Manual*) which provides a general description of the features that are common to processors and cores that implement the PowerPC architecture and indicates those features that are optional or that may be implemented differently in the design of each processor and core.

NOTE

About the Companion *Programming Environments Manual*

The PowerPC architecture definition is flexible to support to a broad range of the processors as well as cores. Note that the *Programming Environments Manual* describes only PowerPC architecture features for 32-bit implementations.

Contact your sales representative for a copy of the *Programming Environments Manual*.

This reference manual and the *Programming Environments Manual* distinguish between the three levels, or programming environments, of the PowerPC architecture, which are as follows:

- User instruction set architecture (UISA)—The UISA defines the architecture level to which user-level software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, memory conventions, and the memory and programming models seen by application programmers.
- Virtual environment architecture (VEA)—The VEA, which is the smallest component of the PowerPC architecture, defines additional user-level functionality

that falls outside typical user-level software requirements. The VEA describes the memory model for an environment in which multiple processors or other devices can access external memory, and defines aspects of the cache model and cache control instructions from a user-level perspective. The resources defined by the VEA are particularly useful for optimizing memory accesses and for managing resources in an environment in which other processors and devices can access external memory.

- Operating environment architecture (OEA)—The OEA defines supervisor-level resources typically required by an operating system. The OEA defines the memory management model, supervisor-level registers, and exception model.

Implementations that conform to the OEA also conform to the UISA and VEA.

Note that some resources are defined more generally at one level in the architecture and more specifically at another. For example, conditions that cause a floating-point exception are defined by the UISA, while the exception mechanism itself is defined by the OEA.

For ease in reference, topics in this book are presented in the same order as the *Programming Environments Manual*. Topics build on one another, beginning with a description and complete summary of the G2 core register model (including the G2_LE core specifics) and followed by the instruction set model and progressing to more specific, architecture-based topics regarding the cache, exception, and memory management models. As such, chapters may include information from multiple levels of the architecture. (For example, the discussion of the cache model uses information from both the VEA and the OEA.)

The PowerPC Architecture: A Specification for a New Family of RISC Processors defines the architecture from the perspective of the three programming environments and remains the defining document for the PowerPC architecture. For information about ordering Freescale documentation, see “Suggested Reading” on page xxxiv.

The information in this book is subject to change without notice, as described in the disclaimers on the title page of this book. As with any technical documentation, it is the readers’ responsibility to be sure they are using the most recent version of the documentation. For more information, contact your sales representative.

To locate any published errata or updates for this reference manual, refer to the world-wide web at <http://www.freescale.com>.

A list of major differences between the MPC603e microprocessor, the G2 core, and the G2_LE core are provided in Table 1-6. Note that the G2 core has similar functionality as the MPC603e. However, the minor differences between them are documented by footnotes.

Audience

This manual is intended to be used as a reference for many semiconductor products targeting a range of markets including automotive, communication, consumer, networking, and computer peripherals. It is intended for system software and hardware developers and applications programmers who want to develop products using the cores. It is assumed that the reader understands operating systems, core system design, and details of the PowerPC architecture.

Organization

Following is a summary and a brief description of the major sections of this manual:

- Chapter 1, “Overview,” is useful for readers who want a general understanding of the features and functions of the PowerPC architecture and the differences between the G2 and G2_LE cores. This chapter describes the flexible nature of the PowerPC architecture definition, and provides an overview of how the PowerPC architecture defines the register set, instruction set and addressing modes, cache model (including instruction and data cache way-locking for the G2 core), exception model, memory management model, instruction timing, system support interface, and debug features for the G2 and G2_LE cores.
- Chapter 2, “Register Model,” provides a brief synopsis of the registers implemented in the G2 core and some registers implemented only in the G2_LE core.
- Chapter 3, “Instruction Set Model,” provides a brief description of the operand conventions, an overview of the PowerPC addressing modes, and a list of the instructions implemented by the G2 core. Note that instructions are organized by functions.
- Chapter 4, “Instruction and Data Cache Operation,” provides a discussion of the cache and memory model as implemented on the G2 core.
- Chapter 5, “Exceptions,” describes the exception model defined in the PowerPC OEA, and the specific exception model implemented on the G2 and G2_LE cores.
- Chapter 6, “Memory Management,” describes the G2 core’s implementation of the memory management unit specifications provided by the OEA.
- Chapter 7, “Instruction Timing,” provides information about latencies, interlocks, special situations, and various conditions to help make programming more efficient. This chapter is of special interest to software engineers and system designers.
- Chapter 8, “Signal Descriptions,” provides descriptions of individual signals of the G2 core that are candidates for being driven as external device signals. This chapter also describes signals which are only defined in the G2_LE core.

- Chapter 9, “Core Interface Operation,” describes signal timings for various operations. It also provides a detailed description of the 60x bus interface, the multiple bus master capability, and the memory coherency features of the G2 core.
- Chapter 10, “Power Management,” provides information about the power saving modes for the G2 core.
- Chapter 11, “Debug Features,” provides information about the debug features of the G2_LE core. This chapter also describes trace facility debug features for both the G2 and G2_LE cores.
- Appendix A, “PowerPC Instruction Set Listings,” lists all the PowerPC instructions while indicating those instructions that are not implemented by the G2 and G2_LE cores; it also includes the instructions that are specific to the G2 and G2_LE cores. Instructions are grouped according to mnemonic, opcode, function, and form. Also included is a quick reference table that contains general information, such as the architecture level, privilege level, and form, and indicates if the instruction is 64-bit and optional.
- Appendix B, “Revision History,” lists the major differences between Revision 1 and Revision 2 of the *G2 Core Reference Manual*.
- This reference manual also includes a glossary and an index.

Suggested Reading

This section lists additional reading that provides background for the information in this reference manual, as well as general information about the PowerPC architecture.

General Information

The following documentation, available through Morgan-Kaufmann Publishers, 340 Pine Street, Sixth Floor, San Francisco, CA, provides useful information about the PowerPC architecture and computer architecture in general:

- *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, by International Business Machines, Inc.
Updates to the architecture specification are accessible via the world-wide web at <http://www.austin.ibm.com/tech/ppc-chg.html>.
- *PowerPC Microprocessor Common Hardware Reference Platform: A System Architecture*, by Apple Computer, Inc., International Business Machines, Inc., and Freescale Semiconductor, Inc.
- *Computer Architecture: A Quantitative Approach*, Second Edition, John L. Hennessy and David A. Patterson.
- *Computer Organization and Design: The Hardware/Software Interface*, Second Edition, David A. Patterson and John L. Hennessy.

- *Inside Macintosh: PowerPC System Software*, Addison-Wesley Publishing Company, One Jacob Way, Reading, MA, 01867; Tel. (800) 282-2732 (U.S.A.), (800) 637-0029 (Canada), (716) 871-6555 (International).

Related Documentation

Freescale documentation is available from the sources listed on the back cover of this manual; the document order numbers are included in parentheses for ease in ordering:

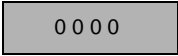
- *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture* (MPCFPE32B/AD)—Describes resources defined by the PowerPC architecture.
- User's and reference manuals—These books provide details about individual implementations and are intended for use with the *Programming Environments Manual*.
- Addenda/errata to user's or reference manuals—Because some processors have follow-on devices, an addendum is provided that describes the additional features and functionality changes. These addenda are intended for use with the corresponding book.
- *Implementation Variances Relative to Rev. 1 of The Programming Environments Manual* is available at <http://www.freescale.com>.
- Technical summaries—Each device has a technical summary that provides an overview of its features. This document is roughly the equivalent to the overview (Chapter 1) of an implementation's user's or reference manual.
- Application notes—These short documents contain useful information about specific design issues useful to programmers and engineers working with Freescale processors.
- Documentation for support chips—These include the following:
 - *MPC106 PCI Bridge/Memory Controller User's Manual* (MPC106UM/AD)
 - *MPC107 PCI Bridge/Memory Controller Technical Summary* (MPC107TS/D)
 - *MPC107 PCI Bridge/Memory Controller User's Manual* (MPC107UM/AD)

Additional literature is published as new processors become available. For a current list of documentation, refer to: <http://www.mot.com/semiconductors>.

Conventions

This document uses the following notational conventions:

mnemonics	Instruction mnemonics are shown in lowercase bold
<i>italics</i>	Italics indicate variable command parameters, for example, bcctrx Book titles in text are set in italics

0x0	Prefix to denote hexadecimal number
0b0	Prefix to denote binary number
rA, rB	Instruction syntax used to identify a source GPR
rA 0	Contents of a specified GPR or the value 0
rD	Instruction syntax used to identify a destination GPR
frA, frB, frC	Instruction syntax used to identify a source FPR
frD	Instruction syntax used to identify a destination FPR
REG[FIELD]	Abbreviations or acronyms for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, MSR[LE] refers to the little-endian mode enable bit in the machine state register.
x	In certain contexts, such as a signal encoding, this indicates a don't care
n	Used to express an undefined numerical value
¬	NOT logical operator
&	AND logical operator
	OR logical operator
	Indicates reserved bits or bit fields in a register. Although these bits may be written to as either ones or zeros, they are always read as zeros.

Acronyms and Abbreviations

Table i contains acronyms and abbreviations that are used in this reference manual.

Table i. Acronyms and Abbreviated Terms

Term	Meaning
ABE	Address bus enable
ALU	Arithmetic logic unit
BAT	Block address translation
BATL	Block address translation lower
BATU	Block address translation upper
BE	Branch trace enable
BIST	Built-in self test
BIU	Bus interface unit
BL	Block size mask

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
BPU	Branch processing unit
BUID	Bus unit ID
CE	Critical interrupt exception enable
CIA	Current instruction address
CMOS	Complementary metal-oxide semiconductor
CMP	IABR compare type
CMP2	IABR2 compare type
COP	Common on-chip processor
CR	Condition register
CSRR0	Critical interrupt save/restore register 0
CSRR1	Critical interrupt save/restore register 1
CTR	Count register
CQ	Completion queue
DAR	Data address register
DABR	Data address breakpoint register
DABR2	Data address breakpoint register 2
DBAT	Data BAT
DBCR	Data address control register
DCE	Data cache enable
DCF1	Data cache flash invalidate
DCMP	Data TLB compare
DEC	Decrementer register
DLOCK	Data cache lock
DMISS	Data TLB miss address
DMMU	Data memory management unit
DPM	Dynamic power management enable
DR	Data address translation enable
DSISR	Register used for determining the source of a DSI exception
DTLB	Data translation lookaside buffer
DWLCK	Data cache way-lock
EA	Effective address
EAR	External access register
ECC	Error checking and correction

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
EE	External interrupt enable
FE0	Floating-point exception model 0
FE1	Floating-point exception model 1
FIFO	First-in-first-out
FP	Floating-point available
FPR	Floating-point register
FPSCR	Floating-point status and control register
FPU	Floating-point unit
GPR	General-purpose register
HBE	High BAT enable
HID0	Hardware implementation register 0
HID1	Hardware implementation register 1
HID2	Hardware implementation register 2
I	Cache-inhibited
IABR	Instruction address breakpoint register 1
IABR2	Instruction address breakpoint register 2
IBAT	Instruction BAT
IBCR	Instruction breakpoint control register
ICE	Instruction cache enable
ICFI	Instruction cache flash invalidate
ICMP	Instruction TLB compare
IEEE	Institute for Electrical and Electronics Engineers
IEE	External interrupt enable
IFEM	Instruction fetch enable M (bit)
ILE	Exception little-endian mode
ILOCK	Instruction cache lock
IMISS	Instruction TLB miss address
IMMU	Instruction memory management unit
IP	Exception prefix
IQ	Instruction queue
IR	Instruction address translation enable
ITLB	Instruction translation lookaside buffer
IU	Integer unit

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
IWLCK	Instruction cache way-lock
L2	Secondary cache
LE	Little-endian mode enable
LET	True little-endian mode bit
LIFO	Last-in-first-out
LR	Link register
LRU	Least recently used
LSB	Least-significant byte
lsb	Least-significant bit
LSU	Load/store unit
M	Memory-coherent
MBAR	System memory base address
ME	Machine check enable
MEI	Modified/exclusive/invalid
MESI	Modified/exclusive/shared/invalid—cache coherency protocol
MFG	Manufacturing revision tag
MJREV	Major processor design revision indicator
MNREV	Minor processor design revision indicator
MMU	Memory management unit
MQ	MQ register
MSB	Most-significant byte
msb	Most-significant bit
MSR	Machine state register
NaN	Not a number
NI	non-IEEE mode bit
No-op	No operation
NOOPTI	No-op the data cache touch instructions
OEA	Operating environment architecture
PID	Processor identification tag
PIR	Processor identification register
PLL	Phase-locked loop
POWER	Performance Optimized with Enhanced RISC architecture
POW	Power management enable

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
POR	Power-on reset
PROC	Processor revision tag
PR	Privilege level
PT	Processor ID type tag
PTE	Page table entry
PTEG	Page table entry group
PVR	Processor version register
RAW	Read-after-write
RI	Recoverable exception
RID	Resource ID
RISC	Reduced instruction set computing
RTL	Register transfer language
RWITM	Read with intent to modify
SDR1	Register that specifies the page table base address for virtual-to-physical address translation
SE	Single-step trace enable
SOC	System-on-a-Chip
SPR	Special-purpose register
SR	Segment register
SRR0	Machine status save/restore register 0
SRR1	Machine status save/restore register 1
SRU	System register unit
SMI	System management interrupt
SVR	System version register
SIG_TYPE	Combinational signal type
T	Translation control bit
TAP	Test access port
TB	Time base facility
TBL	Time base lower register
TBU	Time base upper register
TGPR	Temporary GPR remapping
TLB	Translation lookaside buffer
TTL	Transistor-to-transistor logic
UIMM	Unsigned immediate value

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
UIISA	User instruction set architecture
UTLB	Unified translation lookaside buffer
UUT	Unit under test
VEA	Virtual environment architecture
VPN	Virtual page number
W	Write-through
WAR	Write-after-read
WAW	Write-after-write
WIMG	Write-through/caching-inhibited/memory-coherency enforced/guarded bits
XATC	Extended address transfer code
XER	Register used for indicating conditions such as carries and overflows for integer operations

Terminology Conventions

Table ii describes terminology conventions used in this manual.

Table ii. Terminology Conventions

The Architecture Specification	This Manual
Data storage interrupt (DSI)	DSI exception
Extended mnemonics	Simplified mnemonics
Fixed-point unit (FXU)	Integer unit (IU)
Instruction storage interrupt (ISI)	ISI exception
Interrupt	Exception
Privileged mode (or privileged state)	Supervisor-level privilege
Problem mode (or problem state)	User-level privilege
Real address	Physical address
Relocation	Translation
Storage (locations)	Memory
Storage (the act of)	Access
Store in	Write back
Store through	Write through

Table iii describes instruction field notation used in this manual.

Table iii. Instruction Field Conventions

The Architecture Specification	Equivalent to:
BA, BB, BT	crbA, crbB, crbD (respectively)
BF, BFA	crfD, crfS (respectively)
D	d
DS	ds
FLM	FM
FRA, FRB, FRC, FRT, FRS	frA, frB, frC, frD, frS (respectively)
FXM	CRM
RA, RB, RT, RS	rA, rB, rD, rS (respectively)
SI	SIMM
U	IMM
UI	UIMM
/, //, ///	0...0 (shaded)

Chapter 1

Overview

This chapter provides an overview of features for the embedded G2 processor core, a derivative of the original MPC603e PowerPC™ microprocessor design. The G2 core is an implementation of the PowerPC microprocessor family. This reference manual also describes the G2_LE core, which is a derivative of the G2 core. The G2_LE core implements some enhanced features with a true little-endian mode, an additional critical interrupt signal, and four additional instruction BAT and four additional data BAT registers. This document is written from the perspective of the G2 core and all of the descriptions apply to both the G2 and G2_LE cores, except where explicitly noted. Note that throughout this document, the terms G2 core, core, and processor are used interchangeably.

1.1 Overview

This section describes the details of the G2 core, provides a block diagram showing the major functional units (see Figure 1-1), and briefly describes how these units interact. All differences between the G2 and G2_LE implementations are noted.

The G2 core is a low-power implementation of this microprocessor family of reduced instruction set computing (RISC) microprocessors. The core implements the 32-bit portion of the PowerPC architecture, which defines 32-bit effective addresses, integer data types of 8, 16, and 32 bits, and floating-point data types of 32 and 64 bits.

The G2 core is a superscalar processor that can issue and retire as many as three instructions per clock cycle. Instructions can execute out of program order for increased performance; however, the core makes completion appear sequential.

The G2 core integrates five execution units—an integer unit (IU), a floating-point unit (FPU), a branch processing unit (BPU), a load/store unit (LSU), and a system register unit (SRU). The ability to execute five instructions in parallel and the use of simple instructions with rapid execution times yield high efficiency and throughput for G2-core based systems. Most integer instructions execute in one clock cycle. On the G2 core, the FPU is pipelined so a single-precision multiply-add instruction can be issued and completed every clock cycle. The G2 core provides hardware support for all single- and double-precision floating-point operations for most value representations and all rounding modes.

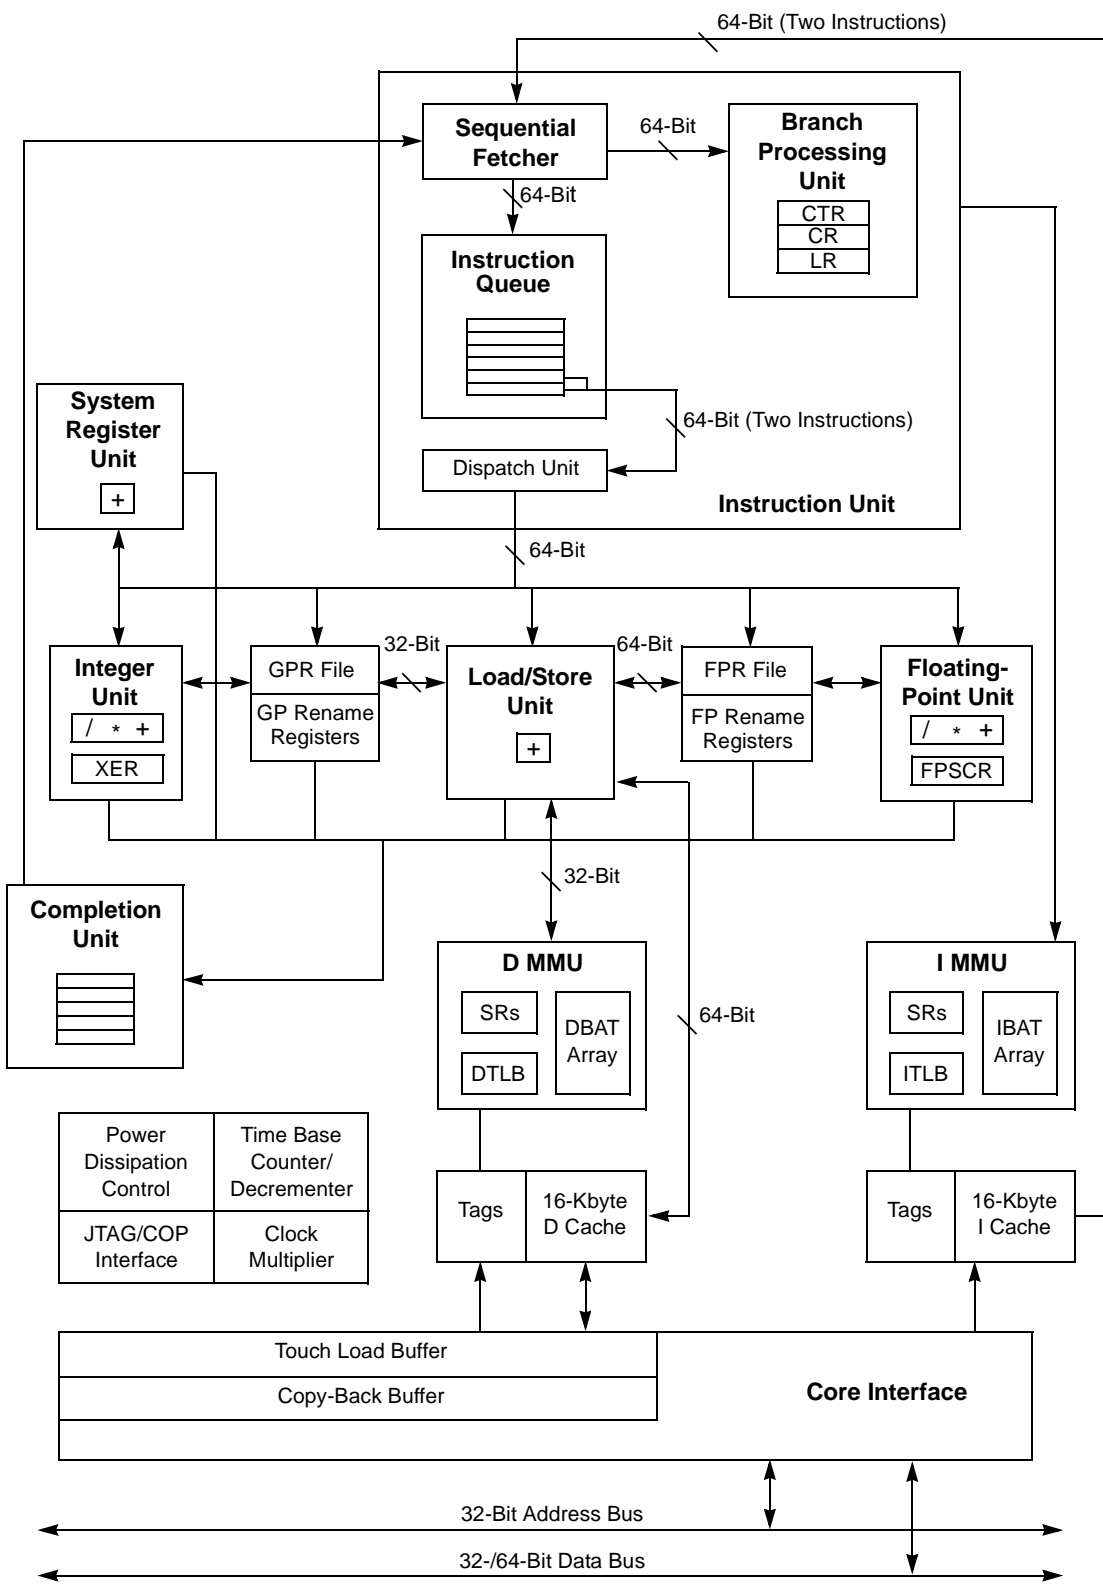


Figure 1-1. G2 Core Block Diagram

The G2 core provides independent on-chip, 16-Kbyte, four-way set-associative, physically-addressed caches for instructions and data, and on-chip instruction and data memory management units (MMUs). The MMUs contain 64-entry, two-way set-associative, data and instruction translation lookaside buffers (DTLB and ITLB) that provide support for demand-paged virtual-memory address translation and variable-sized block translation. The TLBs and caches use a least recently used (LRU) replacement algorithm.

The G2 core also supports block address translation through the use of two independent instruction and data block address translation (IBAT and DBAT) arrays, each containing four pairs of BATs; however the G2_LE core supports block address translation arrays of eight pairs of data BATs and eight pairs of instruction BATs. Effective addresses are compared simultaneously with all four (or eight, for G2_LE) entries in the BAT array during block translation. In accordance with the PowerPC architecture, if an effective address hits in both the TLB and BAT array, the BAT translation takes priority.

The G2 core has a selectable 32- or 64-bit 60x data bus and a 32-bit 60x address bus. The core interface protocol allows multiple masters to compete for system resources through a central external arbiter. The G2 core provides a three-state (exclusive, modified, and invalid) coherency protocol which is a compatible subset of a four-state (modified/exclusive/shared/invalid) MESI protocol. This protocol operates coherently in systems that contain four-state caches. The G2 core supports single-beat and burst data transfers for memory accesses and supports memory-mapped I/O operations.

The G2_LE core has a new MMU with eight additional BATs which provides better performance in protecting accesses on a segment, block, or page basis along with memory accesses and I/O accesses. The true little-endian mode is another enhanced capability of the G2_LE core which is not managed on a page basis through the MMU. Unlike the PowerPC little-endian mode (manipulates the only address bits), the true little-endian mode actually operates on true little-endian instructions and data from memory.

The critical interrupt is an additional exception in the G2_LE core which has higher priority order than the system management interrupt. Also debug feature is improved in the G2_LE. See Section 1.3.8, “Debug Features (G2_LE Only),” for more detail. Additional SPRG exception handling registers are provided for enhancing the use of the operating system.

1.1.1 Features

This section describes the major features of the G2 core noting where the G2 and G2_LE implementations differ:

- High-performance, superscalar microprocessor core
 - As many as three instructions issued and retired per clock
 - As many as five instructions in execution per clock

- Single-cycle execution for most instructions
- Pipelined FPU for all single-precision and most double-precision operations
- Five independent execution units and two register files
 - BPU featuring static branch prediction
 - A 32-bit IU
 - Fully IEEE 754-compliant FPU for both single- and double-precision operations
 - LSU for data transfer between data cache and general-purpose registers (GPRs) and floating-point registers (FPRs)
 - SRU that executes condition register (CR), special-purpose register (SPR), and integer add/compare instructions
 - Thirty-two 32-bit GPRs for integer operands
 - Thirty-two 64-bit FPRs for single- or double-precision operands
- High instruction and data throughput
 - Zero-cycle branch capability (branch folding)
 - Programmable static branch prediction on unresolved conditional branches
 - Instruction fetch unit capable of fetching two instructions per clock from the instruction cache
 - A six-entry instruction queue (IQ) that provides lookahead capability
 - Independent pipelines with feed-forwarding that reduces data dependencies in hardware
 - 16-Kbyte data cache and 16-Kbyte instruction cache—four-way set-associative, physically addressed, LRU replacement algorithm
 - Cache write-back or write-through operation programmable on a per page or per block basis
 - BPU that performs CR lookahead operations
 - Address translation facilities for 4-Kbyte page size, variable block size, and 256-Mbyte segment size
 - A 64-entry, two-way set-associative ITLB and DTLB
 - Four-entry data and instruction BAT arrays (for G2 core), and eight-entry data and instruction BAT arrays (for G2_LE core) providing 128-Kbyte to 256-Mbyte blocks
 - Software table search operations and updates supported through fast trap mechanism
 - 52-bit virtual address; 32-bit physical address
- Facilities for enhanced system performance
 - A 32- or 64-bit split-transaction data bus interface (60x bus) with burst transfers

- Support for one-level address pipelining and out-of-order bus transactions on the 60x interface
- True little-endian mode (for G2_LE only) for compatibility with other true little-endian devices.
- Critical interrupt exception (for G2_LE only) is added
- Hardware support for misaligned little-endian accesses
- Integrated power management
 - Internal processor/bus clock multiplier ratios
 - Three power-saving modes: doze, nap, and sleep
 - Automatic dynamic power reduction when internal functional units are idle
- In-system testability and debugging features through JTAG boundary-scan capability

Features specific to the G2 core not present on the original MPC603e (PID6-603e) processors follow:

- Enhancements to the register set
 - The G2 core has two more additional HID0 bits than the original MPC603e:
 - The address bus enable (ABE) bit, HID0[28], gives the G2 core the ability to broadcast **dcbf**, **dcbi**, and **dcbst** onto the 60x bus.
 - The instruction fetch enable M (IFEM) bit, HID0[24], allows the G2 core to reflect the value of the M bit during instruction translation onto the 60x bus.
 - The G2 core has one more additional HID2 register than the original MPC603e that enables the true little-endian mode, the new additional BAT registers, and cache way-locking for the G2 core.
- Enhancements to cache implementation
 - The instruction cache is blocked only until the critical load completes (hit under reloads allowed)
 - The critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.
 - The HID2 register enables instruction and data cache way-locking.
 - Provides for an optional data cache operation broadcast feature (enabled by HID0[ABE]) that allows for correct system management using an external copy-back L2 cache.
 - All of the cache control instructions (**icbi**, **dcbi**, **dcbf**, and **dcbst**, excluding **dcbz**) require that HID0[ABE] be enabled in order to execute.
- Exceptions
 - The G2 core offers hardware support for misaligned little-endian accesses. Little-endian load/store accesses that are not on a word boundary, with the

- exception of strings and multiples, generate exceptions under the same circumstances as big-endian accesses.
- The G2_LE core supports true little-endian mode to minimize the impact on software porting from true little-endian systems.
- A new input interrupt signal, `core_cint`, is provided to trigger the critical interrupt exception on the G2_LE core.
- The G2 core does not have misalignment support for `eciwx` and `ecowx` graphics instructions. These instructions cause an alignment exception if the access is not on a word boundary.
- Bus clock—New bus multipliers are selected by the encodings of `core_pll_cfg[0:4]`.
- Instruction timing
 - The integer divide instructions, `divwu[o][.]` and `divw[o][.]`, execute in 20 clock cycles; execution of these instructions in the original PID6 MPC603e device takes 37 clock cycles.
 - Support for single-cycle store
 - An adder/comparator added to system register unit that allows dispatch and execution of multiple integer add and compare instructions on each cycle.
- Enhanced debug features
 - Addition of three breakpoint registers—IABR2, DABR, and DABR2
 - Two new breakpoint control registers—DBCR and IBCR
 - Inclusion of four breakpoint signals—`core_iabr`, `core_iabr2`, `core_dabr`, and `core_dabr2`

Figure 1-1 provides a block diagram of the G2 core that shows how the execution units—IU, FPU, BPU, LSU, and SRU—operate independently and in parallel. Note that this is a conceptual diagram and does not attempt to show how these features are physically implemented on the chip.

The G2 core provides address translation and protection facilities, including an ITLB, DTLB, and instruction and data BAT arrays. Instruction fetching and issuing is handled in the instruction unit. Translation of addresses for cache or external memory accesses are handled by the MMUs. Both units are discussed in more detail in Section 1.1.3, “Instruction Unit,” and Section 1.1.6.1, “Memory Management Units (MMUs).”

1.1.2 G2_LE-Specific Features

The following sections describe some of the additional features of the G2_LE core. For a table summary of the differences between the G2 core and the G2_LE cores, see Section 1.4, “Differences Between the MPC603e and the G2 and G2_LE Cores.”

1.1.2.1 True Little-Endian Mode

True little-endian mode is supported in the G2_LE core to minimize the impact on software porting from true little-endian systems. The true little-endian mode applies for all instruction fetches and data load and store operations to and from memory. The G2_LE powers up in one of two endian modes, big-endian mode or true little-endian mode, selected by the `core_tle` signal at the negation of `core_hreset`. Like all the mode control signals, the state of `core_tle` is captured at the negation of `core_hreset`. The state of `MSR[ILE]`, `MSR[LE]`, and `HID2[LET]` are set to the value that is dictated by `core_tle`. The endian mode should be set at the negation of `core_hreset` and should remain unchanged by software for the duration of the system operation.

Bit 4 of `HID2`, (`HID2[LET]`) is used in conjunction with `MSR[LE]` to indicate the endian mode of operation of the G2_LE core as shown in Table 1-1.

Table 1-1. Endian Mode Indication

MSR[LE]	HID2[LET]	Endian Mode
0	x	Big-endian
1	0	Modified (PowerPC) little-endian
1	1	True little-endian

1.1.2.2 Critical Interrupt

A new input interrupt signal, `core_cint`, is provided to trigger the critical interrupt exception on the G2_LE core. This asynchronous exception uses vector offset 0x00A00. `MSR[CE]` is allocated for enabling the critical interrupt, and a new instruction, Return from Critical Interrupt (`rfci`), is implemented to return from these exception handlers. Also, two new registers, `CSRR0` and `CSRR1`, are used to save and restore the processor state for critical interrupts.

Table 1-2 shows the bit allocation of `MSR[CE]`, which enables and disables the critical interrupt.

Table 1-2. Critical Interrupt Enabling Bit

MSR[CE]	Mode
24	0 Disables critical interrupt 1 Enables critical interrupt

1.1.2.3 Other New Signals

There are four additional signals that support the breakpoint state outputs (`core_iabr`, `core_iabr2`, `core_dabr`, and `core_dabr2`) and one additional watchpoint debug feature (`core_tdo_oe`).

1.1.2.4 Additional Supervisor-Level SPRs

The G2_LE core has 29 new/additional supervisor-level SPRs. See Section 1.3.1.7, “Special-Purpose Registers (SPRs),” for more information.

1.1.3 Instruction Unit

As shown in Figure 1-1, the G2 core instruction unit, containing a fetch unit, instruction queue, dispatch unit, and BPU, provides centralized control of instruction flow to the execution units. The instruction unit determines the address of the next instruction to be fetched based on information from the sequential fetcher and from the BPU.

The instruction unit fetches the instructions from the instruction cache into the instruction queue. The BPU receives branch instructions from the fetcher and uses static branch prediction to allow fetching from a predicted instruction stream while a conditional branch is evaluated. The BPU folds out for unconditional branch instructions and conditional branch instructions unaffected by instructions in the execution pipeline.

Instructions issued beyond a predicted branch cannot complete execution until the branch is resolved, preserving the programming model of sequential execution. If any of these are branch instructions, they are decoded but not issued. Instructions to be executed by the FPU, IU, LSU, and SRU are issued and allowed to progress up to the register write-back stage. Write-back is allowed when a correctly predicted branch is resolved, and execution continues along the predicted path.

If branch prediction is incorrect, the instruction unit flushes all predicted path instructions, and instructions are issued from the correct path.

1.1.3.1 Instruction Queue and Dispatch Unit

The instruction queue (IQ), shown in Figure 1-1, holds as many as six instructions and loads up to two instructions from the instruction unit during a single cycle. The instruction fetch unit continuously loads as many instructions as space in the IQ allows. Instructions are dispatched to their respective execution units from the dispatch unit at a maximum rate of two instructions per cycle. Dispatching is facilitated to the IU, FPU, LSU, and SRU by the provision of a reservation station at each unit. The dispatch unit performs source and destination register dependency checking, determines dispatch serializations, and inhibits subsequent instruction dispatching as required.

For a more detailed overview of instruction dispatch, see Section 1.3.6, “Instruction Timing.”

1.1.3.2 Branch Processing Unit (BPU)

The BPU receives branch instructions from the fetch unit and performs CR lookahead operations on conditional branches to resolve them early, achieving the effect of a zero-cycle branch in many cases.

The BPU uses a bit in the instruction encoding to predict the direction of the conditional branch. Therefore, when an unresolved conditional branch instruction is encountered, the core fetches instructions from the predicted target stream until the conditional branch is resolved.

The BPU contains an adder to compute branch target addresses and three user-control registers—the link register (LR), the count register (CTR), and the conditional register (CR). The BPU calculates the return pointer for subroutine calls and saves it into the LR for certain types of branch instructions. The LR also contains the branch target address for the Branch Conditional to Link Register (**bcldr x**) instruction. The CTR contains the branch target address for the Branch Conditional to Count Register (**bcctr x**) instruction. The contents of the LR and CTR can be copied to or from any GPR. Because the BPU uses dedicated registers rather than GPRs or FPRs, execution of branch instructions is largely independent from execution of integer and floating-point instructions.

1.1.4 Independent Execution Units

The PowerPC architecture's support for independent execution units allows implementation of processors with out-of-order instruction execution. For example, because branch instructions do not depend on GPRs or FPRs, branches can often be resolved early, eliminating stalls caused by taken branches.

The four other execution units and the completion unit are described in the following sections.

1.1.4.1 Integer Unit (IU)

The IU executes all integer instructions. The IU executes one integer instruction at a time, performing computations with its arithmetic logic unit (ALU), multiplier, divider, and XER register. Most integer instructions are single-cycle instructions. The 32 GPRs hold integer operands. Stalls due to contention for GPRs are minimized by the automatic allocation of rename registers. The G2 core writes the contents of the rename registers to the appropriate GPR when integer instructions are retired by the completion unit.

1.1.4.2 Floating-Point Unit (FPU)

The FPU contains a single-precision multiply-add array and the floating-point status and control register (FPSCR). The multiply-add array allows the G2 core to efficiently implement multiply and multiply-add operations. The FPU is pipelined so that single- and

double-precision instructions can be issued back-to-back. The 32 FPRs are provided to support floating-point operations. Stalls due to contention for FPRs are minimized by the automatic allocation of rename registers. The G2 core writes the contents of the rename registers to the appropriate FPR when floating-point instructions are retired by the completion unit.

The G2 core supports all IEEE-754 floating-point data types (normalized, denormalized, NaN, zero, and infinity) in hardware, eliminating the latency incurred by software exception routines.

1.1.4.3 Load/Store Unit (LSU)

The LSU executes all load and store instructions and provides the data transfer interface between the GPRs, FPRs, and the cache/memory subsystem. The LSU calculates effective addresses, performs data alignment, and provides sequencing for load/store string and multiple instructions.

Load and store instructions are issued and executed in program order; however, the memory accesses can occur out of order. Synchronizing instructions are provided to enforce strict ordering.

Cacheable loads, when free of data bus dependencies, can execute out of order with a maximum throughput of one per cycle and a two-cycle total latency. Data returned from the cache is held in a rename register until the completion logic commits the value to a GPR or FPR. Stores cannot be executed in a predicted manner and are held in the store queue until the completion logic signals that the store operation is to be completed to memory. The core executes store instructions with a maximum throughput of one per cycle and a three-cycle total latency. The time required to perform the actual load or store depends on whether the operation involves the cache, system memory, or an I/O device.

1.1.4.4 System Register Unit (SRU)

The SRU executes various system-level instructions, including condition register logical operations and move to/from special-purpose register instructions. It also executes integer add/compare instructions. In order to maintain system state, most instructions executed by the SRU are completion-serialized; that is, the instruction is held for execution in the SRU until all prior instructions issued have completed. Results from completion-serialized instructions executed by the SRU are not available or forwarded for subsequent instructions until the instruction completes.

1.1.5 Completion Unit

The completion unit tracks instructions in program order from dispatch through execution and then completes. Completing an instruction commits the core to any architectural

register changes caused by that instruction. In-order completion ensures the correct architectural state when the core must recover from a mispredicted branch or any exception.

Instruction state and other information required for completion is kept in a five-entry FIFO completion queue. A single completion queue entry is allocated for each instruction once it enters the execution unit from the dispatch unit. An available completion queue entry is a required resource for dispatch; if no completion entry is available, dispatch stalls. A maximum of two instructions per cycle are completed in order from the queue.

1.1.6 Memory Subsystem Support

The G2 core provides separate instruction and data caches and MMUs. The core also provides an efficient processor bus interface to facilitate access to main memory and other bus subsystems. The memory subsystem support functions are described in the following sections.

1.1.6.1 Memory Management Units (MMUs)

The G2 core MMUs support up to 4 Petabytes (2^{52}) of virtual memory and 4 Gigabytes (2^{32}) of physical memory (referred to as real memory in the architecture specification) for instruction and data. The MMUs also control access privileges for these spaces on block and page granularities. Referenced and changed status is maintained by the processor for each page to assist implementation of a demand-paged virtual memory system. Note that software assistant is required for the G2 core to maintain reference and changed status. A key bit is implemented to provide information about memory protection violations prior to page table search operations.

The LSU calculates effective addresses (EAs) for data loads and stores, performs data alignment to and from cache memory, and provides the sequencing for load and store string and multiple word instructions. The instruction unit calculates effective addresses for instruction fetching.

After an EA is generated, its higher-order bits are translated by the appropriate MMU into physical address bits. The lower-order EA bits are the same on the physical address which are directed to the on-chip cache and formed the index into a four-way set-associative tag array. After translating the address, the MMU passes the higher-order physical address bits to the cache and the cache lookup completes. For caching-inhibited accesses or accesses that miss in the cache, the untranslated lower-order address bits are concatenated with the translated higher-order address bits; the resulting 32-bit physical address is then used by the memory unit and the system interface to access external memory.

The MMU also directs the address translation and enforces the protection hierarchy programmed by the operating system in relation to the supervisor/user privilege level of the access and in relation to whether the access is a load or store.

For instruction fetches, the IMMU looks for the address in the ITLB and in the IBAT array. If an address hits both, the IBAT array translation is used. Data accesses cause a lookup in the DTLB and DBAT array. In most cases, the translation is in a TLB and the physical address bits are available to the on-chip cache.

The G2_LE core implements four additional IBAT and four additional DBAT entries.

When the EA misses in the TLBs, the core provides hardware assistance for software to perform a search of the translation tables in memory. The hardware assist consists of the following features:

- Automatic storage of the missed effective address in IMISS and DMISS
- Automatic generation of the primary and secondary hashed real address of the page table entry group (PTEG), which are readable from the HASH1 and HASH2 register locations.

The HASH data is generated from the contents of the IMISS or DMISS register. The register that is selected depends on the miss (instruction or data) that was last acknowledged.

- Automatic generation of the first word of the page table entry (PTE) of the tables being searched
- A real page address (RPA) register that matches the format of the lower word of the PTE
- TLB access instructions (**tlbli** and **tlbld**) that are used to load an address translation into the instruction or data TLBs
- Shadow registers for GPR0–GPR3 that allow miss code to execute without corrupting the state of any of the existing GPRs. Shadow registers are used only for servicing a TLB miss.

See Section 1.3.5.2, “Implementation-Specific Memory Management,” for more information about memory management for the core.

1.1.6.2 Cache Units

The G2 core provides independent 16-Kbyte, four-way set-associative instruction and data caches. The cache block is 32 bytes long. The caches adhere to a write-back policy, but the G2 core allows control of cacheability, write policy, and memory coherency at the page and block levels. The caches use an LRU replacement policy.

As shown in Figure 1-1, the caches provide a 64-bit interface to the instruction fetch unit and LSU. The surrounding logic selects, organizes, and forwards the requested information to the requesting unit. Write operations to the cache can be performed on a byte basis, and a complete read-modify-write operation to the cache can occur in each cycle.

The load/store and instruction fetch units provide the caches with the address of the data or instruction to be fetched. In the case of a cache hit, the cache returns two words to the requesting unit.

Because the data cache tags are single-ported, simultaneous load or store and snoop accesses cause resource contention. Snoop accesses have the highest priority and are given first access to the tags, unless the snoop access coincides with a tag write; in this case the snoop is retried and must re-arbitrate for cache access. Loads or stores deferred due to snoop accesses are performed on the clock cycle following the snoop.

1.1.7 Core Interface

Because the caches are on-chip, write-back caches, the most common transactions are burst-read memory operations, burst-write memory operations, and single-beat (noncacheable or write-through) memory read and write operations. There can also be address-only operations, variants of the burst and single-beat operations, (for example, global memory operations that are snooped and atomic memory operations), and address retry activity (for example, when a snooped read access hits a modified cache block).

Memory accesses can occur in single-beat (1–8 bytes) and four-beat burst (32 bytes) data transfers when the 60x bus is configured as 64 bits, and in single-beat (1–4 bytes), two-beat (8 bytes), and eight-beat (32 bytes) data transfers when the bus is configured as 32 bits. The 60x address and data buses operate independently to support pipelining and split transactions during memory accesses. The core can pipeline its own transactions to a depth of one level.

Access to the system interface is granted through an external arbitration mechanism that allows devices to compete for bus mastership. This arbitration is flexible, allowing the core to be integrated into systems that implement various fairness and bus parking procedures to avoid arbitration overhead.

Typically, memory accesses are weakly ordered—sequences of operations, including load/store string and multiple instructions, do not necessarily complete in the order they begin—maximizing the efficiency of the bus without sacrificing coherency of the data. The core allows read operations to precede store operations (except when a dependency exists, or in cases where a noncacheable access is performed), and provides support for a write operation to proceed a previously queued read data tenure (for example, allowing a snoop push to be enveloped by the address and data tenures of a read operation). Because the processor can dynamically optimize run-time ordering of load/store traffic, overall performance is improved.

1.1.8 System Support Functions

The G2 core implements several support functions that include power management, time base/decrementer registers for system timing tasks, an IEEE 1149.1 (JTAG)/common

on-chip processor (COP) test interface, and a phase-locked loop (PLL) clock multiplier. These system support functions are described in the following sections.

1.1.8.1 Power Management

The G2 core provides four power modes, selectable by setting the appropriate control bits in the machine state register (MSR) and hardware implementation register 0 (HID0). The four power modes are as follows:

- **Full-power**—This is the default power state of the G2 core. The G2 core is fully powered and the internal functional units are operating at the full processor clock speed. If the dynamic power management mode is enabled, functional units that are idle will automatically enter a low-power state without affecting performance, software execution, or external hardware.
- **Doze**—All the functional units of the G2 core are disabled except for the time base/decrementer registers and the bus snooping logic. When the processor is in doze mode, an external asynchronous interrupt, system management interrupt, decrementer exception, hard or soft reset, or machine check brings the G2 core into the full-power state. The core in doze mode maintains the PLL in a fully-powered state and locked to the system external clock input (`core_sysclk`) so a transition to the full-power state takes only a few processor clock cycles.
- **Nap**—The nap mode further reduces power consumption by disabling bus snooping, leaving only the time base register and the PLL in a powered state. The core returns to the full-power state upon receipt of an external asynchronous interrupt, system management interrupt, decrementer exception, hard or soft reset, or machine check input (`core_mcp`) signal. A return to full-power state from a nap state takes only a few processor clock cycles.
- **Sleep**—Sleep mode reduces power consumption to a minimum by disabling all internal functional units; then external system logic may disable the PLL and `core_sysclk`. Returning the core to the full-power state requires the enabling of the PLL and `core_sysclk`, followed by the assertion of an external asynchronous interrupt, system management interrupt, hard or soft reset, or `core_mcp` signal after the time required to relock the PLL.

1.1.8.2 Time Base/Decrementer

The time base is a 64-bit register (accessed as two 32-bit registers) that is incremented once every four bus clock cycles; external control of the time base is provided through the time base enable (`core_tben`) signal. The decrementer is a 32-bit register that generates a decrementer interrupt exception after a programmable delay. The contents of the decrementer register are decremented once every four bus clock cycles, and the decrementer exception is generated as the count passes through zero.

1.1.8.3 IEEE 1149.1 (JTAG)/COP Test Interface

The core provides IEEE 1149.1 and COP functions for facilitating board testing and chip debugging. The IEEE 1149.1 test interface provides a means for boundary-scan testing the core and the attached system logic. The COP function shares the IEEE 1149.1 test port, providing a means for executing test routines, and facilitating chip and software debugging.

The G2_LE core has four additional debug interface signals and three additional breakpoint registers (one instruction and two data breakpoint registers) for debugging purposes. These features expand the functionality of breakpoints and watchpoints. The new breakpoint registers are accessible as SPRs. See Section 1.3.8, “Debug Features (G2_LE Only),” for more information.

There are two additional signals, `core_tap_en` and `core_tlmssel`, which allow multiple JTAG blocks. See Section 8.3.12.6, “TLM TAP Enable (`core_tap_en`)—Input,” and Section 8.3.12.7, “Test Linking Module Select (`core_tlmssel`)—Output,” for more information.

1.1.8.4 Clock Multiplier

The internal clocking of the G2 core is generated from and synchronized to the external clock signal, `core_sysclk`, by means of a voltage-controlled oscillator-based PLL. The PLL provides programmable internal processor clock multiplier ratios which multiply the externally supplied clock frequency. The bus clock is the same frequency and is synchronous with `core_sysclk`. The configuration of the PLL can be read by software from the hardware implementation register 1 (HID1).

1.2 PowerPC Architecture Implementation

The PowerPC architecture consists of the following layers, and adherence to the PowerPC architecture can be measured in terms of which of the following levels of the architecture is implemented:

- User instruction set architecture (UISA)—Defines the base user-level instruction set, user-level registers, data types, floating-point exception model, memory models for a uniprocessor environment, and programming model for a uniprocessor environment.
- Virtual environment architecture (VEA)—Describes the memory model for a multiprocessor environment, defines cache control instructions, and describes other aspects of virtual environments. Implementations that conform to the VEA also adhere to the UISA, but may not necessarily adhere to the OEA.
- Operating environment architecture (OEA)—Defines the memory management model, supervisor-level registers, synchronization requirements, and exception model. Implementations that conform to the OEA also adhere to the UISA and VEA.

The PowerPC architecture allows a wide range of designs for such features as cache and system interface implementations.

1.3 Implementation-Specific Information

The PowerPC architecture is derived from the IBM POWER architecture (Performance Optimized with Enhanced RISC architecture). The PowerPC architecture shares the benefits of the POWER architecture optimized for single-chip implementations. The PowerPC architecture design facilitates parallel instruction execution and is scalable to take advantage of future technological gains.

This section describes the PowerPC architecture in general and specific details about the implementation of the G2 core as a low-power, 32-bit member of this G2 core family. The main topics addressed are as follows:

- Section 1.3.1, “Register Model,” describes the registers for the operating environment architecture common among G2 cores that implement the PowerPC architecture and describes the programming model. It also describes the additional registers that are unique to the core.
- Section 1.3.2, “Instruction Set and Addressing Modes,” describes the PowerPC instruction set and addressing modes for the OEA, and defines and describes the instructions implemented in the core.
- Section 1.3.3, “Cache Implementation,” describes the cache model that is defined generally for cores that implement the PowerPC architecture by the VEA. It also provides specific details about the G2 core cache implementation.
- Section 1.3.4, “Exception Model,” describes the exception model of the OEA and the differences in the core exception model.
- Section 1.3.5, “Memory Management,” describes generally the conventions for memory management among these cores. This section also describes the core implementation of the 32-bit PowerPC memory management specification.
- Section 1.3.6, “Instruction Timing,” provides a general description of the instruction timing provided by the superscalar, parallel execution supported by the PowerPC architecture and the G2 core.
- Section 1.3.7, “System Interface,” describes the signals implemented on the core.

The G2 core is a high-performance, superscalar processor core. The PowerPC architecture allows optimizing compilers to schedule instructions to maximize performance through efficient use of the PowerPC instruction set and register model. The multiple, independent execution units allow compilers to optimize instruction throughput. Compilers that take advantage of the flexibility of the PowerPC architecture can additionally optimize system performance.

The following sections summarize the features of the core, including both those that are defined by the architecture and those that are unique to the various core implementations.

Specific features of the core are listed in Section 1.1.1, “Features.”

1.3.1 Register Model

The PowerPC architecture defines register-to-register operations for most computational instructions. Source operands for these instructions are accessed from the registers or are provided as immediate values embedded in the instruction opcode. The three-register instruction format allows specification of a target register distinct from the two-source operands. Load and store instructions transfer data between registers and memory.

The G2 core has two levels of privilege—supervisor mode of operation (typically used by the operating system) and user mode of operation (used by the application software). The programming models incorporate 32 GPRs, 32 FPRs, special-purpose registers (SPRs), and several miscellaneous registers. Each core also has its own unique set of hardware implementation (HID) registers.

Having access to privileged instructions, registers, and other resources allows the operating system to control the application environment (providing virtual memory and protecting operating system and critical machine resources). Instructions that control the state of the G2 core, the address translation mechanism, and supervisor registers can be executed only when the core is operating in supervisor mode.

Figure 1-2 shows all the core registers available at the user and supervisor level. The numbers to the right of the SPRs indicate the number that is used in the syntax of the instruction operands for the move to/from SPR instructions.

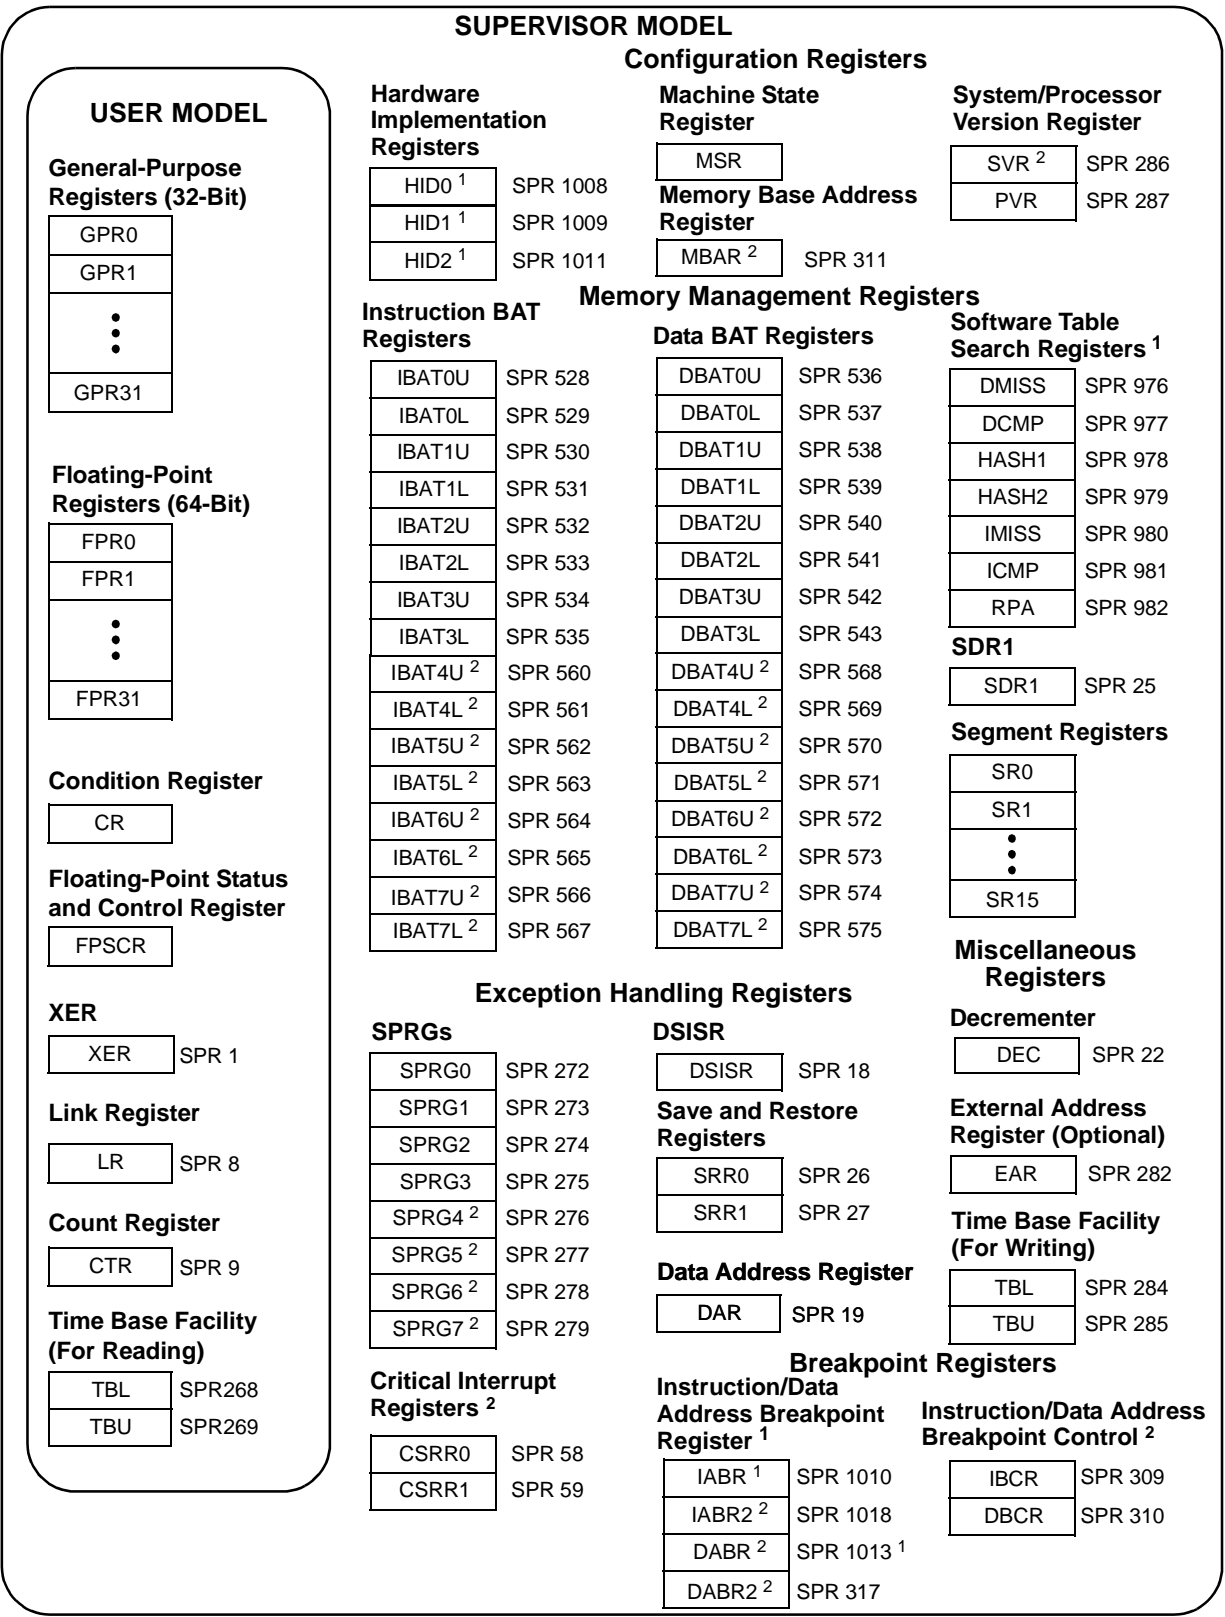
The following sections describe the G2 core implementation-specific features as they apply to registers.

1.3.1.1 General-Purpose Registers (GPRs)

The PowerPC architecture defines 32 user-level GPRs, which are 32 bits wide in 32-bit cores. The GPRs serve as the data source or destination for all integer instructions.

1.3.1.2 Floating-Point Registers (FPRs)

The PowerPC architecture also defines 32 user-level, 64-bit FPRs. The FPRs serve as the data source or destination for floating-point instructions. These registers can contain data objects of either single- or double-precision floating-point formats.



1.3.1.3 Condition Register (CR)

The CR is a 32-bit user-level register that provides a mechanism for testing and branching. It consists of eight 4-bit fields that reflect the results of certain operations, such as move, integer and floating-point comparisons, arithmetic, and logical operations.

1.3.1.4 Floating-Point Status and Control Register (FPSCR)

The user-level FPSCR contains all floating-point exception signal bits, exception summary bits, exception enable bits, and rounding control bits needed for compliance with the IEEE 754 standard.

1.3.1.5 Machine State Register (MSR)

The MSR is a supervisor-level register that defines the state of the core. The contents of this register are saved when an exception is taken and restored when the exception handling completes. A critical interrupt exception is taken only in the G2_LE core when the `core_cint` signal is asserted and MSR[CE] is set. The G2 core implements the MSR as a 32-bit register.

1.3.1.6 Segment Registers (SRs)

For memory management, 32-bit processors implement sixteen 32-bit SRs. To speed access, the core implements the SRs as two arrays; a main array (for data memory accesses) and a shadow array (for instruction memory accesses). Loading a segment entry with the Move to Segment Register (`mtsr`) instruction loads both arrays.

1.3.1.7 Special-Purpose Registers (SPRs)

The OEA defines numerous SPRs that serve a variety of functions, such as providing controls, indicating status, configuring the core, and performing special operations. During normal execution, a program can access the registers, as shown in Figure 1-2, depending on the program's access privilege (supervisor or user, determined by the privilege-level bit, MSR[PR]). Note that GPRs and FPRs are accessed through operands that are part of the instructions. Access to registers can be explicit (that is, through the use of specific instructions for that purpose such as Move to Special-Purpose Register (`mtspr`) and Move from Special-Purpose Register (`mfspir`) instructions) or implicit, as the part of the execution of an instruction. Some registers are accessed both explicitly and implicitly.

The G2_LE core has 29 new/additional supervisor-level SPRs, which are shown in Figure 1-2. Two critical interrupt SPRs (CSRR0 and CSRR1), four additional SPRGs (SPRG4–SPRG7), four pairs of instruction BATs (IBAT4–IBAT7) and four pairs of data BATs (DBAT4–DBAT7), one system version register (SVR), one system memory base address (MBAR), one instruction address breakpoint control (IBCR) and one data address

breakpoint control (DBCR), a new instruction breakpoint register (IABR2), and two data address breakpoint registers (DABR and DABR2) are added to the G2_LE core.

In the G2 core, all SPRs are 32 bits wide.

1.3.1.7.1 User-Level SPRs

The following SPRs are accessible by user-level software:

- Link register (LR)—The LR can be used to provide the branch target address and to hold the return address after branch and link instructions. The LR is 32 bits wide in 32-bit implementations.
- Count register (CTR)—The CTR is decremented and tested automatically as a result of branch-and-count instructions. The CTR is 32 bits wide in 32-bit implementations.
- XER register—The 32-bit XER contains the summary overflow bit, integer carry bit, overflow bit, and a field specifying the number of bytes to be transferred by a Load String Word Indexed (**lswx**) or Store String Word Indexed (**stswx**) instruction.

1.3.1.7.2 Supervisor-Level SPRs

The core also contains SPRs that can be accessed only by supervisor-level software. These registers consist of the following:

- The DSISR defines the cause of data access and alignment exceptions. The cause of a DSI exception for a data breakpoint (match with DABR and DABR2) can be determined by the value of the DSISR[DABR] bit (bit 9).
- The data address register (DAR) holds the address of an access after an alignment or DSI exception. For example, it contains the address of the breakpoint match condition.
- The decremter register (DEC) is a 32-bit decrementing counter that provides a mechanism for causing a decremter exception after a programmable delay.
- SDR1 specifies the page table format used in virtual-to-physical address translation for pages. (Note that physical address is referred to as real address in the architecture specification.)
- The machine status save/restore register 0 (SRR0) is used for saving the address of the instruction that caused the exception, and the address to return to when a Return from Interrupt (**rfi**) instruction is executed.
- The machine status save/restore register 1 (SRR1) is used to save machine status on exceptions and to restore machine status when an **rfi** instruction is executed.
- The SPRG0–SPRG3 registers are provided for operating system use, which reduce the latency that may be incurred because of saving registers to memory while in a handler. Note that G2_LE implements four additional SPRGs.

- The external access register (EAR) controls access to the external control facility through the External Control In Word Indexed (**eciwx**) and External Control Out Word Indexed (**ecowx**) instructions.
- The time base register (TB) is a 64-bit register that maintains the time of day and operates interval timers. It consists of two 32-bit fields—time base upper (TBU) and time base lower (TBL).
- The processor version register (PVR) is a read-only register that identifies the version (model) and revision level of the processor. See Table 1-6 for the version and revision level of the PVR for the G2 and G2_LE processor cores.
- Block address translation (BAT) arrays—The PowerPC architecture defines 16 BAT registers. The G2 core has four pairs of DBAT and IBAT registers. Note that G2_LE supports additional BATs. See Figure 1-2 for a list of the SPR numbers for the BAT arrays.

The following supervisor-level SPRs are implementation-specific (not defined in the PowerPC architecture):

- DMISS and IMISS are read-only registers that are loaded automatically on an instruction or data TLB miss.
- HASH1 and HASH2 contain the physical addresses of the primary and secondary page table entry groups (PTEGs).
- ICMP and DCMP contain a duplicate of the first word in the page table entry (PTE) for which the table search is looking.
- The required physical address (RPA) register is loaded by the core with the second word of the correct PTE during a page table search.
- The hardware implementation (HID0 and HID1) registers provide the means for enabling core checkstops and features, and allows software to read the configuration of the PLL configuration signals. The HID2 register enables the true little-endian mode, cache way-locking, and the additional BAT registers.
- A new system version register (SVR) is added to the G2_LE core, that identifies the specific version (model) and revision level of the system-on-a-chip (SOC) integration.
- System memory base address (MBAR) is a new implementation-specific register for the G2_LE core. It supports a system-level memory map.
- The instruction address breakpoint register (IABR) is loaded with an instruction address that is compared to instruction addresses in the dispatch queue. When an address match occurs, an instruction address breakpoint exception is generated.
- To support critical interrupts, two new registers (CSRR0 and CSRR1) are added to the G2_LE core only.
- Four additional SPRG registers (SPRG4–SPRG7) are in the G2_LE core

- Block address translation (BAT) arrays—The G2_LE core has 16 additional BAT registers (four pairs of DBAT and IBAT registers).
- One additional instruction address breakpoint register (IABR2) and two new data address breakpoint registers (DABR, DABR2) are added to the G2_LE (not in G2 core).
- One instruction breakpoint control (IBCR) and one data breakpoint control (DBCR) are implemented in the G2_LE core (not in G2 core).

1.3.2 Instruction Set and Addressing Modes

The following sections describe the PowerPC instruction set and addressing modes in general.

1.3.2.1 PowerPC Instruction Set and Addressing Modes

All PowerPC instructions are encoded as single-word (32-bit) opcodes. Instruction formats are consistent among all instruction types, permitting efficient decoding to occur in parallel with operand accesses. This fixed instruction length and consistent format simplifies instruction pipelining.

The PowerPC instructions are divided into the following categories:

- Integer instructions—These include computational and logical instructions.
 - Integer arithmetic instructions
 - Integer compare instructions
 - Integer logical instructions
 - Integer rotate and shift instructions
- Floating-point instructions—These include floating-point computational instructions, as well as instructions that affect the FPSCR.
 - Floating-point arithmetic instructions
 - Floating-point multiply/add instructions
 - Floating-point rounding and conversion instructions
 - Floating-point compare instructions
 - Floating-point status and control instructions
- Load/store instructions—These include integer and floating-point load and store instructions.
 - Integer load and store instructions
 - Integer load and store multiple instructions
 - Floating-point load and store

- Primitives used to construct atomic memory operations (**lwarx** and **stwex** instructions)
- Flow control instructions—These include branching instructions, condition register logical instructions, trap instructions, and other instructions that affect the instruction flow.
 - Branch and trap instructions
 - Condition register logical instructions
- Processor control instructions—These instructions are used for synchronizing memory accesses and management of caches, TLBs, and the segment registers.
 - Move to/from SPR instructions
 - Move to/from MSR
 - Synchronize
 - Instruction synchronize
- Memory control instructions—These instructions provide control of caches, TLBs, and segment registers.
 - Supervisor-level cache management instructions
 - Translation lookaside buffer management instructions. Note that there are additional implementation-specific instructions.
 - User-level cache instructions
 - Segment register manipulation instructions
- The G2 core implements the following instructions which are defined as optional by the PowerPC architecture:
 - External Control In Word Indexed (**eciwx**)
 - External Control Out Word Indexed (**ecowx**)
 - Floating Select (**fsel**)
 - Floating Reciprocal Estimate Single-Precision (**fres**)
 - Floating Reciprocal Square Root Estimate (**frsqrte**)
 - Store Floating-Point as Integer Word (**stfiwx**)

Note that this grouping of instructions does not indicate the execution unit that executes a particular instruction or group of instructions.

Integer instructions operate on byte, half-word, and word operands. Floating-point instructions operate on single-precision (one word) and double-precision (one double word) floating-point operands. The PowerPC architecture uses instructions that are 4 bytes long and word-aligned. It provides for byte, half-word, and word operand loads and stores between memory and a set of 32 GPRs. It also provides for word and double-word operand loads and stores between memory and a set of 32 FPRs.

Computational instructions do not modify memory. To use a memory operand in a computation and then modify the same or another memory location, the memory contents must be loaded into a register, modified, and then written back to the target location with distinct instructions.

The G2 core follows the program flow when it is in the normal execution state. However, the flow of instructions can be interrupted directly by the execution of an instruction or by an asynchronous event. Either kind of exception may cause one of several components of the system software to be invoked.

1.3.2.2 Implementation-Specific Instruction Set

The G2 core instruction set is defined as follows:

- The core provides hardware support for all 32-bit PowerPC instructions.
- The core provides two implementation-specific instructions used for software table search operations following TLB misses:
 - Load Data TLB Entry (**tlbld**)
 - Load Instruction TLB Entry (**tlbli**)
- The G2_LE implements the following instruction which is added to support critical interrupts. This is a supervisor-level, context synchronizing instruction.
 - Return from Critical Interrupt (**rftci**)

1.3.3 Cache Implementation

The following sections describe the general cache characteristics as implemented in the PowerPC architecture and the core implementation, specifically. G2_LE-specific information is noted where applicable.

1.3.3.1 PowerPC Cache Characteristics

The PowerPC architecture does not define hardware aspects of cache implementations. The G2 core controls the following memory access modes on a page or block basis:

- Write-back/write-through mode
- Caching-inhibited mode
- Memory coherency

Note that in the core, a cache block is defined as eight words. The VEA defines cache management instructions that provide a means by which the application programmer can affect the cache contents.

1.3.3.2 Implementation-Specific Cache Implementation

The G2 core has two 16-Kbyte, four-way set-associative (instruction and data) caches. The caches are physically addressed, and the data cache can operate in either write-back or write-through mode as specified by the PowerPC architecture.

The data cache is configured as 128 sets of 4 blocks each. Each block consists of 32 bytes, 2 state bits, and an address tag. The two state bits implement the three-state MEI (modified/exclusive/invalid) protocol. Each block contains eight 32-bit words. Note that the PowerPC architecture defines the term ‘block’ as the cacheable unit. For the core, the block size is equivalent to a cache line. A block diagram of the data cache organization is shown in Figure 1-3.

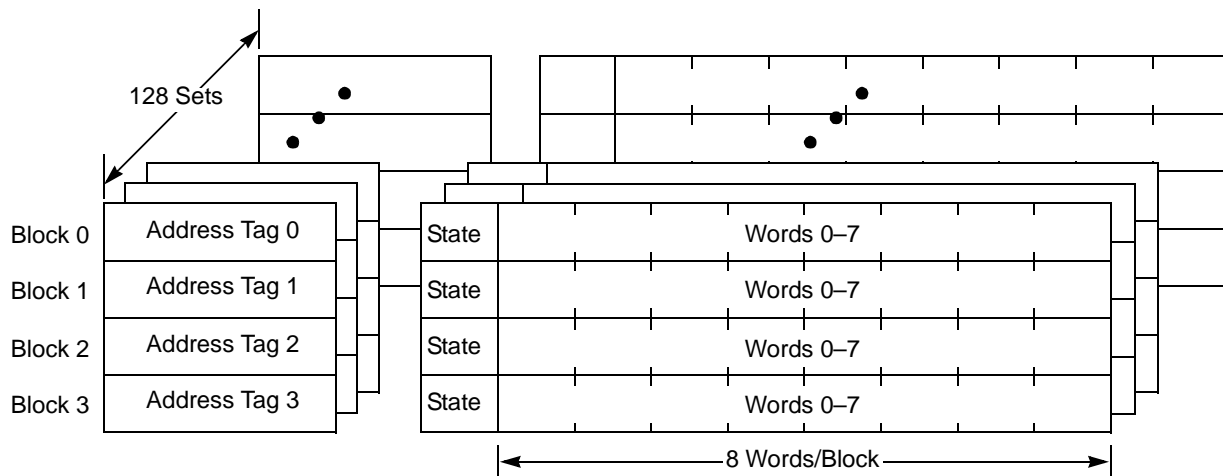


Figure 1-3. Data Cache Organization

The instruction cache also consists of 128 sets of 4 blocks, and each block consists of 32 bytes, an address tag, and a valid bit. The instruction cache may not be written to, except through a block fill operation. In the G2 core, the instruction cache is blocked only until the critical load completes. The G2 core supports instruction fetching from other instruction cache lines following the forwarding of the critical-first-double-word of a cache line load operation. Successive instruction fetches from the cache line being loaded are forwarded, and accesses to other instruction cache lines can proceed during the cache line load operation. The instruction cache is not snooped, and cache coherency must be maintained by software. A fast hardware invalidation capability is provided to support cache maintenance. The organization of the instruction cache is very similar to the data cache shown in Figure 1-3.

Each cache block contains eight contiguous words from memory that are loaded from an 8-word boundary (that is, bits A[27–31] of the effective addresses are zero); thus, a cache block never crosses a page boundary. Misaligned accesses across a page boundary can incur a performance penalty.

The G2 core cache blocks are loaded in four beats of 64 bits each when the core is configured with a 64-bit 60x data bus. When the core is configured with a 32-bit bus, cache block loads are performed with eight beats of 32 bits each. The burst load is performed as critical-double-word-first. The data cache is blocked to internal accesses until the load completes; the instruction cache allows sequential fetching during a cache block load. In the core, the critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.

To ensure coherency among caches in a multiprocessor (or multiple caching-device) implementation, the core implements the MEI protocol. The following three states indicate the state of the cache block:

- **Modified**—The cache block is modified with respect to system memory; that is, data for this address is valid only in the cache and not in system memory.
- **Exclusive**—This cache block holds valid data that is identical to the data at this address in system memory. No other cache has this data.
- **Invalid**—This cache block does not hold valid data.

Cache coherency is enforced by on-chip bus snooping logic. Because the G2 core data cache tags are single-ported, a simultaneous load or store and snoop access represents a resource contention. The snoop access is given first access to the tags. The load or store then occurs on the clock following the snoop.

1.3.3.3 Instruction and Data Cache Way-Locking

The G2 core implements instruction and data cache way-locking, which guarantees that certain memory accesses will hit in the cache. This provides deterministic access times for those accesses. See Chapter 4, “Instruction and Data Cache Operation,” for more information.

1.3.4 Exception Model

This section describes the PowerPC exception model and the G2 core implementation, specifically. G2_LE core-specific information is noted where applicable.

1.3.4.1 PowerPC Exception Model

The PowerPC exception mechanism allows the core to change to supervisor state as a result of external signals, errors, or unusual conditions arising in the execution of instructions, and differs from the arithmetic exceptions defined by the IEEE for floating-point operations. When exceptions occur, information about the state of the core is saved to certain registers and the core begins execution at an address (exception vector) predetermined for each exception type. Processing of exceptions occurs in supervisor mode.

Although multiple exception conditions can map to a single exception vector, a more specific condition may be determined by examining a register associated with the exception—for example, the DSISR and the FPSCR. Additionally, some exception conditions can be explicitly enabled or disabled by software.

The PowerPC architecture requires that exceptions be handled in program order; therefore, although a particular implementation may recognize exception conditions out of order, they are presented strictly in order. When an instruction-caused exception is recognized, any unexecuted instructions that appear earlier in the instruction stream, including any that have not yet entered the execute stage, are required to complete before the exception is taken. Any exceptions caused by those instructions are handled first. Likewise, exceptions that are asynchronous and precise are recognized when they occur, but are not handled until the instruction currently in the completion stage successfully completes execution or generates an exception, and the completed store queue is emptied.

Unless a catastrophic condition causes a system reset or machine check exception, only one exception is handled at a time. If, for example, a single instruction encounters multiple exception conditions, those conditions are handled sequentially. After the exception handler handles an exception, the instruction execution continues until the next exception condition is encountered. However, in many cases there is no attempt to re-execute the instruction. This method of recognizing and handling exception conditions sequentially guarantees that exceptions are recoverable.

Exception handlers should save the information stored in SRR0 and SRR1 early to prevent the program state from being lost due to a system reset or machine check exception or to an instruction-caused exception in the exception handler, and before enabling external interrupts.

The PowerPC architecture supports four types of exceptions:

- Synchronous, precise—These are caused by instructions. All instruction-caused exceptions are handled precisely; that is, the machine state at the time the exception occurs is known and can be completely restored. This means that (excluding the trap and system call exceptions) the address of the faulting instruction is provided to the exception handler and neither the faulting instruction nor subsequent instructions in the code stream will complete execution before the exception is taken. Once the exception is processed, execution resumes at the address of the faulting instruction (or at an alternate address provided by the exception handler). When an exception is taken due to a trap or system call instruction, execution resumes at an address provided by the handler.
- Synchronous, imprecise—The PowerPC architecture defines two imprecise floating-point exception modes, recoverable and nonrecoverable. Even though the G2 core provides a means to enable the imprecise modes, it implements these modes identically to the precise mode (that is, all enabled floating-point enabled exceptions are always precise on the core).

- Asynchronous, maskable—The external, system management interrupt (SMI), and decremter interrupts are maskable asynchronous exceptions. When these exceptions occur, their handling is postponed until the next instruction, and any exceptions associated with that instruction, completes execution. If there are no instructions in the execution units, the exception is taken immediately on determination of the correct restart address (for loading SRR0).
- Asynchronous, nonmaskable—There are two nonmaskable asynchronous exceptions: system reset and the machine check exception. These exceptions may not be recoverable, or may provide a limited degree of recoverability. All exceptions report recoverability through MSR[RI].

1.3.4.2 Implementation-Specific Exception Model

As specified by the PowerPC architecture, all exceptions can be described as either precise or imprecise and either synchronous or asynchronous. Asynchronous exceptions (some of which are maskable) are caused by events external to the processor’s execution; synchronous exceptions, which are all handled precisely by the G2 core, are caused by instructions. A system management interrupt is an implementation-specific exception. The exception classes are shown in Table 1-3. The exceptions are listed in Table 5-3 in order of highest to lowest priority.

Table 1-3. Exception Classifications

Synchronous/Asynchronous	Precise/Imprecise	Exception Type
Asynchronous, nonmaskable	Imprecise	Machine check System reset
Asynchronous, maskable	Precise	External interrupt Decrementer System management interrupt Critical interrupt
Synchronous	Precise	Instruction-caused exceptions

Although exceptions have other characteristics as well, such as whether they are maskable or nonmaskable, the distinctions shown in Table 1-3 define categories of exceptions that the core handles uniquely. Note that Table 1-3 includes no synchronous imprecise instructions. While the PowerPC architecture supports imprecise handling of floating-point exceptions, the core implements floating-point exception modes as precise exceptions.

The G2 core exceptions, and conditions that cause them, are listed in Table 1-4.

Table 1-4. Exceptions and Conditions

Exception Type	Vector Offset (hex)	Causing Conditions
Reserved	00000	—
System reset	00100	A system reset is caused by the assertion of either <code>core_sreset</code> or <code>core_hreset</code> .
Machine check	00200	A machine check is caused by the assertion of the <code>core_tea</code> signal during a data bus transaction, assertion of <code>core_mcp</code> , or an address or data parity error.
DSI	00300	<p>The cause of a DSI exception can be determined by the bit settings in the DSISR, listed as follows:</p> <ol style="list-style-type: none"> 1 Set if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared. 4 Set if a memory access is not permitted by the page or DBAT protection mechanism; otherwise cleared. 5 Set by an <code>eciwX</code> or <code>ecowX</code> instruction if the access is to an address that is marked as write-through, or execution of a load/store instruction that accesses a direct-store segment. 6 Set for a store operation and cleared for a load operation 9 G2_LE core only. Set a data address breakpoint exception when the data (bit 0–28) in the DABR1 or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> • Write breakpoints enabled when DABR[30] is set • Read breakpoints enabled when DABR[31] is set 11 Set if <code>eciwX</code> or <code>ecowX</code> is used and EAR[E] is cleared
ISI	00400	<p>An ISI exception is caused when an instruction fetch cannot be performed for any of the following reasons:</p> <ul style="list-style-type: none"> • The effective (logical) address cannot be translated. That is, there is a page fault for this portion of the translation, so an ISI exception must be taken to load the PTE (and possibly the page) into memory. • The fetch access is to a direct-store segment (indicated by SRR1[3] set). • The fetch access violates memory protection (indicated by SRR1[4] set). If the key bits (Ks and Kp) in the segment register and the PP bits in the PTE are set to prohibit read access, instructions cannot be fetched from this location.
External interrupt	00500	An external interrupt is caused when MSR[EE] = 1 and the <code>core_int</code> signal is asserted.
Alignment	00600	<p>An alignment exception is caused when the core cannot perform a memory access for any of the reasons described below:</p> <ul style="list-style-type: none"> • The operand of a floating-point load or store instruction is not word-aligned. • The operand of <code>lmw</code>, <code>stmw</code>, <code>lwarx</code>, and <code>stwcx</code> instructions are not aligned. • The execution of a floating-point load or store instruction to a direct-store segment. • The operand of a load, store, load multiple, store multiple, load string, or store string instruction crosses a segment boundary into a direct-store segment, or crosses a protection boundary. • Execution of a misaligned <code>eciwX</code> or <code>ecowX</code> instruction. • The instruction is <code>lmw</code>, <code>stmw</code>, <code>lswi</code>, <code>lswx</code>, <code>stswi</code>, <code>stswx</code>, and the G2 core is in little-endian mode. It applies to both PowerPC little-endian and LE mode for G2_LE core. • The operand of <code>dcbz</code> is in memory that is write-through-required or caching-inhibited.

Table 1-4. Exceptions and Conditions (continued)

Exception Type	Vector Offset (hex)	Causing Conditions
Program	00700	<p>A program exception is caused by one of the following exception conditions, which correspond to bit settings in SRR1 and arise during execution of an instruction:</p> <p>Floating-point enabled exception—A floating-point enabled exception condition is generated when the following condition is met:</p> <p>(MSR[FE0] MSR[FE1]) & FPSCR[FEX] is 1.</p> <ul style="list-style-type: none"> • FPSCR[FEX] is set by the execution of a floating-point instruction that causes an enabled exception or by the execution of one of the ‘move to FPSCR’ instructions that results in both an exception condition bit and its corresponding enable bit being set in the FPSCR. • Illegal instruction—An illegal instruction program exception is generated when execution of an instruction is attempted with an illegal opcode or illegal combination of opcode and extended opcode fields (including PowerPC instructions not implemented in the core), or when execution of an optional instruction not provided in the core is attempted (these do not include those optional instructions that are treated as no-ops). • Privileged instruction—A privileged instruction type program exception is generated when the execution of a privileged instruction is attempted and the MSR register user privilege bit, MSR[PR], is set. In the G2 core, this exception is generated for mtspr or mfspir with an invalid SPR field if SPR[0] = 1 and MSR[PR] = 1. This may not be true for all cores that implement the PowerPC architecture. • Trap—A trap type program exception is generated when any of the conditions specified in a trap instruction is met.
Floating-point unavailable	00800	A floating-point unavailable exception is caused by an attempt to execute a floating-point instruction (including floating-point load, store, and move instructions) when the floating-point available bit is cleared (MSR[FP] = 0).
Decrementer	00900	The decrementer exception occurs when DEC[31] changes from 0 to 1. This exception is also enabled with MSR[EE].
Critical interrupt	00A00	A critical interrupt exception is taken when the $\overline{\text{core_cint}}$ signal is asserted and MSR[CE] = 1 (G2_LE only).
Reserved	00B00–00BFF	—
System call	00C00	A system call exception occurs when a System Call (sc) instruction is executed.
Trace	00D00	A trace exception is taken when MSR[SE] = 1 or when the currently completing instruction is a branch and MSR[BE] = 1.
Reserved	00E00	The G2 core does not generate an exception to this vector. Other devices may use this vector for floating-point assist exceptions.
Reserved	00E10–00FFF	—
Instruction translation miss	01000	An instruction translation miss exception is caused when the effective address for an instruction fetch cannot be translated by the ITLB.
Data load translation miss	01100	A data load translation miss exception is caused when the effective address for a data load operation cannot be translated by the DTLB.
Data store translation miss	01200	A data store translation miss exception is caused when the effective address for a data store operation cannot be translated by the DTLB, or where a DTLB hit occurs, and the change bit in the PTE must be set due to a data store operation.

Table 1-4. Exceptions and Conditions (continued)

Exception Type	Vector Offset (hex)	Causing Conditions
Instruction address breakpoint	01300	An instruction address breakpoint exception occurs when the address (bits 0–29) in the IABR matches the next instruction to complete in the completion unit, and IABR[bit 30] is set.
System management interrupt	01400	A system management interrupt is caused when MSR[EE] = 1 and the $\overline{\text{core_smi}}$ input signal is asserted.
Reserved	01500–02FFF	—

1.3.5 Memory Management

The following sections describe the memory management features of the PowerPC architecture and the G2 core implementation, respectively.

1.3.5.1 PowerPC Memory Management

The primary functions of the MMU are to translate logical (effective) addresses to physical addresses for memory accesses and to provide access protection on blocks and pages of memory.

The core generates two types of accesses that require address translation—instruction accesses and data accesses to memory generated by load and store instructions.

The PowerPC MMU and exception model support demand-paged virtual memory. Virtual memory management permits execution of programs larger than the size of physical memory; demand-paged implies that individual pages are loaded into physical memory from system memory only when they are first accessed by an executing program.

The hashed page table is a variable-sized data structure that defines the mapping between virtual page numbers and physical page numbers. The page table size is a power of two, and its starting address is a multiple of its size.

The page table contains a number of page table entry groups (PTEGs). A PTEG contains eight page table entries (PTEs) of 8 bytes each; therefore, each PTEG is 64 bytes long. PTEG addresses are entry points for table search operations.

Address translations are enabled by setting bits in the MSR—MSR[IR] enables instruction address translations and MSR[DR] enables data address translations.

1.3.5.2 Implementation-Specific Memory Management

The instruction and data memory management units in the G2 core provide 4 Gbytes of logical address space accessible to supervisor and user programs with a 4-Kbyte page size and 256-Mbyte segment size. Block sizes range from 128 Kbytes to 256 Mbytes and are

software selectable. In addition, the core uses an interim 52-bit virtual address and hashed page tables for generating 32-bit physical addresses. The MMUs in the G2 core rely on the exception processing mechanism for the implementation of the paged virtual memory environment and for enforcing protection of designated memory areas.

Instruction and data TLBs provide address translation in parallel with the on-chip cache access, incurring no additional time penalty in the event of a TLB hit. A TLB is a cache of the most recently used page table entries. Software is responsible for maintaining the consistency of the TLB with memory. The core TLBs are 64-entry, two-way set-associative caches that contain instruction and data address translations. The core provides hardware assist for software table search operations through the hashed page table on TLB misses. Supervisor software can invalidate TLB entries selectively.

For instructions and data that maintain address translations for blocks of memory, the G2 core and the G2_LE core provide independent four- and eight-entry BAT arrays, respectively. These entries define blocks that can vary from 128 Kbytes to 256 Mbytes. The BAT arrays are maintained by system software. HID2[HBE] is added to the G2_LE for enabling or disabling the four additional pairs of BAT registers. However, regardless of the setting of HID2[HBE], these BATs are accessible by **mfspr** and **mtspr**.

As specified by the PowerPC architecture, the hashed page table is a variable-sized data structure that defines the mapping between virtual page numbers and physical page numbers. The page table size is a power of two, and its starting address is a multiple of its size.

Also as specified by the PowerPC architecture, the page table contains a number of PTEGs. A PTEG contains 8 PTEs of 8 bytes each; therefore, each PTEG is 64 bytes long. PTEG addresses are entry points for table search operations.

1.3.6 Instruction Timing

The G2 core is a pipelined superscalar processor core. Because instruction processing is reduced into a series of stages, an instruction does not require all of the resources of an execution unit at the same time. For example, after an instruction completes the decode stage, it can pass on to the next stage, while the subsequent instruction can advance into the decode stage. This improves the throughput of the instruction flow. For example, it may take three cycles for a single floating-point instruction to execute, but if there are no stalls in the floating-point pipeline, a series of floating-point instructions can have a throughput of one instruction per cycle.

The core instruction pipeline has four major pipeline stages, described as follows:

- The fetch pipeline stage primarily involves retrieving instructions from the memory system and determining the location of the next instruction fetch. Additionally, if possible, the BPU decodes branches during the fetch stage and folds out branch instructions before the dispatch stage.

- The dispatch pipeline stage is responsible for decoding the instructions supplied by the instruction fetch stage, and determining which of the instructions are eligible to be dispatched in the current cycle. In addition, the source operands of the instructions are read from the appropriate register file and dispatched with the instruction to the execute pipeline stage. At the end of the dispatch pipeline stage, the dispatched instructions and their operands are latched by the appropriate execution unit.
- In the execute pipeline stage, each execution unit with an executable instruction executes the selected instruction (perhaps over multiple cycles), writes the instruction's result into the appropriate rename register, and notifies the completion stage when the execution has finished. In the case of an internal exception, the execution unit reports the exception to the completion/write-back pipeline stage and discontinues instruction execution until the exception is handled. The exception is not signaled until that instruction is the next to be completed. Execution of most floating-point instructions is pipelined within the FPU allowing up to three instructions to be executing in the FPU concurrently. The FPU pipeline stages are multiply, add, and round-convert. The LSU has two pipeline stages. The first stage is for effective address calculation and MMU translation, and the second is for accessing data in the cache.
- The complete/write-back pipeline stage maintains the correct architectural machine state and transfers the contents of the rename registers to the GPRs and FPRs as instructions are retired. If the completion logic detects an instruction causing an exception, all following instructions are canceled, their execution results in rename registers are discarded, and instructions are fetched from the correct instruction stream.

A superscalar processor core issues multiple independent instructions into multiple pipelines allowing instructions to execute in parallel. The G2 core has five independent execution units, one each for integer instructions, floating-point instructions, branch instructions, load/store instructions, and system register instructions. The IU and the FPU each have dedicated register files for maintaining operands (GPRs and FPRs, respectively), allowing integer and floating-point calculations to occur simultaneously without interference. Integer division performance of the G2 core has been improved, with the **divwux** and **divwx** instructions executing in 20 clock cycles instead of the 37 cycles required in the MPC603e.

The core provides support for single-cycle store and it provides an adder/comparator in the system register unit that allows the dispatch and execution of multiple integer add and compare instructions on each cycle. Refer to Chapter 7, “Instruction Timing,” for more information.

Because the PowerPC architecture can be applied to such a wide variety of implementations, instruction timing among processor cores varies accordingly.

1.3.7 System Interface

The system interface is specific for each processor core implementation.

The G2 core provides a versatile system interface that allows for a wide range of implementations. The interface includes a 32-bit 60x address bus, a 32- or 64-bit 60x data bus, and 56 control and information signals (see Figure 1-4). The system interface allows for address-only transactions, as well as address and data transactions. The core control and information signals include the address arbitration, address start, address transfer, transfer attribute, address termination, data arbitration, data transfer, data termination, and core state signals. Test and control signals provide diagnostics for selected internal circuits.

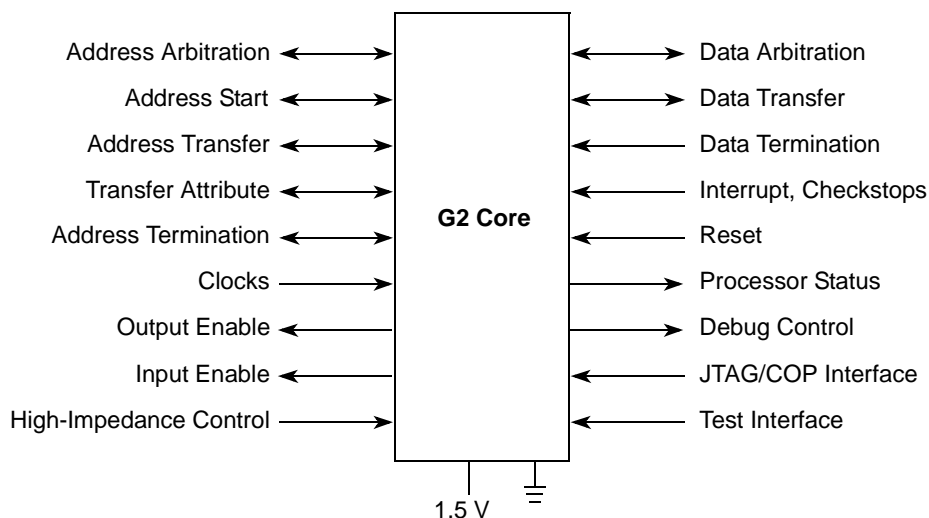


Figure 1-4. System Interface

The system interface supports bus pipelining, allowing the address tenure of one transaction to overlap the data tenure of another. The extent of the pipelining depends on external arbitration and control circuitry. Similarly, the core supports split-bus transactions for systems with multiple potential bus masters—one device can have mastership of the address bus while another has mastership of the data bus. Allowing multiple bus transactions to occur simultaneously increases the available bus bandwidth for other activity, and as a result, improves performance.

The G2 core supports multiple masters through a bus arbitration scheme that allows various devices to compete for the shared bus resource. Arbitration logic can implement priority protocols, such as fairness, and can park masters to avoid arbitration overhead. The MEI protocol ensures coherency among multiple devices and system memory. Also, the core on-chip caches, TLBs, and optional second-level caches can be controlled externally.

The core clocking structure allows the bus to operate at integer multiples of the core cycle time.

The following sections describe the core bus support for memory operations. Note that some signals perform different functions depending on the addressing protocol used.

1.3.7.1 Memory Accesses

The G2 core 60x bus is configured at power-up to either a 32- or 64-bit width.

- When the core is configured with a 32-bit 60x bus, memory accesses allow transfer sizes of 8, 16, 24, or 32 bits in one bus clock cycle. Data transfers occur in either single-beat transactions, two-beat or eight-beat burst transactions, with a single-beat transaction transferring as many as 32 bits. Single- or double-beat transactions are caused by noncached accesses that access memory directly (that is, reads and writes when caching is disabled, caching-inhibited accesses, and stores in write-through mode). Eight-beat burst transactions, which always transfer an entire cache block (32 bytes), are initiated when a line is read from or written to memory.
- When the core is configured with a 64-bit 60x bus, memory accesses allow transfer sizes of 8, 16, 24, 32, 40, 48, 56, or 64 bits in one bus clock cycle. Data transfers occur in either single-beat transactions or four-beat burst transactions. Single-beat transactions are caused by noncached accesses that access memory directly (that is, reads and writes when caching is disabled, caching-inhibited accesses, and stores in write-through mode). Four-beat burst transactions, which always transfer an entire cache block (32 bytes), are initiated when a line is read from or written to memory.

1.3.7.2 Signals

The G2 core signals are grouped as follows:

- Address arbitration signals—The G2 core uses these signals to arbitrate for 60x address bus mastership.
- Address transfer start signals—These signals indicate that a bus master has begun a transaction on the address bus of the 60x bus.
- Address transfer signals—These signals, consisting of the address bus, address parity, and address parity error signals, are used to transfer the address and to ensure the integrity of the transfer.
- Transfer attribute signals—These signals provide information about the type of transfer, such as the transfer size and whether the transaction is bursted, write-through, or caching-inhibited.
- Address transfer termination signals—These signals are used to acknowledge the end of the address phase of the transaction. They also indicate whether a condition exists that requires the address phase to be repeated.
- Data arbitration signals—The G2 core uses these signals to arbitrate for 60x data bus mastership.

- Data transfer signals—These signals, consisting of the data bus, data parity, and data parity error signals, are used to transfer the data and to ensure the integrity of the transfer.
- Data transfer termination signals—Data termination signals are required after each data beat in a data transfer. In a single-beat transaction, the data termination signals also indicate the end of the tenure. In burst accesses, the data termination signals apply to individual beats and indicate the end of the tenure only after the final data beat. They also indicate whether a condition exists that requires the data phase to be repeated.
- Output enable signals—These output signals indicate that the corresponding outputs of the G2 core are driving, provided the corresponding high-impedance control signal is also asserted.
- High-impedance control signals—These input signals (static) enable the operation of the output enable signals.
- Input enable signals—These output signals indicate that the corresponding input signals are being received by the core, provided the corresponding high-impedance control signal is also asserted.
- System status signals—These signals include the external interrupt signals, checkstop signals, and both soft- and hard-reset signals. These signals are used to interrupt and, under various conditions, to reset the core.
- Reset configuration signals—These signals are sampled while `core_hreset` is asserted and they control certain modes of operation.
- JTAG/COP interface signals—The JTAG (IEEE 1149.1) interface and common on-chip processor (COP) unit provides a serial interface to the system for performing monitoring and boundary tests.
- Processor status—These signals include the memory reservation signal, machine quiesce control signals, time base enable signal, and `core_tlbisync` signal.
- Clock signals—These signals provide for system clock input and frequency control.
- Test interface signals—Signals like address matching, combinational matching and watchpoint are used in the `G2_LE` for production testing.

Seven additional signals are added to the `G2_LE` core to support true little-endian mode (`core_tle`), the critical interrupt function (`core_cint`), the breakpoint state outputs (`core_iabr`, `core_iabr2`, `core_dabr`, and `core_dabr2`), and the debug features (`core_tdo_oe`).

NOTE

A bar over a signal name indicates that the signal is active low—for example, core_artry_in (address retry) and core_ts_in (transfer start). Active-low signals are referred to as asserted (active) when they are low and negated when they are high. Signals that are not active low, such as core_ap_in[0:3] (address bus parity signals) and core_tt_in[0:4] (transfer type signals) are referred to as asserted when they are high and negated when they are low.

1.3.8 Debug Features (G2_LE Only)

Some new debug features are specific to the G2_LE core. Accesses to the debug facilities are available only in supervisor mode by using the **mtspr** and **mfspr** instructions. The G2_LE provides the following additional features in the JTAG/COP interface:

- Addition of three breakpoint registers—IABR2, DABR, and DABR2
- Two new breakpoint control registers—DBCR and IBCR
- Inclusion of four breakpoint signals—core_iabr, core_iabr2, core_dabr, and core_dabr2

If instruction or data breakpoints are set to match with any exception vector, an unrecoverable state occurs. Also, instruction or data breakpoints must not be set to match any address used in the breakpoint exception handlers. A breakpoint that matches within an exception handler can cause an indeterminate or unrecoverable processor state.

1.3.8.1 Instruction Address Breakpoint Registers (IABR and IABR2)

IABR and IABR2 can be used to cause a breakpoint exception if a specified instruction address is encountered. IABR and IABR2 control the instruction address breakpoint exception. IABR[CEA] holds an effective address to which each instruction's address is compared. The exception is enabled by setting IABR[30]. The exception is taken when there is an instruction address breakpoint match on the next instruction to complete.

The instruction address match does not complete before the breakpoint exception is taken but the address of that instruction is stored in SRR0. Upon execution of an **rfi** instruction, the instruction addressed in SRR0 is retired, meaning that the results are committed to the destination registers or memory address.

Note that IABR is implemented in both the G2 core and the G2_LE core; IABR2 is an implementation-specific register for the G2_LE core only. Also, note that IBCR gives further control of instruction breakpoints for the G2_LE core.

1.3.8.2 Data Address Breakpoint Registers (DABR and DABR2)

DABR and DABR2 cause a breakpoint exception (subset of the DSI exception) if there is a match between the CEA field and the address of any data access and the data breakpoint is enabled. DABR[CEA] and DABR2[CEA] hold an effective address to which each data access address is compared. In addition, data breakpoints are enabled for write and read accesses individually by setting bit 30 and bit 31 of the DABR, respectively. Finally, the data address breakpoint translation bit (DABR[BT]) must match MSR[DR] for a match to occur.

The data access that causes a match is not performed before the data breakpoint exception is taken. When the exception occurs, the DAR is set to the address of the data access that caused the breakpoint, and DSISR[9] is set. The address of the instruction associated with the matching data access is saved in SRR0. Upon execution of an **rfi** instruction, the instruction addressed in SRR0 is retired, and all results are committed to the destination address in memory.

1.3.8.3 Breakpoint Signaling

The breakpoint signaling provided on the G2_LE core allows observability of breakpoint matches external to the core. The `core_iabr`, `core_iabr2`, `core_dabr`, and `core_dabr2` breakpoint signals are asserted for at least one bus clock cycle when the respective breakpoint occurs.

- When DBCR and IBCR are configured for an OR combinational signal type, the breakpoint signals `core_iabr`, `core_iabr2` and `core_dabr`, `core_dabr2` reflect their respective breakpoints.
- When the DBCR and IBCR are configured for AND combinational signal type, only the `core_iabr2` and `core_dabr2` breakpoint signals are asserted after the AND condition is met (both instruction breakpoints occurred or both data breakpoints occurred).

The breakpoint signaling conditions are described in Chapter 11, “Debug Features.”

1.3.8.4 Other Debug Resources

In addition to the four breakpoint registers and two breakpoint control registers, other internal register values control and observe the effects of breakpoint conditions. Table 1-5 shows these registers and their bits.

Table 1-5. Other Debug and Support Register Bits

Register	Bits	Name	Description
MSR	17	PR	Privilege level. Breakpoint registers can only be accessed when this bit is cleared (supervisor mode).
	21	SE	Single-step trace enable 0 The processor executes instructions normally 1 The processor generates a trace exception on the successful completion of the next instruction
	22	BE	Branch trace enable 0 The processor executes branch instructions normally 1 The processor generates a trace exception on the successful completion of a branch instruction
HIDO	0–31	—	See Table 2-5 for details
DAR	0–31	—	Data address register. DAR is loaded with the effective address of a data breakpoint condition that matches.
DSISR	9	DABR	Set if DABR exception occurs

1.4 Differences Between the MPC603e and the G2 and G2_LE Cores

Table 1-6 describes the differences between the MPC603e and the G2 and G2_LE cores. Note that the G2 core has similar functionality to the MPC603e processor. However, the minor differences between them are documented by footnotes.

Table 1-6. Differences Between G2 and G2_LE Cores

G2 Core	G2_LE Core	Impact
New PVR register value ¹	New PVR register value	The G2 core version number is 0x8081 and the revision level starts at 0x1010 and changes for each revision of the core. The G2_LE core version number is 0x8082 and the revision level starts at 0x1010 and changes for each revision of the core.
Big-endian or modified little-endian modes	core_tle is a new signal for enabling true little-endian mode at reset	True little-endian mode (for G2_LE only) for compatibility with other true little-endian devices. True little-endian mode is supported in the G2_LE core to minimize the impact on software porting from true little-endian systems. Unlike other devices that implement the PowerPC architecture, G2_LE supports true big-endian, true little-endian, and modified little-endian mode of operations.
Only one external interrupt signal (core_int)	An additional input interrupt signal, core_cint, implements a critical interrupt function.	MSR[CE] is allocated for enabling the critical interrupt
—	A new instruction is implemented for critical interrupt	Return from Critical Interrupt (rfci) is implemented to return from these exception handlers

Table 1-6. Differences Between G2 and G2_LE Cores (continued)

G2 Core	G2_LE Core	Impact
—	Vector offset for critical interrupt	An exception vector offset of 0x00A00 is defined for critical interrupt
—	Two new registers are implemented for saving processor state for critical interrupts	CSRR0 and CSRR1 have the same bit assignments as SRR0 and SRR1, respectively.
Supports instruction cache way-locking in addition to entire instruction cache locking	Supports instruction cache way-locking in addition to entire instruction cache locking	HID2 register controls instruction cache way-locking. The instruction cache way-locking is useful for locking blocks of instructions into the instruction cache for time-critical applications that require deterministic behavior.
Supports data cache way-locking in addition to entire data cache locking	Supports data cache way-locking in addition to entire data cache locking	HID2 register controls data cache way-locking. It is useful for locking blocks of data into the data cache for time-critical applications where deterministic behavior is required.
SPRG0–SPRG3 are the four SPRG registers in the MPC603e and the G2 core	Four additional SPRG registers are implemented in G2_LE core only	The additional SPRGs reduce latencies that may be incurred from saving registers to memory while in an exception handler
G2 core has five JTAG/COP interface signals	One additional JTAG/COP interface signal is implemented in the G2_LE	The core_tdo_oe output signal is used for debugging. Note that core_tdo is always driven, regardless of the state of core_tdo_oe.
Instruction address breakpoint exception is controlled by IABR	Instruction address breakpoint exception is controlled by IABR and IABR2	Instruction address breakpoint exceptions in both the G2 and the G2_LE cores use the 0x01300 vector offset
—	Two new data address breakpoint registers are implemented in the G2_LE	The two new data address breakpoint registers (DABR and DABR2) expand the debug functionality of the breakpoints. The new breakpoint registers are accessible as SPRs with mtspr and mfspir .
—	One instruction register and one data breakpoint control register are implemented	IBCR and DBCR are implemented to support the additional debug features. These registers are accessible as SPRs with mtspr and mfspir .
—	Breakpoint signals are implemented for debug	Breakpoint signals— <u>core_iabr</u> , <u>core_iabr2</u> , <u>core_dabr</u> , <u>core_dabr2</u> —are asserted to indicate a breakpoint condition as programmed in DBCR and IBCR. These signals may be OR'd or AND'd to reflect the respective breakpoints.
—	Vector offset for data address breakpoint exception is 0x00300	Data address breakpoint exception is a DSI exception. The cause of a DSI exception can be determined by the bit settings of DSISR[9]. DAR contains the address of the breakpoint match condition.
—	One new register is implemented for supporting system level memory map	System memory base address register (MBAR) can be accessed with mtspr or mfspir using SPR311 in supervisor mode. It can store the present memory base address for the system memory map.

Table 1-6. Differences Between G2 and G2_LE Cores (continued)

G2 Core	G2_LE Core	Impact
—	One new register is implemented for identifying specific version and revision level of the system-on-a-chip (SOC)	The system version register (SVR) can be accessed with mf spr using SPR286. This register is programmed externally by the chip-integrator.
The G2 core has four pairs of data and four pairs of instruction BAT registers	The G2_LE has eight pairs of data and eight pairs of instruction BAT registers	IBAT4–IBAT7 are the four additional pairs of instruction BATs and DBAT4–DBAT7 are the four additional data BATs in G2_LE only. HID2[HBE] is added to the G2_LE for enabling or disabling the four additional pairs of BAT registers. These BATs are accessible by the mf spr and mts pr instructions regardless of the setting of HID2[HBE].
HID0–HID2 are the three unique hardware implementation registers for the G2 core ²	New bits are defined in HID2 for enabling the high BATs and true little-endian mode	HID0 and HID1 provide the means for enabling core checkstops and features and allows software to read the configuration of PLL configuration signals. HID2 enables cache way-locking; it also enables the true little-endian mode and the new additional BAT registers, for the G2_LE core.
—	The LSSD test control and the scan chain connections are rearranged in the G2_LE	New test integration requirements
—	G2_LE has seven additional signals for address matching, combinational matching, and breakpoints.	The G2_LE core implements the following additional features: <ul style="list-style-type: none"> • To support true little-endian mode <code>core_tle</code> is implemented • To support critical interrupt function <code>core_cint</code> is implemented • To support breakpoint state output, <code>core_iabr</code>, <code>core_iabr2</code>, <code>core_dabr</code>, and <code>core_dabr2</code> are implemented • To support additional debug features <code>core_tdo_oe</code> is added

¹ The MPC603e processor version number is 6 for PID6-603e and 7 for the PID7t-603e. The revision level starts at 0x0100 and changes for each revision of the MPC603e.

² HID0–HID1 are the two unique hardware implementation registers for the MPC603e.



Chapter 2

Register Model

This chapter describes the PowerPC register model and specific implementation on the G2 and G2_LE core.

2.1 Register Set

This section describes the register organization in the G2 core as defined by the three levels of the PowerPC architecture—user instruction set architecture (UISA), virtual environment architecture (VEA), and operating environment architecture (OEA), as well as the core implementation-specific registers. Full descriptions of the basic register set defined by the PowerPC architecture are provided in Chapter 2, “Register Set,” in the *Programming Environments Manual*.

The PowerPC architecture defines register-to-register operations for all computational instructions. Source data for these instructions is accessed from the on-chip registers or is provided as an immediate value embedded in the opcode. The three-register instruction format allows specification of a target register distinct from the two source registers, thus preserving the original data for use by other instructions and reducing the number of instructions required for certain operations. Data is transferred between memory and registers with explicit load and store instructions only.

Note that there may be registers common to other processors of this family that are not implemented in the G2 core. When the core detects special-purpose register (SPR) encodings other than those defined in this document, it either takes an exception or it treats the instruction as a no-op. (Note that exceptions are referred to as interrupts in the architecture specification.) Conversely, some SPRs in the G2 core may not be implemented in other processors or may not be implemented in the same way.

2.1.1 PowerPC Register Set

The UISA registers, shown in Figure 2-1, can be accessed by either user- or supervisor-level instructions (the architecture specification refers to user- and supervisor-level as problem state and privileged state, respectively). The general-purpose registers (GPRs) and floating-point registers (FPRs) are accessed through instruction operands. Access to registers can be explicit (that is, through the use of specific instructions

for that purpose, such as the **mtspr** and **mfspir** instructions) or implicit as part of the execution (or side effect) of an instruction. Some registers are accessed both explicitly and implicitly.

Figure 2-1 describes both the registers in the G2 core and the additional registers of the G2_LE core. All G2 core registers are present in the G2_LE core. Also note that the implementation-specific registers for the G2 and G2_LE cores are shown in Figure 2-1.

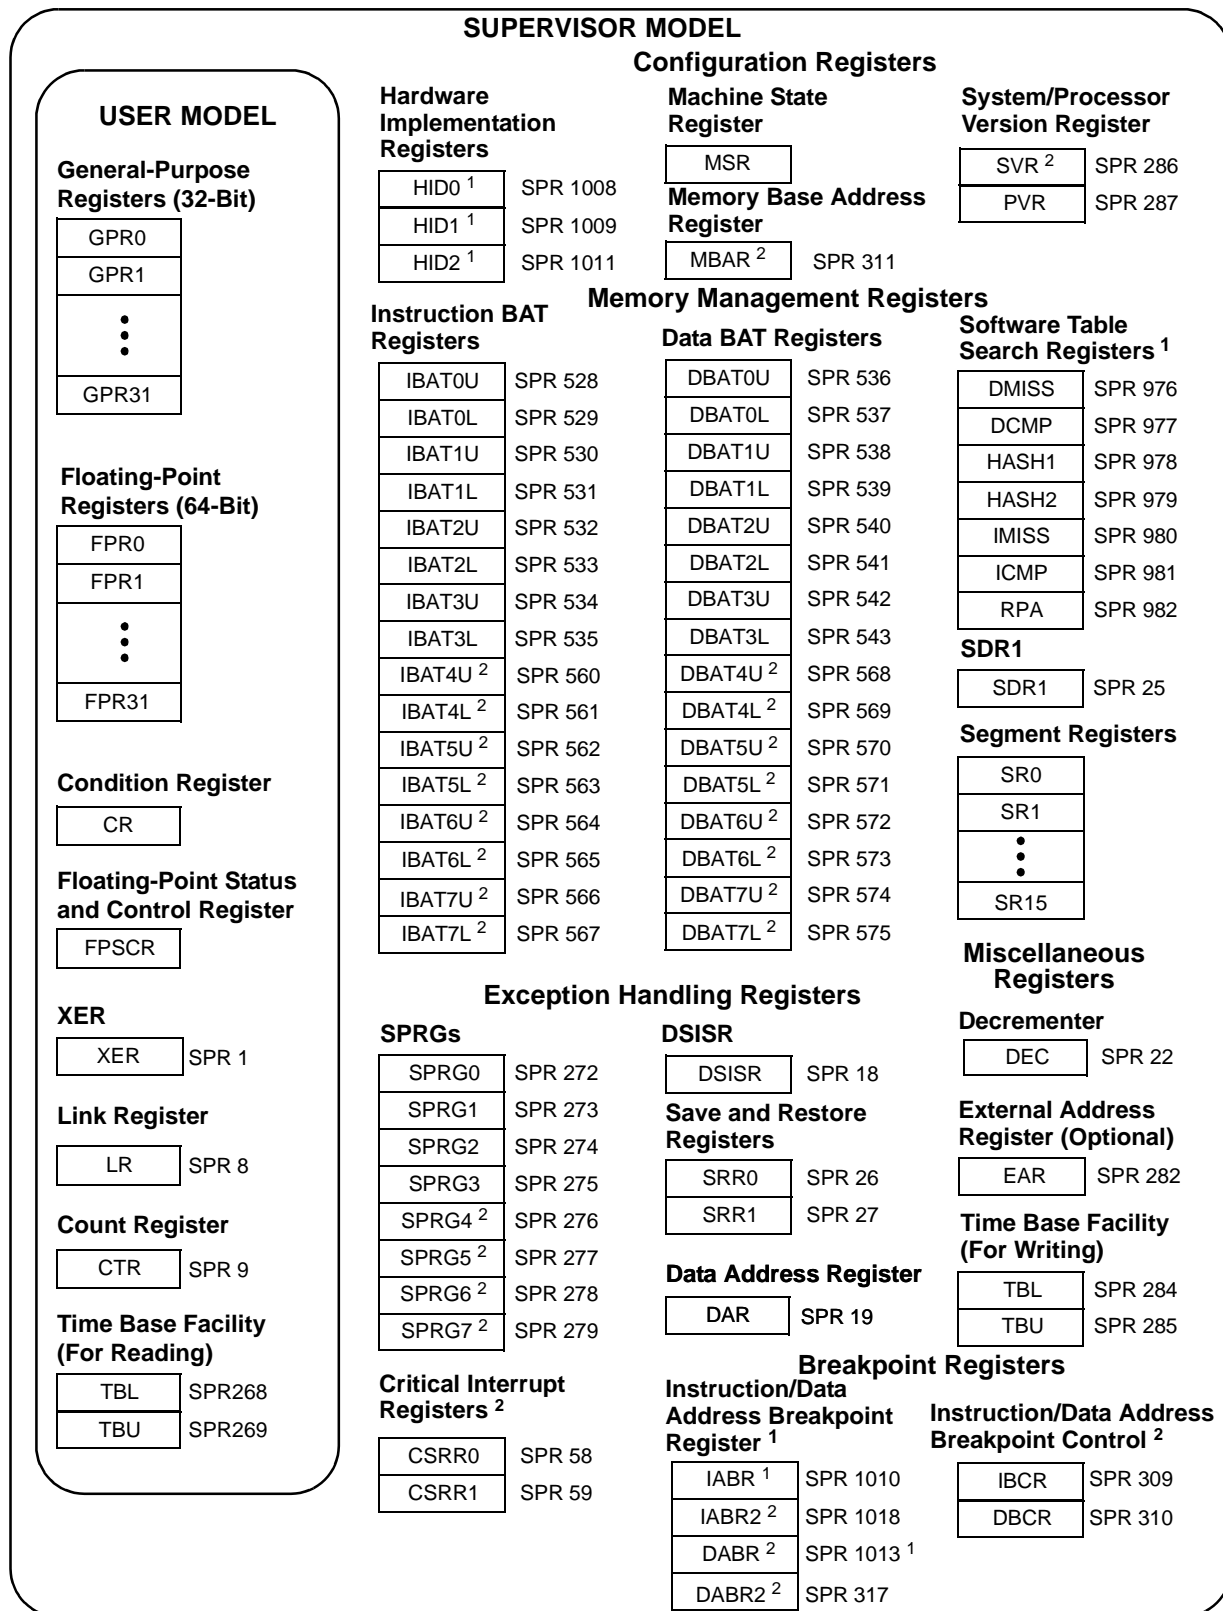
The number to the right of the register name indicates the number that is used in the syntax of the instruction operands to access the register (for example, the number used to access the XER is SPR1).

For more information on the PowerPC register set, refer to Chapter 2, “Register Set,” in the *Programming Environments Manual*.

The G2 core user-level registers are described as follows:

- User-level registers (UISA)—The user-level registers can be accessed by all software with either user or supervisor privileges. The user-level register set includes the following:
 - General-purpose registers (GPRs). The GPR file consists of thirty-two 32-bit GPRs designated as GPR0–GPR31. This register file serves as the data source or destination for all integer instructions and provides data for generating addresses.
 - Floating-point registers (FPRs). The FPR file consists of thirty-two 64-bit FPRs designated as FPR0–FPR31, which serves as the data source or destination for all floating-point instructions. These registers can contain data objects of either single- or double-precision floating-point format.

Before the **stfd** instruction is used to store the contents of an FPR to memory, the FPR must have been initialized after reset (explicitly loaded with any value) by using a floating-point load instruction.
 - Condition register (CR). The CR consists of eight 4-bit fields, CR0–CR7, that reflect the results of certain arithmetic operations and provides a mechanism for testing and branching.
 - Floating-point status and control register (FPSCR). The FPSCR contains all floating-point exception signal bits, exception summary bits, exception enable bits, and rounding control bits needed for compliance with the IEEE 754 standard.



The remaining user-level registers are SPRs. Note that the PowerPC architecture provides a separate mechanism for accessing SPRs (the **mtspr** and **mfspir** instructions). These instructions are commonly used to explicitly access certain registers, while other SPRs may be accessed as the side effect of executing other instructions.

- XER register (XER). The 32-bit XER indicates overflow and carries for integer operations. It is set implicitly by many instructions.
- Link register (LR). The 32-bit LR provides the branch target address for the Branch Conditional to Link Register (**bclr**) instruction and can optionally be used to hold the logical address (referred to as the effective address in the architecture specification) of the instruction that follows a branch and link instruction, typically used for linking to subroutines.
- Count register (CTR). The 32-bit CTR can be used to hold a loop count that can be decremented during execution of appropriately coded branch instructions. It can also provide the branch target address for the Branch Conditional to Count Register (**bcctr**) instruction.
- User-level registers (VEA)—The VEA introduces the time base facility (TB) for reading. The TB is a 64-bit register pair whose contents are incremented once every four bus clock cycles. The TB consists of two 32-bit registers—time base upper (TBU) and time base lower (TBL). Note that the time base registers are read-only in user state.

The core supervisor-level registers are described as follows:

- Supervisor-level registers (OEA)—The OEA defines the registers an operating system uses for memory management, configuration, and exception handling. The PowerPC architecture defines the following supervisor-level registers:
 - Configuration registers
 - Processor version register (PVR). This read-only register identifies the version (model) and revision level of this processor core. The contents of the PVR can be copied to a GPR by the **mfspir** instruction. Read access to the PVR is supervisor-level only; write access is not provided. The PVR consists of the fields as described in Table 2-1.

Table 2-1. PVR Field Descriptions

Bits	Name	Description
0–3	CID	Company or manufacturer ID number. For Freescale and Freescale licensees, bit 0 is set to one. Freescale's code is 0b1000.
4–5	—	Reserved
6–9	PT	Processor ID type. Optional field to identify different versions of the same processor [PID]; must read as zero if unused.

Table 2-1. PVR Field Descriptions (continued)

Bits	Name	Description
10–15	PID ¹	Processor identification. This field is used to indicate different implementations of the PowerPC architecture.
16–19	PROC	Process revision. This field identifies the relative process changes and revisions.
20–23	MFG	Manufacturing revision. This optional field identifies relative manufacturing revisions and changes. This was formerly the major processor design revision indicator (in the MPC603e).
24–27	MJREV	Major processor design revision indicator
28–31	MNREV	Minor processor design revision indicator

¹ The PID values are assigned by the PowerPC architecture group.

Note that the PowerPC architecture defines this register in more general terms than defined in Table 2-1. Architecturally, the PVR consists of two 16-bit fields as described in Table 2-2.

Table 2-2. Architectural PVR Field Descriptions

Bits	Name	Description
0–15	Version	A 16-bit number that uniquely identifies a particular processor version. This number can be used to determine the version of a processor; it may not distinguish between different end product models if more than one model uses the same processor.
16–31	Revision	A 16-bit number that distinguishes between various releases of a particular version (that is, an engineering change level). The value of the revision portion of the PVR is implementation-specific. The processor revision level is changed for each revision of the device.

Implementation Note—The G2 core version number is 0x8081 and the revision level starts at 0x1010 and changes for each revision of the core. The G2_LE core version number is 0x8082 and the revision level starts at 0x2010 and changes for each revision of the core. Table 2-3 describes some of the PVR values for G2-related devices.

Table 2-3. Assigned PVR Values

Device Name	Version No.	Revision No.
MPC603r (PID7)	0x0007	0x1201
G2 core—original	0x0081	0x0011
G2 core (G2H4)	0x8081	0x1010
G2 core (general-purpose)	0x8082	0x1010
G2 core (licensee-specific)	0x9081	0x0010
G2_LE core (licensee-specific)	0x8082	0x0010
G2_LE core (general-purpose)	0x8082	0x2010
G2_LE core (licensee-specific)	0xA082	0x2010
MPC603e (PID6)	0x0006	0x0101

Table 2-3. Assigned PVR Values (continued)

Device Name	Version No.	Revision No.
MPC603e (PID7v)	0x0007	0x0100, 0x0201
Space for future versions		

- Machine state register (MSR). The MSR defines the state of the processor. The MSR can be modified by the Move to Machine State Register (**mtmsr**), System Call (**sc**), and Return from Exception (**rfi**) and Return from Critical Exception (**rfci**) instructions. It can be read by the Move from Machine State Register (**mfmsr**) instruction.

Implementation Note—The G2 core defines MSR[13] as the power management enable (POW) bit and MSR[14] as the temporary GPR remapping (TGPR) bit. The G2_LE allocates MSR[24] for enabling the critical interrupt and **rfci**, the return from critical interrupt instruction. MSR[31] is used in conjunction with HID2[LET] to indicate the endian mode of operation of the G2_LE core. These bits are described in Table 2-4.

Table 2-4. MSR Bit Settings

Bit	Name	Description
0	—	Reserved. Full function.
1–4	—	Reserved. Partial function.
5–9	—	Reserved. Full function.
10–12	—	Reserved. Partial function.
13	POW	Power management enable (implementation-specific) 0 Disables programmable power modes (normal operation mode) 1 Enables programmable power modes (nap, doze, or sleep mode). This bit controls the programmable power modes only; it has no effect on dynamic power management (DPM). MSR[POW] may be altered with an mtmsr instruction only. Also, when altering the POW bit, software may alter only this bit in the MSR and no others. The mtmsr instruction must be followed by a context-synchronizing instruction. See Chapter 10, “Power Management,” for more information.
14	TGPR	Temporary GPR remapping (implementation-specific) 0 Normal operation 1 TGPR mode. GPR0–GPR3 are remapped to TGPR0–TGPR3 for use by TLB miss routines. The contents of GPR0–GPR3 remain unchanged while MSR[TGPR] = 1. Attempts to use GPR4–GPR31 with MSR[TGPR] = 1 yield undefined results. Temporarily replaces TGPR0–TGPR3 with GPR0–GPR3 for use by TLB miss routines. The TGPR bit is set when either an instruction TLB miss, data read miss, or data write miss exception is taken. The TGPR bit is cleared by an rfi instruction.

Table 2-4. MSR Bit Settings (continued)

Bit	Name	Description
15	ILE	Exception little-endian mode. When an exception occurs, this bit is copied into MSR[LE] to select the endian mode for the context established by the exception.
16	EE	External interrupt enable 0 The processor ignores external interrupts, system management interrupts, and decrementer interrupts. 1 The processor is enabled to take an external interrupt, system management interrupt, or decrementer interrupt.
17	PR	Privilege level 0 The processor can execute both user- and supervisor-level instructions 1 The processor can only execute user-level instructions
18	FP	Floating-point available 0 The processor prevents dispatch of floating-point instructions, including floating-point loads, stores, and moves. 1 The processor can execute floating-point instructions, and can take floating-point enabled exception type program exceptions.
19	ME	Machine check enable 0 Machine check exceptions are disabled 1 Machine check exceptions are enabled
20	FE0	Floating-point exception mode 0 (see Table 5-8)
21	SE	Single-step trace enable 0 The processor executes instructions normally 1 The processor generates a trace exception upon the successful completion of the next instruction
22	BE	Branch trace enable 0 The processor executes branch instructions normally 1 The processor generates a trace exception upon the successful completion of a branch instruction
23	FE1	Floating-point exception mode 1 (see Table 5-8)
24	CE	Critical interrupt exception enable (G2_LE core-only) 0 Critical interrupts disabled 1 Critical interrupts enabled; critical interrupt exception and rfci instruction enabled The critical interrupt is an asynchronous implementation-specific exception. The critical interrupt exception vector offset is 0x00A00. The rfci instruction is implemented to return from these exception handlers. Also, CSRR0 and CSRR1 are used to save and restore the processor state for critical interrupts.
25	IP	Exception prefix. The setting of this bit specifies whether an exception vector offset is prepended with Fs or 0s. In the following description, <i>nnnn</i> is the offset of the exception. See Table 5-2. 0 Exceptions are vectored to the physical address 0x000 <i>n_nnnn</i> 1 Exceptions are vectored to the physical address 0xFFFF <i>n_nnnn</i>
26	IR	Instruction address translation 0 Instruction address translation is disabled 1 Instruction address translation is enabled See Chapter 6, “Memory Management”

Table 2-4. MSR Bit Settings (continued)

Bit	Name	Description
27	DR	Data address translation 0 Data address translation is disabled 1 Data address translation is enabled See Chapter 6, “Memory Management”
28–29	—	Reserved. Full function.
30	RI	Recoverable exception (for system reset and machine check exceptions) 0 Exception is not recoverable 1 Exception is recoverable
31	LE	Little-endian mode enable 0 The processor runs in big-endian mode 1 The processor runs in little-endian mode. For the G2_LE core, see Section 1.1.2.1, “True Little-Endian Mode,” for a definition of whether the core is operating in true little-endian mode or modified little-endian mode.

— Memory management registers

- Block-address translation (BAT) registers. The G2 core also supports eight block-address translation registers (BATs) through the use of two independent instruction and data block address translation (IBAT and DBAT) arrays, each containing four pairs of BATs. However, the G2_LE core supports block address translation arrays of eight pairs of data BATs and eight pairs of instruction BATs, which are implementation-specific. Effective addresses are compared simultaneously with all four (or eight, for G2_LE) entries in the BAT array during block translation. Figure 2-1 lists SPR numbers for the BAT registers.
- SDR1. The SDR1 register specifies the page table base address used in virtual-to-physical address translation. (Note that physical address is referred to as real address in the architecture specification.)
- Segment registers (SRs). The OEA defines sixteen 32-bit segment registers (SR0–SR15). The fields in the segment register are interpreted differently depending on the value of bit 0.

— Exception handling registers

- Data address register (DAR). After a data access or an alignment exception, the DAR is set to the effective address generated by the faulting instruction.
- The SPRG0–SPRG3 registers are provided for operating system use, which reduce the latency that may be incurred because of saving registers to memory while in a handler and also assist in searching the page tables in software. If software table searching is not enabled, then these registers may be used for any supervisor purpose. Note that the G2_LE core implements four additional SPRGs (SPRG4–SPRG7), which are not defined by the PowerPC architecture. The format of these registers is the same as that of

SPRG0–SPRG3 defined in Section 2.1.2.11, “SPRG4–SPRG7 (G2_LE Only).”

- DSISR. The DSISR defines the cause of data access and alignment exceptions.
- Machine status save/restore register [0–1] (SRR0, SRR1). The SRR0 and SRR1 are used to save machine status on exceptions and to restore machine status when an **rfi** instruction is executed.

Implementation Note—The G2 core implements the KEY bit (bit 12) in the SRR1 register to simplify the table search software. For more information refer to Chapter 6, “Memory Management.”

Note that to support critical interrupts, two new registers, CSRR0 and CSRR1, are implemented on the G2_LE core, which are not defined by the PowerPC architecture. These registers have same bit assignments as SRR0 and SRR1, and are described in Section 2.1.2, “Implementation-Specific Registers.”

— Miscellaneous registers

- The time base facility (TB) for writing. The TB is a 64-bit register pair that can be used to provide time-of-day or interval timing. It consists of two 32-bit registers—time base upper (TBU) and time base lower (TBL). The TB is incremented once every four clock cycles on the core.
- Decrementer (DEC). The DEC register is a 32-bit decrementing counter that provides a mechanism for causing a decrementer exception after a programmable delay. The DEC is decremented once every four bus clock cycles.
- External access register (EAR). The EAR is a 32-bit register used in conjunction with the **eciwx** and **ecowx** instructions. Although the PowerPC architecture specifies that EAR26–EAR31 are used to select a device, the G2 core implements only bits 28–31. Note that EAR and the **eciwx** and **ecowx** instructions are optional in the PowerPC architecture and may not be supported in all processors that implement the OEA.

2.1.2 Implementation-Specific Registers

The G2 core defines the DMISS, IMISS, DCMP, ICMP, HASH1, HASH2, and RPA registers for software table search operations. These registers should be accessed only when address translation is disabled (MSR[IR] and MSR[DR] are both zero). For a complete discussion, refer to Section 6.5.2, “Implementation-Specific Table Search Operation.” Also, HID0, HID1, and IABR SPRs are defined and described in this section. These registers can be accessed by supervisor-level instructions only using the SPR numbers shown in Figure 2-1.

Note that the G2_LE core defines the following:

- Two new critical interrupt registers (CSRR0, CSRR1), which are implementation-specific. The CSRR0 and CSRR1 registers support the critical interrupt function, which have the same bit assignments as SRR0 and SRR1, respectively. The effective address for resuming program execution is saved into CSRR0 and the content of the MSR is saved into CSRR1. An additional **rftci** instruction is implemented for supporting the return from a critical interrupt, selecting the CSRR0 and CSRR1 registers.
- Four additional exception handling SPRG registers, which are provided for operating system use.
- A new system version register (SVR). See Section 2.1.2.12, “System Version Register (SVR)—G2_LE Only,” for bit definitions.
- System memory base address (MBAR) is a new implementation-specific register for the G2_LE core. It supports a system-level memory map. See Section 2.1.2.13, “System Memory Base Address (MBAR)—G2_LE Only,” for more information.
- Eight additional BATs (IBAT4–IBAT7 and DBAT4–DBAT7), providing better performance in protecting accesses on a segment, block, or page basis along with memory accesses and I/O accesses. See Figure 2-1 for a list of the SPR numbers for the BAT arrays.
- One additional address breakpoint register (IABR2), one new instruction address breakpoint control register (IBCR), two new data breakpoint registers (DABR, DABR2), and one new data address breakpoint control register (DBCR) are implemented in the G2_LE processor core. All these registers are implementation-specific and they are described in the Section 2.1.2.14, “Instruction Address Breakpoint Registers (IABR and IABR2),” and Section 2.1.2.15, “Data Address Breakpoint Register (DABR and DABR2)—G2_LE Only.”

2.1.2.1 Hardware Implementation Register 0 (HID0)

The HID0 register, shown in Figure 2-2, defines enable bits for various G2 core-specific features.

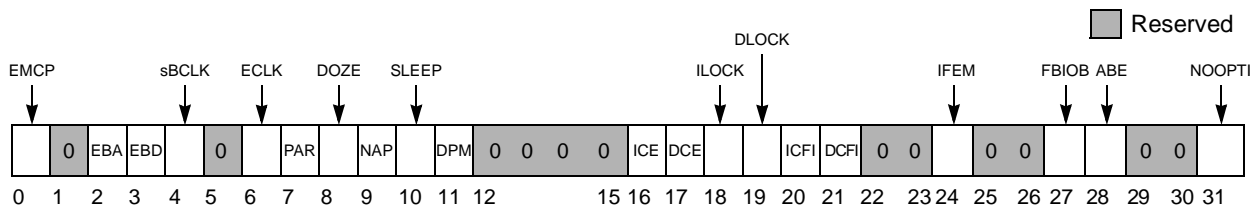


Figure 2-2. Hardware Implementation Register 0 (HID0)

Table 2-5 shows the bit definitions for HID0.

Table 2-5. HID0 Bit Functions

Bits	Name	Function
0	EMCP	Enable <code>core_mcp</code> . The primary purpose of this bit is to mask out further machine check exceptions caused by assertion of <code>core_mcp</code> , similar to how MSR[EE] can mask external interrupts. 0 Masks <code>core_mcp</code> . Asserting <code>core_mcp</code> does not generate a machine check exception or a checkstop. 1 Asserting <code>core_mcp</code> causes checkstop if MSR[ME] = 0 or a machine check exception if ME = 1
1	—	Reserved
2	EBA	Enable <code>core_ap_in[0:3]</code> and <code>core_ape</code> for address parity checking. EBA and EBD allow the processor to operate with memory subsystems that do not generate parity. 0 Disables address parity checking during a snoop operation 1 Allows an address parity error during snoop operations to cause a checkstop if MSR[ME] = 0 or a machine check exception if MSR[ME] = 1
3	EBD	Enable <code>core_dpe</code> for data parity checking. EBA and EBD allow the processor to operate with memory subsystems that do not generate parity. 0 Disables data parity checking 1 Allows a data parity error during reads to cause a checkstop if MSR[ME] = 0 or a machine check exception if MSR[ME] = 1
4	SBCLK	<code>core_clk_out</code> output enable. Used in conjunction with HID0[ECLK] and <code>core_hreset</code> to configure <code>core_clk_out</code> . See Table 2-6.
5	—	Reserved
6	ECLK	<code>core_clk_out</code> output enable. Used in conjunction with HID0[SBCLK] and the <code>core_hreset</code> signal to configure <code>core_clk_out</code> . See Table 2-6.
7	PAR	Disable precharge of <code>core_artry_out</code> 0 Precharge of <code>core_artry_out</code> enabled 1 Alters bus protocol slightly by preventing the processor from driving <code>core_artry_out</code> to high (negated) state. If this is done, the integrated device must restore the signals to the high state.
8	DOZE ¹	Doze mode enable. Operates in conjunction with MSR[POW]. 0 Doze mode disabled 1 Doze mode enabled. Doze mode is invoked by setting MSR[POW] while this bit is set. In doze mode, the PLL, time base, and snooping remain active.
9	NAP ¹	Nap mode enable. Operates in conjunction with MSR[POW]. 0 Nap mode disabled 1 Nap mode enabled. Doze mode is invoked by setting MSR[POW] while this bit is set. In nap mode, the PLL and time base remain active.
10	SLEEP ¹	Sleep mode enable. Operates in conjunction with MSR[POW]. 0 Sleep mode disabled 1 Sleep mode enabled. Sleep mode is invoked by setting MSR[POW] while this bit is set. <code>core_qreq</code> is asserted to indicate that the processor is ready to enter sleep mode. If the system logic determines that the processor may enter sleep mode, the quiesce acknowledge signal, <code>core_qack</code> , is asserted back to the processor. Once <code>core_qack</code> assertion is detected, the processor enters sleep mode after several processor clocks. At this point, the system logic may turn off the PLL by first configuring <code>core_pll_cfg[0:4]</code> to PLL bypass mode, then disabling <code>core_sysclk</code> .

Table 2-5. HID0 Bit Functions (continued)

Bits	Name	Function
11	DPM ¹	Dynamic power management enable 0 Dynamic power management is disabled 1 Functional units enter a low-power mode automatically if the unit is idle. This does not affect operational performance and is transparent to software or any external hardware.
12–15	—	Reserved, should be cleared.
16	ICE ²	Instruction cache enable 0 The instruction cache is neither accessed nor updated. All pages are accessed as if they were marked cache-inhibited (WIM = x1x). Potential cache accesses from the bus (snoop and cache operations) are ignored. In the disabled state for the L1 caches, the cache tag state bits are ignored and all accesses are propagated to the 60x bus as single-beat transactions. For those transactions, however, <code>core_ci</code> reflects the state of the I bit in the MMU for that page regardless of cache disabled status. ICE is zero at power-up. 1 The instruction cache is enabled
17	DCE	Data cache enable 0 The data cache is neither accessed nor updated. All pages are accessed as if they were marked cache-inhibited (WIM = x1x). Potential cache accesses from the 60x bus (snoop and cache operations) are ignored. In the disabled state for the L1 caches, the cache tag state bits are ignored and all accesses are propagated to the 60x bus as single-beat transactions. For those transactions, however, <code>core_ci</code> reflects the state of the I bit in the MMU for that page regardless of cache disabled status. DCE is zero at power-up. 1 The data cache is enabled
18	ILOCK	Instruction cache lock 0 Normal operation 1 Instruction cache is locked. A locked cache supplies data normally on a hit, but the access is treated as a cache-inhibited transaction on a miss. On a miss, the transaction to the 60x bus is single-beat; however, <code>core_ci</code> still reflects the state of the I bit in the MMU for that page independent of cache locked or disabled status. To prevent locking during a cache access, an isync instruction must precede the setting of ILOCK.
19	DLOCK	Data cache lock 0 Normal operation 1 Data cache is locked. A locked cache supplies data normally on a hit, but is treated as a cache-inhibited transaction on a miss. On a miss, the transaction to the 60x bus is single-beat; however, <code>core_ci</code> still reflects the state of the I bit in the MMU for that page independent of cache locked or disabled status. A snoop hit to a locked L1 data cache performs as if the cache were not locked. A cache block invalidated by a snoop remains invalid until the cache is unlocked. To prevent locking during a cache access, a sync instruction must precede the setting of DLOCK.
20	ICFI	Instruction cache flash invalidate 0 The instruction cache is not invalidated. The bit is cleared when the invalidation operation begins (usually the next cycle after the write operation to the register). The instruction cache must be enabled for the invalidation to occur. 1 An invalidate operation is issued that marks the state of each instruction cache block as invalid without writing back modified cache blocks to memory. Cache access is blocked during this time. Bus accesses to the cache are signaled as a miss during invalidate-all operations. Setting ICFI clears all the valid bits of the blocks and the PLRU bits to point to way L0 of each set. For the G2 core, the proper use of the ICFI and DCFI bits is to set and clear them with two consecutive mtspr operations.

Table 2-5. HID0 Bit Functions (continued)

Bits	Name	Function
21	DCFI	Data cache flash invalidate 0 The data cache is not invalidated. The bit is cleared when the invalidation operation begins (usually the next cycle after the write operation to the register). The data cache must be enabled for the invalidation to occur. 1 An invalidate operation is issued that marks the state of each data cache block as invalid without writing back modified cache blocks to memory. Cache access is blocked during this time. Bus accesses to the cache are signaled as a miss during invalidate-all operations. Setting DCFI clears all the valid bits of the blocks and the PLRU bits to point to way L0 of each set. For the G2 core, the proper use of the ICFI and DCFI bits is to set and clear them with two consecutive mtspr operations.
22–23	—	Reserved, should be cleared.
24	IFEM	Enable M bit on 60x bus for instruction fetches 0 M bit not reflected on bus for instruction fetches. Instruction fetches are treated as nonglobal on the bus. 1 Instruction fetches reflect the M bit from the WIM settings
25–26	—	Reserved, should be cleared.
27	FBIOB	Force branch indirect on bus 0 Register indirect branch targets are fetched normally 1 Forces register indirect branch targets to be fetched externally
28	ABE	Address broadcast enable. Controls whether certain address-only operations (such as cache operations) are broadcast on the 60x bus. 0 Address-only operations affect only local caches and are not broadcast 1 Address-only operations are broadcast on the 60x bus Affected instructions are dcbi , dcbf , and dcbst . Note that these cache control instruction broadcasts are not snooped by the G2 core. Refer to Section 4.3.3, “Data Cache Control,” for more information.
29–30	—	Reserved, should be cleared.
31	NOOPTI	No-op the data cache touch instructions 0 The dcbt and dcbtst instructions are enabled 1 The dcbt and dcbtst instructions are no-oped globally

¹ See Chapter 10, “Power Management.”

² See Chapter 4, “Instruction and Data Cache Operation.”

Table 2-6 shows how HID0[BCLK], HID0[ECLK], and core_hreset are used to configure core_clk_out. See Section 8.3.15.2, “Test Clock Output (core_clk_out),” for more information.

Table 2-6. HID0[SBCLK] and HID0[ECLK] core_clk_out Configuration

core_hreset	HID0[ECLK]	HID0[SBCLK]	core_clk_out
Asserted	x	x	Core
Negated	0	0	Core
Negated	0	1	Core clock frequency/2
Negated	1	0	Core
Negated	1	1	Bus

HID0 can be accessed with **mtspr** and **mfspir** using SPR1008.

2.1.2.2 Hardware Implementation Register 1 (HID1)

The HID1 register, shown in Figure 2-3, defines enable bits for various G2 core-specific features.

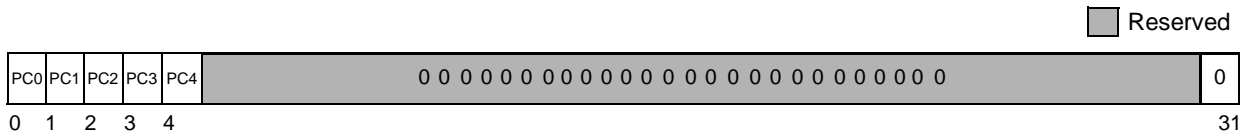


Figure 2-3. Hardware Implementation Register 1 (HID1)

Table 2-7 shows the bit definitions for HID1.

Table 2-7. HID1 Bit Settings

Bits	Name	Description
0	PC0	PLL configuration bit 0 (read-only)
1	PC1	PLL configuration bit 1 (read-only)
2	PC2	PLL configuration bit 2 (read-only)
3	PC3	PLL configuration bit 3 (read-only)
4	PC4	PLL configuration bit 4 (read-only)
5–30	—	Reserved, should be cleared
31	0	Tied to zero

Note: The clock configuration bits reflect the state of the core_pll_cfg[0:4] signals.

HID1 can be accessed with **mfspir** using SPR1009.

2.1.2.3 Hardware Implementation Register 2 (HID2)

The G2 core implements an additional hardware implementation-dependent HID2 register, shown in Figure 2-4, which enables cache way-locking; the G2_LE core also enables true

little-endian mode and the new additional BAT registers. It is a supervisor-only, read/write, implementation-specific special purpose register (SPR) which is accessed as SPR1011 (decimal). The HID2 bits are shown in Table 2-8.

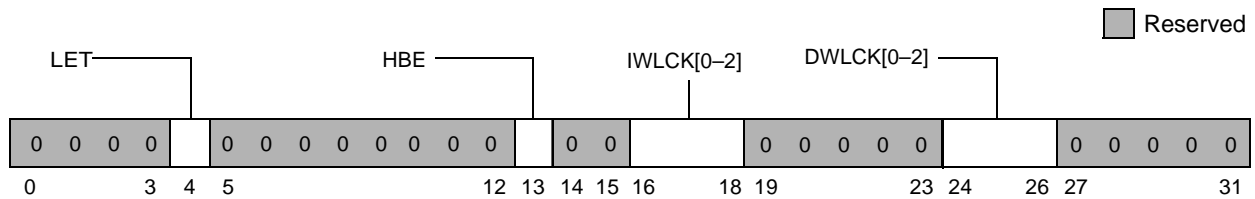


Figure 2-4. Hardware Implementation-Dependent Register 2 (HID2)

Table 2-8 describes the HID2 fields.

Table 2-8. HID2 Bit Descriptions

Bits	Name	Description
0–3	—	Reserved
4	LET	True little-endian. This bit enables true little-endian mode operation for instruction and data accesses. This bit is set to reflect the state of the core_tle signal at the negation of core_hreset. This bit is used in conjunction with MSR[LE] to determine the endian mode of operation as described in Table 1-1. 0 Modified (PowerPC) little-endian mode 1 True little-endian mode Changing the value of this bit during normal operation is discouraged
5–12	—	Reserved
13	HBE	High BAT enable. Regardless of the setting of HID2[HBE], these BATs are accessible by mf spr and mts pr . 0 IBAT[4–7] and DBAT[4–7] are disabled 1 IBAT[4–7] and DBAT[4–7] are enabled
14	—	Reserved
15	—	Reserved
16–18	IWLCK[0–2]	Instruction cache way-lock. Useful for locking blocks of instructions into the instruction cache for time-critical applications that require deterministic behavior. See Chapter 4, “Instruction and Data Cache Operation.” 000 = no ways locked 001 = way 0 locked 010 = way 0 through way 1 locked 011 = way 0 through way 2 locked 100 = way 0 through way 3 locked 101 = way 0 through way 4 locked 110 = way 0 through way 5 locked 111 = Reserved
19–23	—	Reserved

Table 2-8. HID2 Bit Descriptions (continued)

Bits	Name	Description
24–26	DWLCK[0–2]	Data cache way-lock. Useful for locking blocks of data into the data cache for time-critical applications where deterministic behavior is required. See Chapter 4, “Instruction and Data Cache Operation.” 000 = no ways locked 001 = way 0 locked 010 = way 0 through way 1 locked 011 = way 0 through way 2 locked 100 = way 0 through way 3 locked 101 = way 0 through way 4 locked 110 = way 0 through way 5 locked 111 = Reserved
27–31	—	Reserved

2.1.2.4 Data and Instruction TLB Miss Address Registers (DMISS and IMISS)

DMISS and IMISS, shown in Figure 2-5, are loaded automatically on a data or instruction TLB miss. DMISS and IMISS contain the effective address of the access that caused the TLB miss exception. The contents are used by the core when calculating the values of HASH1 and HASH2 and by the **tlbld** and **tlbli** instructions when loading a new TLB entry. Note that the G2 core always loads DMISS with a big-endian address, even when MSR[LE] is set. These registers are both read- and write-accessible. However, caution should be used when writing to these registers.

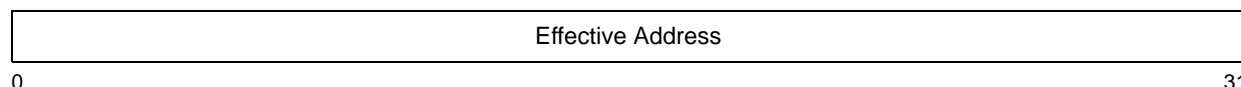


Figure 2-5. DMISS and IMISS Registers

2.1.2.5 Data and Instruction TLB Compare Registers (DCMP and ICMP)

DCMP and ICMP, shown in Figure 2-6, contain the first word in the required PTE. The contents are constructed automatically from the contents of the segment registers and the effective address (DMISS or IMISS) when a TLB miss exception occurs. Each PTE read from the tables during the table search process should be compared with this value to determine if the PTE is a match. Upon execution of a **tlbld** or **tlbli** instruction, the upper 25 bits of the DCMP or ICMP register and 11 bits of the effective address are loaded into the first word of the selected TLB entry. These registers are read and write to the software.

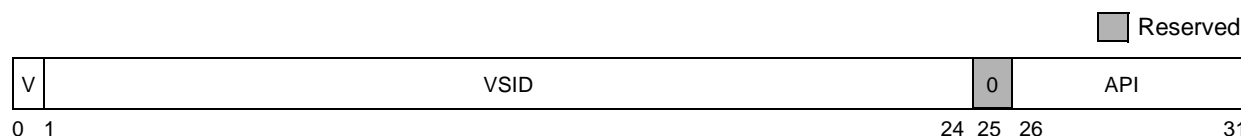


Figure 2-6. DCMP and ICMP Registers

Table 2-9 describes the bit settings for the DCOMP and ICMP registers.

Table 2-9. DCOMP and ICMP Bit Settings

Bit	Name	Description
0	V	Valid bit. Set by the processor on a TLB miss exception.
1–24	VSID	Virtual segment ID. Copied from VSID field of corresponding segment register.
25	—	Reserved, should be cleared.
26–31	API	Abbreviated page index. Copied from API of effective address.

2.1.2.6 Primary and Secondary Hash Address Registers (HASH1 and HASH2)

HASH1 and HASH2, shown in Figure 2-7, contain the physical addresses of the primary and secondary PTEGs for the access that caused the TLB miss exception. For convenience, the G2 core automatically constructs the full physical address by routing SDR1 bits 0–6 into HASH1 and HASH2 and clearing the lower 6 bits. These read-only registers are constructed from the DMISS or IMISS contents (the register choice is determined by which miss most recently occurred).

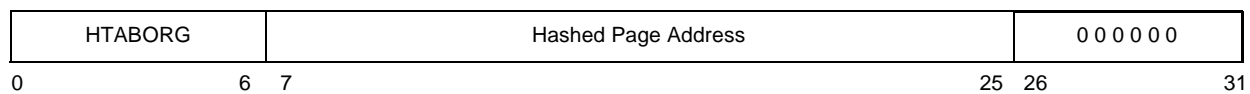


Figure 2-7. HASH1 and HASH2 Registers

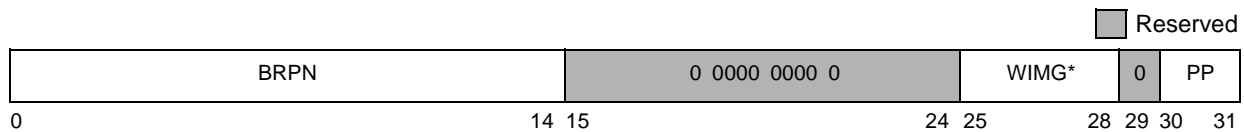
Table 2-10 describes the bit settings of the HASH1 and HASH2 registers.

Table 2-10. HASH1 and HASH2 Bit Settings

Bit	Name	Description
0–6	HTABORG	Copy of the upper 7 bits of the HTABORG field from SDR1
7–25	Hashed page address	Address bits 7–25 of the PTEG to be searched
26–31	—	Reserved

2.1.2.7 Required Physical Address Register (RPA)

During a page table search operation, the software must load the RPA, shown in Figure 2-8, with the second word of the correct PTE. When the **tlbld** or **tlbli** instruction is executed, the RPA and DMISS or IMISS register are merged and loaded into the selected TLB entry. The referenced (R) bit is ignored when the write occurs (no location exists in the TLB entry for this bit). The RPA register is read and write accessible to the software.



*W and G bits are not defined for IBAT registers. Attempting to write to these bits causes boundedly-undefined results.

Figure 2-10. Lower BAT Register

The BAT registers contain the effective-to-physical address mappings for blocks of memory. This mapping includes the effective address bits that are compared with the effective address of the access, the memory/cache access mode bits (WIMG), and the protection bits for the block. The size of the block and the starting address of the block are defined by the physical block number (BRPN) and block size mask (BL) fields.

The sixteen new BAT registers are enabled by HID2[HBE]. However, regardless of the setting of this bit, the BAT registers are accessible by the **mf spr** and **mt spr** instructions and are only accessible to supervisor-level programs. See Section 2.1.2.3, “Hardware Implementation Register 2 (HID2),” for more information on the HBE bit.

2.1.2.9 Critical Interrupt Save/Restore Register 0 (CSRR0)—G2_LE Only

CSRR0 is used to save machine status on critical interrupt exceptions and restore machine status when an **r fci** instruction is executed. The format of CSRR0 is shown in Figure 2-11.

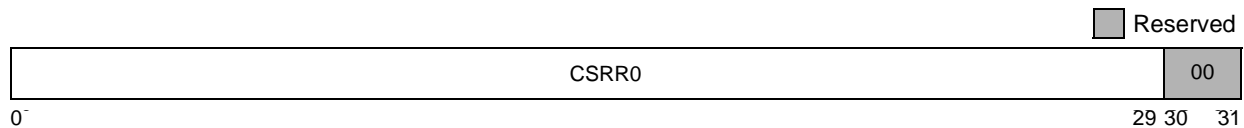


Figure 2-11. Critical Interrupt Save/Restore Register 0 (CSRR0)

For information on how specific exceptions affect CSRR0, refer to the descriptions of individual exceptions in Chapter 5, “Exceptions.”

2.1.2.10 Critical Interrupt Save/Restore Register 1 (CSRR1)—G2_LE Only

CSRR1 is used to save machine status on exceptions and to restore machine status when an **r fci** instruction is executed. Figure 2-12 shows the CSRR1 format.

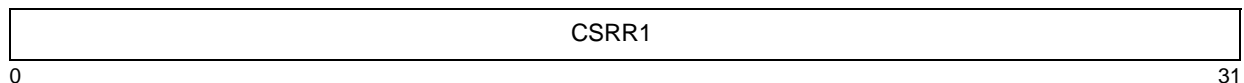


Figure 2-12. Critical Interrupt Save/Restore Register 1 (CSRR1)

For information on how specific exceptions affect CSRR1, refer to the individual exceptions in Chapter 5, “Exceptions.”

2.1.2.11 SPRG4–SPRG7 (G2_LE Only)

The G2_LE core provides four additional SPRG (SPRG4–SPRG7) registers for general operating system use, such as performing a fast state save or for supporting multiprocessor implementations. The formats of SPRG4–SPRG7 are shown in Figure 2-13.

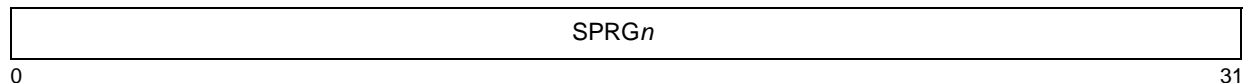


Figure 2-13. SPRG0–SPRG7 Registers

For information on conventional uses for SPRG4–SPRG7, refer to Section 5.2.1.3, “SPRG4–SPRG7 (G2_LE Only).”

2.1.2.12 System Version Register (SVR)—G2_LE Only

The system version register (SVR) is 32-bit (G2_LE specific), read-only register that identifies the specific version (model) and revision level of the system on a chip (SOC), including the processor core identification by the PVR. Supervisor mode write access is reserved for future use. The SVR can be accessed with **mfspr** using SPR286. The bits in SVR are defined in Table 2-12.

Note that all bits within this register must be programmed by the SOC and unused bits must be set to zero. Also, SVR4–SVR15 are control fields for this register.

Table 2-12. System Version Register (SVR) Bit Settings

Bits	Name	Description
0–3	CID	Company or manufacturer ID. These bits are required. Bit 0 must set to 1.
4–9	SOCOP ¹	SOC Integration options. This optional field identifies the SOC device specific options that are integrated within the SOC. The field reads 0 when it is not used.
10–15	SID ²	SOC ID. This required field is used to identify the SOC device.
16–19	PROC	Process revision field. This optional field is used to indicate different process revisions of the SOC.
20–23	MFG	Manufacturing revision. This optional field identifies uniquely different manufacturing revisions of the SOC.
24–27	MJREV	Major SOC design revision indicator. This is a required field.
28–31	MNREV	Minor SOC design revision indicator. This is a required field.

¹ The SID values are assigned by the PowerPC architecture.

² The SOC value is an optional field assigned by the SOC design integrator.

2.1.2.13 System Memory Base Address (MBAR)—G2_LE Only

The G2_LE core implements a new memory base address register (MBAR) to support the system level memory map. The MBAR can be accessed with **mtspr** or **mfspr** using

SPR311 in supervisor mode. The present memory base address for the system memory map is stored in this register. It is important to ensure that the present value of the base offset is current in the system memory.

2.1.2.14 Instruction Address Breakpoint Registers (IABR and IABR2)

The IABR, shown in Figure 2-14, controls the instruction address breakpoint exception. In the G2_LE core, an additional address breakpoint register (IABR2) is implemented. IABR[CEA] holds an effective address to which each instruction's address is compared. The exception is enabled by setting IABR[BE]. The exception is taken when there is an instruction address breakpoint match on the next instruction to complete. The instruction tagged with the match cannot complete before the breakpoint exception is taken. The address of the instruction which matches the breakpoint condition is stored in SRR0. The tagged instruction is completed and retired on return from the exception (**rfi** or **rfti**). The results are then committed to the destination registers and address.

Note that if the IABR/IABR2 register values are set to any exception vector, an unrecoverable processor state will occur.

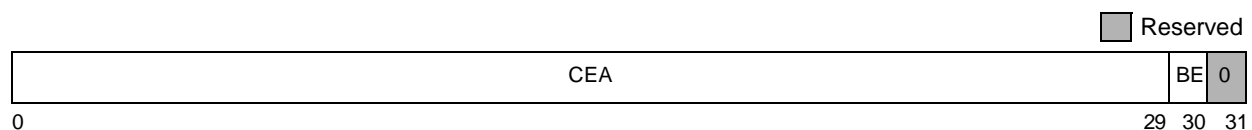


Figure 2-14. Instruction Address Breakpoint Registers (IABR and IABR2)

The bits in the IABR and IABR2 are defined in Table 2-13. For more information about the instruction breakpoint exception, see Section 5.5.16, “Instruction Address Breakpoint Exception (0x01300).”

Table 2-13. Instruction Address Breakpoint Register (IABR and IABR2) Bit Settings

Bits	Name	Description
0–29	CEA	Compare effective address. Word address to be compared.
30	BE	Breakpoint enable. IABR (or IABR2) enabled. Setting this bit enables the IABR exception.
31	—	Reserved

2.1.2.14.1 Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only

The IBCR, shown in Figure 2-15, is a supervisor-level register with SPR309 on the G2_LE core, which is accessible only by using an **mtspr** or **mfspir** instruction. The IBCR controls the compare and match type conditions for IABR and IABR2. Note that IABR and IABR2 must be enabled before the effects of IBCR are realized.

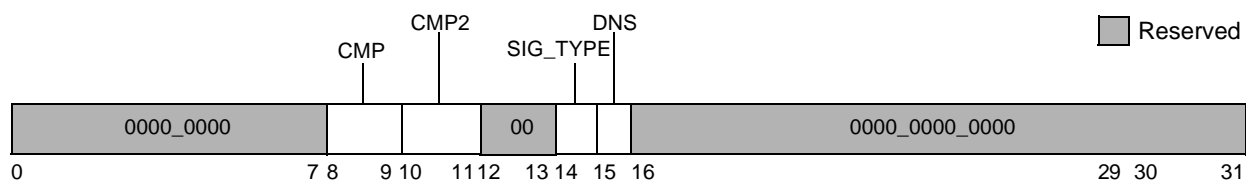


Figure 2-15. Instruction Address Breakpoint Control Register (IBCR)

Table 2-14 describes the IBCR fields.

Table 2-14. Instruction Address Breakpoint Control Registers (IBCR)

Bits	Name	Description
0–7	—	Reserved
8–9	CMP	IABR breakpoint compare type 00 Match if instruction's EA equals IABR[CEA] 01 Reserved 10 Match if instruction's EA is less than IABR[CEA] 11 Match if instruction's EA is greater than or equal to IABR[CEA]
10–11	CMP2	IABR2 breakpoint compare type 00 Match if instruction's EA equals IABR2[CEA] 01 Reserved 10 Match if instruction's EA less than IABR2[CEA] 11 Match if instruction's EA greater than or equal to IABR2[CEA]
12	—	Reserved
13	—	Reserved
14	SIG_TYPE	Combinational signal type 0 Instruction's EA matches IABR[CEA] OR instruction's EA matches IABR2[CEA] 1 Instruction's EA matches IABR[CEA] AND instruction's EA matches IABR2[CEA]
15	DNS	Do not signal. Disable $\overline{\text{core_iabr}}$ and $\overline{\text{core_iabr2}}$ output signals 0 Allow signal to toggle on a match 1 Do not toggle signal on match

2.1.2.15 Data Address Breakpoint Register (DABR and DABR2)—G2_LE Only

The optional data address breakpoint facility on the G2_LE core is controlled by optional SPRs, DABR and DABR2. The data address breakpoint facility provides a means to detect data accesses to a designated double-word address. The breakpoint address is compared to the effective address of all data accesses; it does not apply to instruction fetches.

DABR and DABR2, the two data address breakpoint registers shown in Figure 2-16, can both cause the data address breakpoint exception.

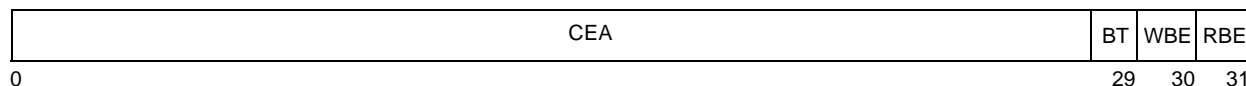


Figure 2-16. Data Address Breakpoint Registers (DABR and DABR2)

When an enabled data breakpoint condition matches with the address of a data access, a DSI exception occurs. When a DSI exception is taken to indicate a data breakpoint condition, DAR is set to the data address that causes the breakpoint and DSISR[9] is set. The address of the instruction associated with the breakpoint condition is stored in SRR0.

Note that if the DABR/DABR2 register values are set to match on any exception vector, an indeterminate or unrecoverable processor state may occur.

Table 2-15 describes the fields in DABR and DABR2.

Table 2-15. Data Address Breakpoint Registers (DABR and DABR2) Bit Settings

Bits	Name	Description
0–28	CEA	Data address breakpoint
29	BT	Breakpoint translation enable. Match if MSR[DR] = DABR[BT].
30	WBE	Data write enable. Matching on data writes enabled.
31	RBE	Data read enable. Matching on data reads enabled.

A data address breakpoint match is detected for a load or store instruction if the following conditions are met for any byte accessed:

- EA0–EA28 = DABR[CEA]
- MSR[DR] = DABR[BT]
- The instruction is a store and DABR[WBE] = 1 or the instruction is a load and DABR[RBE] = 1

Even if the above conditions are satisfied, it is undefined whether a match occurs in the following cases:

- A store conditional indexed instruction (**stwcx.**) in which the store is not performed
- A load or store string instruction (**lswx** or **stswx**) with a zero length
- A **dcbz**, **dcba**, **eciwx**, or **ecowx** instruction. For the purpose of determining whether a match occurs, **eciwx** is treated as a load and **dcbz**, **dcba**, and **ecowx** are treated as stores.

The cache management instructions other than **dcbz** and **dcba** never cause a match. If **dcbz** or **dcba** causes a match, some or all of the target memory locations may have been updated.

When a match occurs, a DSI exception is generated. Refer to Section 5.5.3, “DSI Exception (0x00300),” more information on the data address breakpoint facility.

2.1.2.15.1 Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only

The DBCR is a supervisor-level register with SPR310 on the G2_LE core, which is accessible only by using `mtspr` and `mfspir`. The DBCR controls the compare and match type conditions for DABR1 and DABR2. Figure 2-17 shows the format of the DBCR.

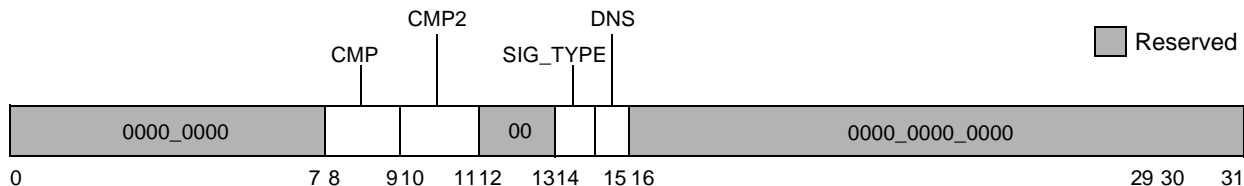


Figure 2-17. Data Address Breakpoint Control Register (DBCR)

Table 2-16 provides the description of DBCR bit settings.

Table 2-16. Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only

Bits	Name	Description
0–7	—	Reserved
8–9	CMP	DABR1 breakpoint compare type 00 Match if data's EA equals DABR[CEA] 01 Reserved 10 Match if data's EA less than DABR[CEA] 11 Match if data's EA greater than or equal to DABR[CEA]
10–11	CMP2	DABR2 breakpoint compare type 00 Match if data's EA equals DABR2[CEA] 01 Reserved 10 Match if data's EA less than DABR2[CEA] 11 Match if data's EA greater than or equal to DABR2[CEA]
12–13	—	Reserved
14	SIG_TYPE	Combinational signal type 0 Data access EA matches DABR[CEA] OR EA matches DABR2[CEA] 1 Data access EA matches DABR[CEA] AND EA matches DABR2[CEA]
15	DNS	Do not signal. Disable <code>core_dabr</code> and <code>core_dabr2</code> output signals. 0 Allow signal to toggle on a match 1 Do not toggle signal on match
16–31	—	Reserved

Chapter 3

Instruction Set Model

This chapter describes the operand conventions as they are represented in two levels of the PowerPC architecture. It also provides detailed descriptions of conventions used for storing values in registers and memory, accessing the core registers, and the representation of data in these registers.

- Operand conventions
- G2 core instruction set

3.1 Operand Conventions

This section describes the integer and floating-point operand conventions. It also describes the big- and little-endian byte ordering for the G2 and G2_LE cores.

3.1.1 Data Organization in Memory and Memory Operands

Bytes in memory are numbered consecutively starting with 0. Each number is the address of the corresponding byte.

Memory operands may be bytes, half words, words, or double words, or, for the load/store multiple and move assist instructions, a sequence of bytes or words. The address of a memory operand is the address of its first byte (that is, of its lowest-numbered byte). Operand length is implicit for each instruction.

3.1.2 Endian Modes and Byte Ordering

The PowerPC architecture supports both big- and little-endian byte ordering. The default byte and bit ordering is big-endian. See Section 3.1.2, “Byte Ordering,” in the *Programming Environments Manual*, for more information about big- and little-endian byte ordering.

True little-endian mode is supported in the G2_LE core to minimize the impact on software porting from true little-endian systems. The true little-endian mode applies for all instruction fetches and data load and store operations to and from memory. The G2_LE powers up in one of two endian modes, big-endian mode or true little-endian mode, selected by the core_tle signal at the negation of core_hreset. The endian mode should be set at the

negation of `core_hreset`, and should remain unchanged by software for the duration of the system operation.

Bit 4 of HID2, (HID2[LET]) is used in conjunction with MSR[LE] to indicate the endian mode of operation of the G2_LE core as shown in Table 3-1.

Table 3-1. Endian Mode Indication

MSR[LE]	HID2[LET]	Endian Mode
0	x	Big-endian
1	0	Modified (PowerPC) little-endian
1	1	True little-endian

When the G2_LE core is in true little-endian mode, memory and I/O subsystems are treated as true little-endian. The following occurs when operating in true little-endian mode:

- The byte reversing for instruction occurs before the instruction is decoded.
- The byte reversing for data occurs when the data item is being moved to or from the GPR.

Therefore, the byte reversal in little-endian mode for load or store accesses occurs between memory or the data cache, and the register files for the G2_LE core.

3.1.3 Alignment and Misaligned Accesses

The operand of a single-register memory access instruction has a natural alignment boundary equal to the operand length. In other words, the natural address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is misaligned. For a detailed discussion about memory operands, see Chapter 3, “Operand Conventions,” in the *Programming Environments Manual*.

Operands for single-register memory access instructions have the characteristics shown in Table 3-2. (Although not permitted as memory operands, quad words are shown because quad-word alignment is desirable for certain memory operands.)

Table 3-2. Memory Operands

Operand	Length	Addr[28–31] If Aligned
Byte	8 bits	xxxx
Half word	2 bytes	xxx0
Word	4 bytes	xx00
Double word	8 bytes	x000
Quad word	16 bytes	0000

Note: An x in an address bit position indicates that the bit can be 0 or 1 independent of the state of other address bits.

The concept of alignment is also applied more generally to data in memory. For example, a 12-byte data item is said to be word-aligned if its address is a multiple of four.

Implementation Notes—The following describes how the G2 core handles alignment and misaligned accesses:

- The G2 core provides hardware support for some misaligned memory accesses. However, misaligned accesses suffer a performance degradation compared to aligned accesses of the same type.
- The G2 core does not provide hardware support for floating-point load/store operations that are not word-aligned. In such a case, the core invokes an alignment exception and the exception handler must break up the misaligned access. For this reason, floating-point single- and double-word accesses should always be word-aligned. Note that a floating-point double-word access on a word-aligned boundary requires an extra cycle to complete.

Any half-word, word, double-word, and string reference access that crosses an alignment boundary must be broken into multiple discrete accesses. For string accesses, the hardware makes no attempt to get aligned to reduce the number of accesses. (Multiple word accesses are architecturally required to be aligned.) The resulting performance degradation depends on how well each individual access behaves with respect to the memory hierarchy. At a minimum, additional cache access cycles are required. More dramatically, each discrete access to a noncacheable page involves an individual bus operation that reduces the effective bus bandwidth.

The frequent use of misaligned accesses is discouraged because they can compromise the overall performance.

3.1.4 Floating-Point Execution Model

The G2 core provides hardware support for all single- and double-precision floating-point operations for most value representations and all rounding modes. The PowerPC

architecture provides for hardware to implement a floating-point system as defined in ANSI/IEEE Standard 754-1985, *IEEE Standard for Binary Floating Point Arithmetic*. For detailed information about the floating-point execution model, refer to Chapter 3, “Operand Conventions,” in the *Programming Environments Manual*.

The IEEE 754 standard includes 64- and 32-bit arithmetic. The standard requires that single-precision arithmetic be provided for single-precision operands. The standard permits double-precision arithmetic instructions to have either (or both) single-precision or double-precision operands, but states that single-precision arithmetic instructions should not accept double-precision operands.

The UISA follows these guidelines:

- Double-precision arithmetic instructions may have single-precision operands but always produce double-precision results.
- Single-precision arithmetic instructions require all operands to be single-precision and always produce single-precision results.

For arithmetic instructions, conversions from double- to single-precision must be done explicitly by software, while conversions from single- to double-precision are done implicitly.

All PowerPC implementations provide the equivalent of the following execution models to ensure that identical results are obtained. The definition of the arithmetic instructions for infinities, denormalized numbers, and NaNs follow conventions described in the following sections.

Although the double-precision format specifies an 11-bit exponent, exponent arithmetic uses two additional bit positions to avoid potential transient overflow conditions. An extra bit is required when denormalized double-precision numbers are prenormalized. A second bit is required to permit computation of the adjusted exponent value in the following examples when the corresponding exception enable bit is one:

- Underflow during multiplication using a denormalized factor
- Overflow during division using a denormalized divisor

3.1.5 Effect of Operand Placement on Performance

The VEA states that the placement (location and alignment) of operands in memory affect the relative performance of memory accesses. The best performance is guaranteed if memory operands are aligned on natural boundaries. To obtain the best performance from the core, the programmer should assume the performance model described in Chapter 3, “Operand Conventions,” in the *Programming Environments Manual*.

3.2 Instruction Set Summary

This section describes instructions and addressing modes defined for the G2 core. These instructions are divided into the following functional categories:

- Integer instructions—These include arithmetic and logical instructions. For more information, see Section 3.2.4.1, “Integer Instructions.”
- Floating-point instructions—These include floating-point arithmetic instructions, as well as instructions that affect the floating-point status and control register (FPSCR). For more information, see Section 3.2.4.2, “Floating-Point Instructions.”
- Load and store instructions—These include integer and floating-point load and store instructions. For more information, see Section 3.2.4.3, “Load and Store Instructions.”
- Flow control instructions—These include branching instructions, condition register logical instructions, and other instructions that affect the instruction flow. For more information, see Section 3.2.4.4, “Branch and Flow Control Instructions.”
- Trap instructions—These are used to test for a specified set of conditions; see Section 3.2.4.5, “Trap Instructions.”
- Processor control instructions—These are used for synchronizing memory accesses and managing caches, TLBs, and segment registers. For more information, see Section 3.2.4.6, “Processor Control Instructions,” Section 3.2.5.1, “Processor Control Instructions,” and Section 3.2.6.2, “Processor Control Instructions—OEA.”
- Memory synchronization instructions—These are used for synchronizing memory accesses. See Section 3.2.4.7, “Memory Synchronization Instructions—UISA” and Section 3.2.5.2, “Memory Synchronization Instructions—VEA.”
- Memory control instructions—These provide control of caches, TLBs, and segment registers. For more information, see Section 3.2.5.3, “Memory Control Instructions—VEA” and Section 3.2.6.3, “Memory Control Instructions—OEA.”
- System linkage instructions—These include the System Call (**sc**) and Return from Interrupt (**rfi**) instructions. See Section 3.2.6.1, “System Linkage Instructions.”
- External control instructions—These include instructions for use with special input/output devices. See Section 3.2.5.4, “External Control Instructions.”

Note that this grouping of instructions does not necessarily indicate the execution unit that processes a particular instruction or group of instructions. This information, which is useful in taking full advantage of the G2 core superscalar parallel instruction execution, is provided in Chapter 8, “Instruction Set,” of the *Programming Environments Manual*.

Integer instructions operate on word operands. Floating-point instructions operate on single- and double-precision floating-point operands. PowerPC instructions are 4-byte words. The UISA provides for byte, half-word, and word operand loads and stores between

memory and a set of 32 GPRs. It also provides for word and double-word operand loads and stores between memory and a set of 32 FPRs.

Arithmetic and logical instructions do not read or modify memory. To use the contents of a memory location in a computation and then modify the same or another memory location, the memory contents must be loaded into a register, modified, and then written to the target location using load and store instructions.

The description of each instruction includes the mnemonic and a formatted list of operands. To simplify assembly language programming, a set of simplified mnemonics (extended mnemonics in the architecture specification) and symbols is provided for some of the frequently-used instructions; see Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, for a complete list of simplified mnemonic examples.

3.2.1 Classes of Instructions

The G2 core instructions belong to one of the following three classes:

- Defined
- Illegal
- Reserved

Note that although the definitions of these terms are consistent among the processors of this family, the assignment of these classifications is not. For example, an instruction that is specific to 64-bit implementations is considered defined for 64-bit implementations but illegal for 32-bit implementations such as the G2 core.

The class is determined by examining the primary opcode and the extended opcode, if any. If either is not that of a defined instruction or of a reserved instruction, the instruction is illegal.

In future versions of the PowerPC architecture, instruction codings that are now illegal may become assigned to instructions in the architecture or may be reserved by being assigned to processor-specific instructions.

3.2.1.1 Definition of Boundedly Undefined

If instructions are encoded with incorrectly set bits in reserved fields, the results on execution can be said to be boundedly undefined. If a user-level program executes the incorrectly coded instruction, the resulting undefined results are bounded in that a spurious change from user to supervisor state is not allowed, and the level of privilege exercised by the program in relation to memory access and other system resources cannot be exceeded. Boundedly undefined results for a given instruction may vary between implementations, and between execution attempts in the same implementation.

3.2.1.2 Defined Instruction Class

Defined instructions are guaranteed to be supported in all PowerPC implementations, except as stated in the instruction descriptions in Chapter 8, “Instruction Set,” in the *Programming Environments Manual*. The G2 core provides hardware support for all instructions defined for 32-bit implementations.

A processor of this family invokes the illegal instruction error handler (part of the program exception) when the unimplemented PowerPC instructions are encountered so they can be emulated in software, as required.

A defined instruction can have invalid forms, as described in the following section.

3.2.1.3 Illegal Instruction Class

Illegal instructions are grouped into the following categories:

- Instructions not defined in the PowerPC architecture. These opcodes are available for future extensions of the PowerPC architecture; that is, future versions of the PowerPC architecture may define any of these instructions to perform new functions.

The following primary opcodes are defined as illegal but may be used in future extensions to the architecture:

1, 4, 5, 6, 9, 22, 56, 57, 60, 61

- Instructions defined in the PowerPC architecture but not implemented in a specific PowerPC implementation. For example, instructions that can be executed on 64-bit processors are considered illegal by 32-bit processor cores.

The following primary opcodes are defined for 64-bit implementations only and are illegal on the core:

2, 30, 58, 62

- All unused extended opcodes are illegal. The unused extended opcodes can be determined from information in Appendix A.2, “Instructions Sorted by Opcode,” and Section 3.2.1.4, “Reserved Instruction Class.” Notice that extended opcodes for instructions that are defined only for 64-bit implementations are illegal in 32-bit implementations, and vice versa.

The following primary opcodes have unused extended opcodes:

17, 19, 31, 59, 63 (primary opcodes 30 and 62 are illegal for all 32-bit implementations, but as 64-bit opcodes they have some unused extended opcodes)

- An instruction consisting entirely of zeros is guaranteed to be an illegal instruction. This increases the probability that an attempt to execute data or uninitialized memory invokes the system illegal instruction error handler (a program exception). Note that if only the primary opcode consists of all zeros, the instruction is

considered a reserved instruction. This is further described in Section 3.2.1.4, “Reserved Instruction Class.”

An attempt to execute an illegal instruction invokes the illegal instruction error handler (a program exception) but has no other effect. Section 5.5.7, “Program Exception (0x00700),” describes illegal and invalid instruction exceptions.

Except for an instruction consisting entirely of binary zeros, illegal instructions are available for further additions to the PowerPC architecture.

3.2.1.4 Reserved Instruction Class

Reserved instructions are allocated to specific implementation-dependent purposes not defined by the PowerPC architecture. An attempt to execute an unimplemented reserved instruction invokes the illegal instruction error handler (a program exception). See Section 5.5.7, “Program Exception (0x00700),” for additional information about illegal and invalid instruction exceptions.

The following types of instructions are included in this class:

- Implementation-specific instructions (for example, Load Data TLB Entry (**tlbld**) and Load Instruction TLB Entry (**tlbli**) instructions).
- Optional instructions defined by the PowerPC architecture but not implemented by the core (for example, Floating Square Root (**fsqrt**) and Floating Square Root Single (**fsqrts**) instructions).

3.2.2 Addressing Modes

This section provides an overview of conventions for addressing memory and calculating effective addresses as defined by the PowerPC architecture for 32-bit implementations. For more detailed information, see “Conventions” in Chapter 4, “Addressing Modes and Instruction Set Summary,” of the *Programming Environments Manual*.

3.2.2.1 Memory Addressing

A program references memory using the effective (logical) address computed by the processor when it executes a memory access or branch instruction or when it fetches the next sequential instruction.

As described in Section 3.1.1, “Data Organization in Memory and Memory Operands,” bytes in memory are numbered consecutively starting with zero. Each number is the address of the corresponding byte.

3.2.2.2 Memory Operands

Memory operands may be bytes, half words, words, or double words, or, for the load/store multiple and load/store string instructions, a sequence of bytes or words. The address of a memory operand is the address of its first byte (that is, of its lowest-numbered byte). Operand length is implicit for each instruction. The PowerPC architecture supports both big- and little-endian byte ordering. The default byte and bit ordering is big-endian. See Section 3.1.2, “Byte Ordering,” in the *Programming Environments Manual*, for more information about big- and little-endian byte ordering.

The operand of a single-register memory access instruction has a natural alignment boundary equal to the operand length. In other words, the “natural” address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is misaligned. For a detailed discussion about memory operands, see Chapter 3, “Operand Conventions,” in the *Programming Environments Manual*.

3.2.2.3 Effective Address Calculation

An effective address (EA) is the 32-bit sum computed by the processor core when executing a memory access or branch instruction or when fetching the next sequential instruction. For a memory access instruction, if the sum of the effective address and the operand length exceeds the maximum effective address, the memory operand is considered to wrap around from the maximum effective address through effective address 0, as described in the following paragraphs.

Effective address computations for both data and instruction accesses use 32-bit unsigned binary arithmetic. A carry from bit 0 is ignored.

Load and store operations have three categories of effective address generation:

- Register indirect with immediate index mode
- Register indirect with index mode
- Register indirect mode

Section 3.2.4.3.2, “Integer Load and Store Address Generation,” describes effective address generation for load and store operations.

Branch instructions have three categories of effective address generation:

- Immediate
- Link register indirect
- Count register indirect

Section 3.2.4.4.1, “Branch Instruction Address Calculation,” describes branch instruction effective address generation.

3.2.2.4 Synchronization

The synchronization described in this section refers to the state of the core performing the synchronization.

3.2.2.4.1 Context Synchronization

The System Call (**sc**) and Return from Interrupt (**rfi**) instructions perform context synchronization by allowing previously issued instructions to complete before performing a change in context. Execution of one of these instructions ensures the following:

- No higher priority exception exists (**sc**).
- All previous instructions have completed to a point where they can no longer cause an exception. If a prior memory access instruction causes direct-store error exceptions, the results are guaranteed to be determined before this instruction is executed.
- Previous instructions complete execution in the context (privilege, protection, and address translation) under which they were issued.
- The instructions following the **sc** or **rfi** instruction execute in the context established by these instructions.

3.2.2.4.2 Execution Synchronization

An instruction is execution synchronizing if all previously initiated instructions appear to have completed before the instruction is initiated or, in the case of the Synchronize (**sync**) and Instruction Synchronize (**isync**) instructions, before the instruction completes. For example, the Move to Machine State Register (**mtmsr**) instruction is execution synchronizing. It ensures that all preceding instructions have completed execution and will not cause an exception before the instruction executes but does not ensure subsequent instructions execute in the newly established environment. For example, if the **mtmsr** sets MSR[PR], unless an **isync** immediately follows the **mtmsr** instruction, a privileged instruction could be executed or privileged access could be performed without causing an exception even though MSR[PR] indicates user mode.

3.2.2.4.3 Instruction-Related Exceptions

There are two kinds of exceptions in the G2 core—those caused directly by the execution of an instruction and those caused by an asynchronous event. Either may cause components of the system software to be invoked.

Exceptions can be caused directly by the execution of an instruction as follows:

- An attempt to execute an illegal instruction causes the illegal instruction (program exception) handler to be invoked. An attempt by a user-level program to execute the supervisor-level instructions listed below causes the privileged instruction (program

exception) handler to be invoked. The core provides the following supervisor-level instructions: **dcbi** (this instruction should never be used on the G2 core), **mfmsr**, **mfmspr**, **mfsr**, **mfsrin**, **mtmsr**, **mtspr**, **mtsr**, **mtsrin**, **rfi**, **tlbie**, **tlbsync**, **tlbld**, and **tlbli**. Note that the privilege level of the **mfmspr** and **mtspr** instructions depends on the SPR encoding.

- An attempt to access memory that is not available (page fault) causes the ISI exception handler to be invoked.
- An attempt to access memory with an effective address alignment that is invalid for the instruction causes the alignment exception handler to be invoked.
- The execution of an **sc** instruction invokes the system call exception handler that permits a program to request the system to perform a service.
- The execution of a trap instruction invokes the program exception trap handler.
- The execution of a floating-point instruction when floating-point instructions are disabled or unavailable invokes the floating-point unavailable exception handler.
- The execution of an instruction that causes a floating-point exception while exceptions are enabled in the MSR invokes the program exception handler.

Exceptions caused by asynchronous events are described in Chapter 5, “Exceptions.”

3.2.3 Instruction Set Overview

This section provides a brief overview of the PowerPC instructions implemented in the core and highlights any special information with respect to how the G2 core implements a particular instruction. Note that the categories used in this section correspond to those used in Chapter 4, “Addressing Modes and Instruction Set Summary,” in the *Programming Environments Manual*. These categorizations are somewhat arbitrary and are provided for the convenience of the programmer and do not necessarily reflect the PowerPC architecture specification.

Note that some of the instructions have the following optional features:

- CR Update—The dot (.) suffix on the mnemonic enables the update of the CR.
- Overflow option—The **o** suffix indicates that the overflow bit in the XER is enabled.

3.2.4 PowerPC UISA Instructions

The UISA includes the base user-level instruction set (excluding a few user-level cache control, synchronization, and time base instructions), user-level registers, programming model, data types, and addressing modes. This section discusses the instructions defined in the UISA.

3.2.4.1 Integer Instructions

This section describes the integer instructions. These consist of the following:

- Integer arithmetic instructions
- Integer compare instructions
- Integer logical instructions
- Integer rotate and shift instructions

Integer instructions use the content of the GPRs as source operands and place results into GPRs, into the XER, and into condition register (CR) fields.

3.2.4.1.1 Integer Arithmetic Instructions

Table 3-3 lists the integer arithmetic instructions for the core.

Table 3-3. Integer Arithmetic Instructions

Name	Mnemonic	Operand Syntax
Add	add (add. addo addo.)	rD,rA,rB
Add Carrying	addc (addc. addco addco.)	rD,rA,rB
Add Extended	adde (adde. addeo addeo.)	rD,rA,rB
Add Immediate	addi	rD,rA,SIMM
Add Immediate Carrying	addic	rD,rA,SIMM
Add Immediate Carrying and Record	addic.	rD,rA,SIMM
Add Immediate Shifted	addis	rD,rA,SIMM
Add to Minus One Extended	addme (addme. addmeo addmeo.)	rD,rA
Add to Zero Extended	addze (addze. addzeo addzeo.)	rD,rA
Divide Word	divw (divw. divwo divwo.)	rD,rA,rB
Divide Word Unsigned	divwu (divwu. divwuo divwuo.)	rD,rA,rB
Multiply High Word	mulhw (mulhw.)	rD,rA,rB
Multiply High Word Unsigned	mulhwu (mulhwu.)	rD,rA,rB
Multiply Low	mullw (mullw. mullwo mullwo.)	rD,rA,rB
Multiply Low Immediate	mulli	rD,rA,SIMM
Negate	neg (neg. nego nego.)	rD,rA
Subtract From	subf (subf. subfo subfo.)	rD,rA,rB
Subtract From Carrying	subfc (subfc. subfco subfco.)	rD,rA,rB
Subtract From Extended	subfe (subfe. subfeo subfeo.)	rD,rA,rB
Subtract From Immediate Carrying	subfic	rD,rA,SIMM
Subtract From Minus One Extended	subfme (subfme. subfmeo subfmeo.)	rD,rA
Subtract From Zero Extended	subfze (subfze. subfzeo subfzeo.)	rD,rA

Although there is no Subtract Immediate instruction, its effect can be achieved by using an **addi** instruction with the immediate operand negated. Simplified mnemonics are provided that include this negation. The **subf** instructions subtract the second operand (**rA**) from the third operand (**rB**). Simplified mnemonics are provided in which the third operand is subtracted from the second operand. See Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, for examples.

3.2.4.1.2 Integer Compare Instructions

The integer compare instructions algebraically or logically compare the contents of **rA** with either the UIMM operand, the SIMM operand, or the contents of **rB**. The comparison is signed for the **cmpi** and **cmp** instructions, and unsigned for the **cmpli** and **cmpl** instructions. Table 3-4 lists the integer compare instructions.

Table 3-4. Integer Compare Instructions

Name	Mnemonic	Operand Syntax
Compare	cmp	crfD ,L, rA , rB
Compare Immediate	cmpi	crfD ,L, rA ,SIMM
Compare Logical	cmpl	crfD ,L, rA , rB
Compare Logical Immediate	cmpli	crfD ,L, rA ,UIMM

The **crfD** operand can be omitted if the result of the comparison is to be placed in CR0. Otherwise, the target CR field must be specified in the instruction **crfD** field.

For more information refer to Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*.

3.2.4.1.3 Integer Logical Instructions

The logical instructions shown in Table 3-5 perform bit-parallel operations. Logical instructions with the CR update enabled and instructions **andi**. and **andis**. set CR field CR0 to characterize the result of the logical operation. These fields are set as if the sign-extended low-order 32 bits of the result were algebraically compared to zero. Logical instructions without CR update and the remaining logical instructions do not modify the CR. Logical instructions do not affect the XER[SO], XER[OV], and XER[CA] bits.

For simplified mnemonics examples for the integer logical operations see Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*.

Table 3-5. Integer Logical Instructions

Name	Mnemonic	Operand Syntax
AND	and (and.)	rA,rS,rB
AND Immediate	andi .	rA,rS,UIMM
AND Immediate Shifted	andis .	rA,rS,UIMM
AND with Complement	andc (andc.)	rA,rS,rB
Count Leading Zeros Word	cntlzw (cntlzw.)	rA,rS
Equivalent	eqv (eqv.)	rA,rS,rB
Extend Sign Byte	extsb (extsb.)	rA,rS
Extend Sign Half Word	extsh (extsh.)	rA,rS
NAND	nand (nand.)	rA,rS,rB
NOR	nor (nor.)	rA,rS,rB
OR	or (or.)	rA,rS,rB
OR Immediate	ori	rA,rS,UIMM
OR Immediate Shifted	oris	rA,rS,UIMM
OR with Complement	orc (orc.)	rA,rS,rB
XOR	xor (xor.)	rA,rS,rB
XOR Immediate	xori	rA,rS,UIMM
XOR Immediate Shifted	xoris	rA,rS,UIMM

3.2.4.1.4 Integer Rotate and Shift Instructions

Rotation operations are performed on data from a GPR, and the result, or a portion of the result, is returned to a GPR. See Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, for a complete list of simplified mnemonics that allows simpler coding of often-used functions such as clearing the leftmost or rightmost bits of a register, left justifying or right justifying an arbitrary field, and simple rotates and shifts.

Integer rotate instructions rotate the contents of a register. The result of the rotation is either inserted into the target register under control of a mask (if a mask bit is 1, the associated bit of the rotated data is placed into the target register; and if the mask bit is 0, the associated bit in the target register is unchanged), or ANDed with a mask before being placed into the target register.

The integer rotate instructions are listed in Table 3-6.

Table 3-6. Integer Rotate Instructions

Name	Mnemonic	Operand Syntax
Rotate Left Word Immediate then AND with Mask	rlwinm (rlwinm.)	rA,rS,SH,MB,ME
Rotate Left Word Immediate then Mask Insert	rlwimi (rlwimi.)	rA,rS,SH,MB,ME
Rotate Left Word then AND with Mask	rlwnm (rlwnm.)	rA,rS,rB,MB,ME

The integer shift instructions perform left and right shifts. Immediate-form logical (unsigned) shift operations are obtained by specifying masks and shift values for certain rotate instructions. Simplified mnemonics are provided, making coding of such shifts simpler and easier to understand.

Multiple-precision shifts can be programmed as shown in Appendix C, “Multiple-Precision Shifts,” in the *Programming Environments Manual*.

The integer shift instructions are listed in Table 3-7.

Table 3-7. Integer Shift Instructions

Name	Mnemonic	Operand Syntax
Shift Left Word	slw (slw.)	rA,rS,rB
Shift Right Algebraic Word	sraw (sraw.)	rA,rS,rB
Shift Right Algebraic Word Immediate	srawi (srawi.)	rA,rS,SH
Shift Right Word	srw (srw.)	rA,rS,rB

3.2.4.2 Floating-Point Instructions

This section describes the floating-point instructions, which include the following:

- Floating-point arithmetic instructions
- Floating-point multiply-add instructions
- Floating-point rounding and conversion instructions
- Floating-point compare instructions
- Floating-point status and control register instructions
- Floating-point move instructions

See Section 3.2.4.3, “Load and Store Instructions,” for information about floating-point loads and stores.

The PowerPC architecture supports a floating-point system as defined in the IEEE 754 standard, but requires software support to conform with that standard. All floating-point operations conform to the IEEE 754 standard, except if software sets the non-IEEE mode bit (NI) in the FPSCR. The G2 core is in the nondenormalized mode when the NI bit is set in the FPSCR. If a denormalized result is produced, a default result of zero is generated.

The generated zero has the same sign as the denormalized number. The core performs single- and double-precision floating-point operations compliant with the IEEE 754 floating-point standard.

Implementation Note—Single-precision denormalized results require two additional processor clock cycles to round. When loading or storing a single-precision denormalized number, the load/store unit may take up to 24 processor clock cycles to convert between the internal double-precision format and the external single-precision format.

3.2.4.2.1 Floating-Point Arithmetic Instructions

The floating-point arithmetic instructions are listed in Table 3-8.

Table 3-8. Floating-Point Arithmetic Instructions

Name	Mnemonic	Operand Syntax
Floating Add (Double-Precision)	fadd (fadd.)	frD,frA,frB
Floating Add Single	fadds (fadds.)	frD,frA,frB
Floating Divide (Double-Precision)	fdiv (fdiv.)	frD,frA,frB
Floating Divide Single	fdivs (fdivs.)	frD,frA,frB
Floating Multiply (Double-Precision)	fmul (fmul.)	frD,frA,frC
Floating Multiply Single	fmuls (fmuls.)	frD,frA,frC
Floating Reciprocal Estimate Single	fres (fres.)	frD,frB
Floating Reciprocal Square Root Estimate	frsqrte (frsqrte.)	frD,frB
Floating Select	fsel (fsel.)	frD,frA,frC,frB
Floating Subtract (Double-Precision)	fsub (fsub.)	frD,frA,frB
Floating Subtract Single	fsubs (fsubs.)	frD,frA,frB

3.2.4.2.2 Floating-Point Multiply-Add Instructions

These instructions combine multiply and add operations without an intermediate rounding operation. The fractional part of the intermediate product is 106 bits wide, and all 106 bits take part in the add/subtract portion of the instruction.

The floating-point multiply-add instructions are listed in Table 3-9.

Table 3-9. Floating-Point Multiply-Add Instructions

Name	Mnemonic	Operand Syntax
Floating Multiply-Add (Double-Precision)	fmadd (fmadd.)	frD,frA,frC,frB
Floating Multiply-Add Single	fmadds (fmadds.)	frD,frA,frC,frB
Floating Multiply-Subtract (Double-Precision)	fmsub (fmsub.)	frD,frA,frC,frB
Floating Multiply-Subtract Single	fmsubs (fmsubs.)	frD,frA,frC,frB

Table 3-9. Floating-Point Multiply-Add Instructions (continued)

Name	Mnemonic	Operand Syntax
Floating Negative Multiply-Add (Double-Precision)	fnmadd (fnmadd.)	frD,frA,frC,frB
Floating Negative Multiply-Add Single	fnmadds (fnmadds.)	frD,frA,frC,frB
Floating Negative Multiply-Subtract (Double-Precision)	fnmsub (fnmsub.)	frD,frA,frC,frB
Floating Negative Multiply-Subtract Single	fnmsubs (fnmsubs.)	frD,frA,frC,frB

Implementation Note—Single-precision multiply-type instructions operate faster than their double-precision equivalents. See Chapter 7, “Instruction Timing,” for more information.

3.2.4.2.3 Floating-Point Rounding and Conversion Instructions

The Floating Round to Single-Precision (**frsp**) instruction is used to truncate a 64-bit double-precision number to a 32-bit single-precision floating-point number. The floating-point conversion instructions convert a 64-bit double-precision floating-point number to a 32-bit signed integer number.

The PowerPC architecture defines bits 0–31 of floating-point register **frD** as undefined when executing the Floating Convert to Integer Word (**fctiw**) and Floating Convert to Integer Word with Round Toward Zero (**fctiwz**) instructions.

Examples of uses of these instructions to perform various conversions can be found in Appendix D, “Floating-Point Models,” in the *Programming Environments Manual*. The floating-point rounding instructions are shown in Table 3-10.

Table 3-10. Floating-Point Rounding and Conversion Instructions

Name	Mnemonic	Operand Syntax
Floating Convert to Integer Word	fctiw (fctiw.)	frD,frB
Floating Convert to Integer Word with Round Toward Zero	fctiwz (fctiwz.)	frD,frB
Floating Round to Single-Precision	frsp (frsp.)	frD,frB

3.2.4.2.4 Floating-Point Compare Instructions

Floating-point compare instructions compare the contents of two floating-point registers. The comparison ignores the sign of zero (that is $+0 = -0$). The floating-point compare instructions are listed in Table 3-11.

Table 3-11. Floating-Point Compare Instructions

Name	Mnemonic	Operand Syntax
Floating Compare Ordered	fcmpo	crD,frA,frB
Floating Compare Unordered	fcmpu	crD,frA,frB

3.2.4.2.5 Floating-Point Status and Control Register Instructions

Every FPSCR instruction appears to synchronize the effects of all floating-point instructions executed by a given processor. Executing an FPSCR instruction ensures that all floating-point instructions previously initiated by the given processor appear to have completed before the FPSCR instruction is initiated and that no subsequent floating-point instructions appear to be initiated by the given processor until the FPSCR instruction has completed. The FPSCR instructions are listed in Table 3-12.

Table 3-12. Floating-Point Status and Control Register Instructions

Name	Mnemonic	Operand Syntax
Move from FPSCR	mffs (mffs.)	frD
Move to Condition Register from FPSCR	mcrfs	crD,crfS
Move to FPSCR Bit 0	mtfsb0 (mtfsb0.)	crbD
Move to FPSCR Bit 1	mtfsb1 (mtfsb1.)	crbD
Move to FPSCR Field Immediate	mtfsfi (mtfsfi.)	crD,IMM
Move to FPSCR Fields	mtfsf (mtfsf.)	FM,frB

Implementation Note—The architecture notes that, in some implementations, the Move to FPSCR Fields (**mtfsf**) instruction may perform more slowly when only a portion of the fields are updated as opposed to all of the fields. This is not the case in the G2 core.

3.2.4.2.6 Floating-Point Move Instructions

Floating-point move instructions copy data from one floating-point register to another. The floating-point move instructions do not modify the FPSCR. The CR update option in these instructions controls the placing of result status into CR1. Floating-point move instructions are listed in Table 3-13.

Table 3-13. Floating-Point Move Instructions

Name	Mnemonic	Operand Syntax
Floating Absolute Value	fabs (fabs.)	frD,frB
Floating Move Register	fmr (fmr.)	frD,frB
Floating Negate	fneg (fneg.)	frD,frB
Floating Negative Absolute Value	fnabs (fnabs.)	frD,frB

3.2.4.3 Load and Store Instructions

Load and store instructions are issued and translated in program order; however, the accesses can occur out of order. Synchronizing instructions are provided to enforce strict ordering. This section describes the load and store instructions of the G2 core, which consist of the following:

- Integer load instructions
- Integer store instructions
- Integer load and store with byte-reverse instructions
- Integer load and store multiple instructions
- Integer load and store string instructions
- Floating-point load instructions
- Floating-point store instructions

3.2.4.3.1 Self-Modifying Code

When a processor modifies a memory location that may be contained in the instruction cache, software must ensure that memory updates are visible to the instruction fetching mechanism. This can be achieved by the following instruction sequence:

```
dcbst | update memory
sync  | wait for update
icbi  | remove (invalidate) copy in instruction cache
isync | remove copy in own instruction buffer
```

These operations are required because the data cache is a write-back cache. Since instruction fetching bypasses the data cache, changes to items in the data cache may not be reflected in memory until the fetch operations complete.

Special care must be taken to avoid coherency paradoxes in systems that implement unified secondary caches, and designers should carefully follow the guidelines for maintaining cache coherency that are provided in the VEA, and discussed in Chapter 5, “Cache Model and Memory Coherency,” in the *Programming Environments Manual*. Because the core does not broadcast the M bit for instruction fetches, external caches are subject to coherency paradoxes.

3.2.4.3.2 Integer Load and Store Address Generation

Integer load and store operations generate effective addresses using register indirect with immediate index mode, register indirect with index mode, or register indirect mode. See Section 3.2.2.3, “Effective Address Calculation.” Note that the core is optimized for load and store operations that are aligned on natural boundaries, and operations that are not naturally aligned may suffer performance degradation. Refer to Section 5.5.6.1, “Integer Alignment Exceptions.”

3.2.4.3.3 Register Indirect Integer Load Instructions

For integer load instructions, the byte, half word, word, or double word addressed by the EA is loaded into **rD**. Many integer load instructions have an update form, in which **rA** is updated with the generated effective address. For these forms, the EA is placed into **rA** and the memory element (byte, half word, word, or double word) addressed by EA is loaded into **rD**.

Implementation Note—In some implementations, the load half word algebraic instructions (**lha** and **lhax**) and the load with update (**lbzu**, **lbzux**, **lhzu**, **lhzux**, **lhau**, **lhaux**, **lwu**, and **lwux**) instructions may execute with greater latency than other types of load instructions. In the G2 core, these instructions operate with the same latency as other load instructions.

Table 3-14 lists the integer load instructions.

Table 3-14. Integer Load Instructions

Name	Mnemonic	Operand Syntax
Load Byte and Zero	lbz	rD,d(rA)
Load Byte and Zero Indexed	lbzx	rD,rA,rB
Load Byte and Zero with Update	lbzu	rD,d(rA)
Load Byte and Zero with Update Indexed	lbzux	rD,rA,rB
Load Half Word Algebraic	lha	rD,d(rA)
Load Half Word Algebraic Indexed	lhax	rD,rA,rB
Load Half Word Algebraic with Update	lhau	rD,d(rA)
Load Half Word Algebraic with Update Indexed	lhaux	rD,rA,rB
Load Half Word and Zero	lhz	rD,d(rA)
Load Half Word and Zero Indexed	lhzx	rD,rA,rB
Load Half Word and Zero with Update	lhzu	rD,d(rA)
Load Half Word and Zero with Update Indexed	lhzux	rD,rA,rB
Load Word and Zero	lwz	rD,d(rA)
Load Word and Zero Indexed	lwzx	rD,rA,rB
Load Word and Zero with Update	lwzu	rD,d(rA)
Load Word and Zero with Update Indexed	lwzux	rD,rA,rB

3.2.4.3.4 Integer Store Instructions

For integer store instructions, the contents of **rS** are stored into the byte, half word, word, or double word in memory addressed by the effective address. Many store instructions have

an update form, in which **rA** is updated with the EA. For these forms, the following rules apply:

- If **rA** \neq 0, the EA is placed into **rA**.
- If **rS** = **rA**, the contents of **rS** are copied to the target memory element, then the generated EA is placed into **rA** (**rS**).

The G2 core defines store with update instructions with **rA** = 0 and integer store instructions with the CR update option enabled (Rc field, bit 31, in the instruction encoding = 1) to be invalid forms. Table 3-15 provides a list of the integer store instructions for the core.

Table 3-15. Integer Store Instructions

Name	Mnemonic	Operand Syntax
Store Byte	stb	rS,d(rA)
Store Byte Indexed	stbx	rS,rA,rB
Store Byte with Update	stbu	rS,d(rA)
Store Byte with Update Indexed	stbux	rS,rA,rB
Store Half Word	sth	rS,d(rA)
Store Half Word Indexed	sthx	rS,rA,rB
Store Half Word with Update	sthv	rS,d(rA)
Store Half Word with Update Indexed	sthvx	rS,rA,rB
Store Word	stw	rS,d(rA)
Store Word Indexed	stwx	rS,rA,rB
Store Word with Update	stwu	rS,d(rA)
Store Word with Update Indexed	stwux	rS,rA,rB

3.2.4.3.5 Integer Load and Store with Byte-Reverse Instructions

Table 3-16 describes integer load and store with byte-reverse instructions. When used in a system operating with the default big-endian byte order, these instructions have the effect of loading and storing data in little-endian order. Likewise, when used in a system operating with little-endian byte order, these instructions have the effect of loading and storing data in big-endian order. When used in a G2_LE core-based system operating with true little-endian byte order, these instructions have the effect of loading and storing data in true little-endian order. For more information about big- and little-endian byte ordering, see Section 3.1.2, “Byte Ordering,” in the *Programming Environments Manual*. For more information about true little-endian operation, see Section 3.1.2, “Endian Modes and Byte Ordering.”

The G2_LE core supports the true little-endian mode. In true little-endian mode, the core treats the memory and I/O subsystems as little-endian memory. In this case, instruction and data bytes are reserved as follows:

- The byte reversing for instruction accesses occurs before the instruction is decoded.
- The byte reversing occurs for data accesses when the data item is being moved to or from the GPR.

Therefore, byte reversal during the load or store accesses is performed between memory or the data cache, and the register files.

Table 3-16. Integer Load and Store with Byte-Reverse Instructions

Name	Mnemonic	Operand Syntax
Load Half Word Byte-Reverse Indexed	lhbrx	rD,rA,rB
Load Word Byte-Reverse Indexed	lwbrx	rD,rA,rB
Store Half Word Byte-Reverse Indexed	sthbrx	rS,rA,rB
Store Word Byte-Reverse Indexed	stwbrx	rS,rA,rB

Implementation Note—In some implementations, load byte-reverse instructions (**lhbrx** and **lwbrx**) may have greater latency than other load instructions; however, these instructions operate with the same latency as other load instructions in the core.

3.2.4.3.6 Integer Load and Store Multiple Instructions

The integer load/store multiple instructions are used to move blocks of data to and from the GPRs. In some implementations, these instructions are likely to have greater latency and take longer to execute, perhaps much longer, than a sequence of individual load or store instructions that produce the same results.

Implementation Notes—The following describes the G2 core implementation of the load/store multiple instruction:

- The load multiple and store multiple instructions may have operands that require memory accesses crossing a 4-Kbyte page boundary. As a result, these instructions may be interrupted by a DSI exception associated with the address translation of the second page. In this case, the core performs some or all of the memory references from the first page, and none of the memory references from the second page before taking the exception. On return from the DSI exception, the load or store multiple instruction will re-execute from the beginning. For additional information, refer to “DSI Exception (0x00300)” in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.
- The PowerPC architecture defines the load multiple word (**lmw**) instruction with **rA** in the range of registers to be loaded as an invalid form. It defines the load multiple and store multiple instructions with misaligned operands (that is, the EA is not a multiple of four) to cause an alignment exception. The core defines the load multiple

word (**lmw**) instruction with **rA** in the range of registers to be loaded as an invalid form.

- The PowerPC architecture describes some preferred instruction forms for the integer load and store multiple instructions that may perform better than other forms in some implementations. None of these preferred forms affect instruction performance in the G2 core.
- When the core is operating with little-endian byte order, execution of a load or store multiple instruction causes the system alignment error handler to be invoked; see Section 3.1.2, “Byte Ordering,” in the *Programming Environments Manual*, for more information. Table 3-17 lists the integer load and store multiple instructions for the G2 core.

Table 3-17. Integer Load and Store Multiple Instructions

Name	Mnemonic	Operand Syntax
Load Multiple Word	lmw	rD,d(rA)
Store Multiple Word	stmw	rS,d(rA)

3.2.4.3.7 Integer Load and Store String Instructions

The integer load and store string instructions allow movement of data from memory to registers or from registers to memory without concern for alignment. These instructions can be used for a short move between arbitrary memory locations or to initiate a long move between misaligned memory fields.

When the core is operating with little-endian byte order, execution of a load or store string instruction causes the system alignment error handler to be invoked; see Section 3.1.2, “Byte Ordering,” in the *Programming Environments Manual*, for more information.

Table 3-18 lists the integer load and store string instructions.

Table 3-18. Integer Load and Store String Instructions

Name	Mnemonic	Operand Syntax
Load String Word Immediate	lswi	rD,rA,NB
Load String Word Indexed	lswx	rD,rA,rB
Store String Word Immediate	stswi	rS,rA,NB
Store String Word Indexed	stswx	rS,rA,rB

Load string and store string instructions may involve operands that are not word-aligned. As described in “Alignment Exception (0x00600)” in Chapter 6, “Exceptions,” in the *Programming Environments Manual*, a misaligned string operation suffers a performance penalty compared to a word-aligned operation of the same type.

When a string operation crosses a 4-Kbyte boundary, the instruction may be interrupted by a DSI exception associated with the address translation of the second page. In this case, the core performs some or all memory references from the first page and none from the second before taking the exception. On return from the DSI exception, the load or store string instruction will re-execute from the beginning. For more information, refer to “DSI Exception (0x00300)” in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

Implementation Note—If **rA** is in the range of registers to be loaded for a Load String Word Immediate (**lswi**) instruction or if either **rA** or **rB** is in the range of registers to be loaded for a Load String Word Indexed (**lswx**) instruction, the PowerPC architecture defines the instruction to be of an invalid form. In addition, the **lswx** and **stswx** instructions that specify a string length of zero are defined to be invalid by the PowerPC architecture. However, none of these cases hold true for the G2 core—the core treats these cases as valid forms.

3.2.4.3.8 Floating-Point Load and Store Address Generation

Floating-point load and store operations generate effective addresses using the register indirect with immediate index addressing mode and register indirect with index addressing mode (details are described below). Floating-point loads and stores are not supported for direct-store accesses. The use of the floating-point load and store operations for direct-store accesses results in a DSI exception.

3.2.4.3.9 Floating-Point Load Instructions

Separate floating-point load instructions are used for single-precision and double-precision operands. Because FPRs support only double-precision format, the FPU converts single-precision data to double-precision format before loading the operands into the target FPR. This conversion is described fully in “Floating-Point Load Instructions” in Appendix D, “Floating-Point Models,” in the *Programming Environments Manual*.

Implementation Note—The PowerPC architecture defines load with update instructions with **rA = 0** as an invalid form; however, the core treats this case as a valid form.

Table 3-19 provides a list of the floating-point load instructions.

Table 3-19. Floating-Point Load Instructions

Name	Mnemonic	Operand Syntax
Load Floating-Point Double	lfd	frD,d(rA)
Load Floating-Point Double Indexed	lfdx	frD,rA,rB
Load Floating-Point Double with Update	lfdU	frD,d(rA)
Load Floating-Point Double with Update Indexed	lfdUX	frD,rA,rB

Table 3-19. Floating-Point Load Instructions (continued)

Name	Mnemonic	Operand Syntax
Load Floating-Point Single	lfs	frD,d(rA)
Load Floating-Point Single Indexed	lfsx	frD,rA,rB
Load Floating-Point Single with Update	lfsu	frD,d(rA)
Load Floating-Point Single with Update Indexed	lfsux	frD,rA,rB

3.2.4.3.10 Floating-Point Store Instructions

There are three basic forms of the store instruction—single-precision, double-precision, and integer. The integer form is supported by the optional **stfiwx** instruction. Because the FPRs support only double-precision format for floating-point data, the FPU converts double-precision data to single-precision format before storing the operands. The conversion steps are described in “Floating-Point Store Instructions” in Appendix D, “Floating-Point Models,” in the *Programming Environments Manual*.

Implementation Note—The PowerPC architecture defines store with update instructions with **rA = 0** as an invalid form; however, the core treats this case as valid.

Table 3-20 lists the floating-point store instructions.

Table 3-20. Floating-Point Store Instructions

Name	Mnemonic	Operand Syntax
Store Floating-Point as Integer Word Indexed	stfiwx	frS,rA,rB
Store Floating-Point Double	stfd	frS,d(rA)
Store Floating-Point Double Indexed	stfdx	frS,rA,rB
Store Floating-Point Double with Update	stfdu	frS,d(rA)
Store Floating-Point Double with Update Indexed	stfdux	frS,rA,rB
Store Floating-Point Single	stfs	frS,d(rA)
Store Floating-Point Single Indexed	stfsx	frS,rA,rB
Store Floating-Point Single with Update	stfsu	frS,d(rA)
Store Floating-Point Single with Update Indexed	stfsux	frS,rA,rB

3.2.4.4 Branch and Flow Control Instructions

Branch instructions are executed by the branch processing unit (BPU). The BPU receives branch instructions from the fetch unit and performs CR lookahead operations on conditional branches to resolve them early, achieving the effect of a zero-cycle branch in many cases.

Some branch instructions can redirect instruction execution conditionally based on the value of bits in the CR. When the branch processor encounters one of these instructions, it scans the execution pipelines to determine whether an instruction in progress may affect the particular CR bit. If no interlock is found, the branch can be resolved immediately by checking the bit in the CR and taking the action defined for the branch instruction.

If an interlock is detected, the branch is considered unresolved and the direction of the branch is predicted using static branch prediction as described in “Conditional Branch Control” in Chapter 4, “Addressing Modes and Instruction Set Summary,” in the *Programming Environments Manual*. The interlock is monitored while instructions are fetched for the predicted branch. When the interlock is cleared, the branch processor determines whether the prediction was correct, based on the value of the CR bit. If the prediction is correct, the branch is considered completed and instruction fetching continues. If the prediction is incorrect, the fetched instructions are purged, and instruction fetching continues along the alternate path. See Chapter 8, “Instruction Timing,” in the *Programming Environments Manual*, for more information about how branches are executed.

3.2.4.4.1 Branch Instruction Address Calculation

Branch instructions can change the instruction sequence. Instruction addresses are always assumed to be word aligned; the processor ignores the two low-order bits of the generated branch target address.

Branch instructions compute the effective address (EA) of the next instruction address using the following addressing modes:

- Branch relative
- Branch conditional to relative address
- Branch to absolute address
- Branch conditional to absolute address
- Branch conditional to link register
- Branch conditional to count register

3.2.4.4.2 Branch Instructions

Table 3-21 lists the branch instructions provided by the processors that implement the PowerPC architecture. To simplify assembly language programming, a set of simplified mnemonics and symbols is provided for the most frequently used forms of branch conditional, compare, trap, rotate and shift, and certain other instructions. See Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, for a list of simplified mnemonic examples.

Table 3-21. Branch Instructions

Name	Mnemonic	Operand Syntax
Branch	b (ba bl bla)	target_addr
Branch Conditional	bc (bca bcl bcla)	BO,BI,target_addr
Branch Conditional to Count Register	bcctr (bcctrl)	BO,BI
Branch Conditional to Link Register	bclr (bclrl)	BO,BI

3.2.4.4.3 Condition Register Logical Instructions

Condition register logical instructions, shown in Table 3-22, and the Move Condition Register Field (**mcrf**) instruction are also defined as flow control instructions, although they are executed by the system register unit (SRU). Most instructions executed by the SRU are completion-serialized to maintain system state; that is, the instruction is held for execution in the SRU until all prior instructions issued have completed.

Table 3-22. Condition Register Logical Instructions

Name	Mnemonic	Operand Syntax
Condition Register AND	crand	crbD,crbA,crbB
Condition Register AND with Complement	crandc	crbD,crbA,crbB
Condition Register Equivalent	creqv	crbD,crbA,crbB
Condition Register NAND	crnand	crbD,crbA,crbB
Condition Register NOR	crnor	crbD,crbA,crbB
Condition Register OR	cror	crbD,crbA,crbB
Condition Register OR with Complement	crorc	crbD,crbA,crbB
Condition Register XOR	crxor	crbD,crbA,crbB
Move Condition Register Field	mcrf	crfD,crfS

Note that if the LR update option is enabled for any of these instructions, these forms of the instructions are invalid in the G2 core.

3.2.4.5 Trap Instructions

The trap instructions shown in Table 3-23 are provided to test for a specified set of conditions. If any of the conditions tested by a trap instruction are met, the system trap handler is invoked. If the tested conditions are not met, instruction execution continues normally.

Table 3-23. Trap Instructions

Name	Mnemonic	Operand Syntax
Trap Word	tw	TO,rA,rB
Trap Word Immediate	twi	TO,rA,SIMM

See Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, for a complete set of simplified mnemonics.

3.2.4.6 Processor Control Instructions

UIA-level processor control instructions are used to read from and write to the condition register (CR).

3.2.4.6.1 Move To/From Condition Register Instructions

Table 3-24 lists the instructions provided by the G2 core for reading from or writing to the CR.

Table 3-24. Move To/From Condition Register Instructions

Name	Mnemonic	Operand Syntax
Move from Condition Register	mfcr	rD
Move to Condition Register Fields	mtrcf	CRM,rS
Move to Condition Register from XER	mcrxr	crfD

3.2.4.7 Memory Synchronization Instructions—UIA

Memory synchronization instructions control the order in which memory operations are completed with respect to asynchronous events and the order in which memory operations are seen by other processors or memory access mechanisms. See Chapter 4, “Instruction and Data Cache Operation,” for additional information about these instructions and about related aspects of memory synchronization.

The **sync** instruction delays execution of subsequent instructions until previous instructions have completed to the point that they can no longer cause an exception and until all previous memory accesses are performed globally; the **sync** operation is not broadcast onto the G2 core 60x bus interface. Additionally, all load and store cache/bus activities initiated by prior instructions are completed. Touch load operations (**dcbt** and **dcbtst**) are required to complete at least through address translation but are not required to complete on the bus.

The functions performed by the **sync** instruction normally take a significant amount of time to complete; as a result, frequent use of this instruction may adversely affect performance. In addition, the number of cycles required to complete a **sync** instruction depends on system parameters and on the processor's state when the instruction is issued.

The proper paired use of the **lwarx** and **stwcx.** instructions allows programmers to emulate common semaphore operations such as test and set, compare and swap, exchange memory, and fetch and add. Examples of these operations can be found in Appendix E, “Synchronization Programming Examples,” in the *Programming Environments Manual*. Typically, the **lwarx** instruction should be paired with an **stwcx.** instruction with the same

effective address used for both instructions of the pair. Note that the reservation granularity is 32 bytes.

The concept behind the use of the **lwarx** and **stwcx.** instructions is that a processor may load a semaphore from memory, compute a result based on the value of the semaphore, and conditionally store it back to the same location (only if that location has not been modified since it was first read), and determine if the store was successful. The conditional store is performed, based on the existence of a reservation established by the preceding **lwarx** instruction. If the reservation exists when the store is executed, the store is performed which sets a bit in the CR. If the reservation does not exist when the store is executed, the target memory location is not modified and a bit is cleared in the CR.

If the store was successful, the sequence of instructions from the read of the semaphore to the store that updated the semaphore appear to have been executed atomically (that is, no other processor or mechanism modified the semaphore location between the read and the update), thus providing the equivalent of a real atomic operation. However, in reality, other cores may have read from the location during this operation. In the G2 core, the reservations are made on behalf of aligned 32-byte sections of the memory address space.

The **lwarx** and **stwcx.** instructions require the EA to be aligned. Exception handling software should not attempt to emulate a misaligned **lwarx** or **stwcx.** instruction, because there is no correct way to define the address associated with the reservation.

In general, the **lwarx** and **stwcx.** instructions should be used only in system programs, which can be invoked by application programs as needed.

At most, one reservation exists simultaneously on any processor. The address associated with the reservation can be changed by a subsequent **lwarx** instruction. The conditional store is performed, based on the existence of a reservation established by the preceding **lwarx** regardless of whether the address generated by the **lwarx** matches that generated by the **stwcx.** instruction. A reservation held by the processor is cleared by one of the following:

- Executing an **stwcx.** instruction to any address
- Attempt by some other device to modify a location in the reservation granularity (32 bytes)

The **lwarx** and **stwcx.** instructions to write-through memory do not cause a DSI exception.

Table 3-25 lists the UISA memory synchronization instructions for the G2 core.

Table 3-25. Memory Synchronization Instructions—UISA

Name	Mnemonic	Operand Syntax
Load Word and Reserve Indexed	lwarx	rD,rA,rB
Store Word Conditional Indexed	stwcx.	rS,rA,rB
Synchronize	sync	—

3.2.5 PowerPC VEA Instructions

The VEA describes the semantics of the memory model that can be assumed by software processes, and includes descriptions of the cache model, cache-control instructions, address aliasing, and other related issues.

3.2.5.1 Processor Control Instructions

The VEA defines the Move from Time Base (**mftb**) instruction for reading the contents of the time base register. The **mftb** is a user-level instruction, as shown in Table 3-26.

Table 3-26. Move From Time Base Instruction

Name	Mnemonic	Operand Syntax
Move from Time Base	mftb	rD, TBR

Simplified mnemonics are provided for the **mftb** instruction so it can be coded with the TBR name as part of the mnemonic rather than requiring it to be coded as an operand. The **mftb** instruction serves as both a basic and simplified mnemonic. Assemblers recognize an **mftb** mnemonic with two operands as the basic form, and an **mftb** mnemonic with one operand as the simplified form. Simplified mnemonics are also provided for Move from Time Base Upper (**mftbu**), a variant of the **mftb** instruction rather than of **mfspr**. The core ignores the extended opcode differences between **mftb** and **mfspr** by ignoring bit 25 of both instructions and treating them identically. Refer to Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*.

3.2.5.2 Memory Synchronization Instructions—VEA

Memory synchronization instructions control the order in which memory operations are performed with respect to asynchronous events, and the order in which memory operations are seen by other processors or memory access mechanisms. See Chapter 4, “Instruction and Data Cache Operation,” for additional information about these instructions and about related aspects of memory synchronization.

Implementation Notes—The following describes how the core handles memory synchronization in the VEA.

- The Instruction Synchronize (**isync**) instruction causes the core to discard all prefetched instructions, wait for any preceding instructions to complete, and then branch to the next sequential instruction (having the effect of clearing the pipeline behind the **isync** instruction).
- The Enforce In-Order Execution of I/O (**eieio**) instruction is used to ensure memory reordering of noncacheable memory access. Because the core does not reorder noncacheable memory accesses, the **eieio** instruction is treated as a no-op.

Table 3-27 lists the VEA memory synchronization instructions for the G2 core.

Table 3-27. Memory Synchronization Instructions—VEA

Name	Mnemonic	Operand Syntax
Enforce In-Order Execution of I/O	eieio	—
Instruction Synchronize	isync	—

3.2.5.3 Memory Control Instructions—VEA

Memory control instructions include the following types:

- Cache management instructions
- Segment register manipulation instructions
- Translation lookaside buffer management instructions

This section describes the user-level cache management instructions defined by the VEA. See Section 3.2.6.3, “Memory Control Instructions—OEA,” for information about supervisor-level cache, segment register manipulation, and translation lookaside buffer management instructions.

The instructions listed in Table 3-28 provide user-level programs the ability to manage on-chip caches when they exist.

Table 3-28. User-Level Cache Instructions

Name	Mnemonic	Operand Syntax
Data Cache Block Flush	dcbf	rA,rB
Data Cache Block Set to Zero	dcbz	rA,rB
Data Cache Block Store	dcbst	rA,rB
Data Cache Block Touch	dcbt	rA,rB
Data Cache Block Touch for Store	dcbtst	rA,rB
Instruction Cache Block Invalidate	icbi	rA,rB

As with other memory-related instructions, the effect of the cache management instructions on memory are weakly ordered. If the programmer needs to ensure that cache or other instructions have been performed with respect to all other processors and system mechanisms, a **sync** instruction must be placed in the program following those instructions.

Note that when data address translation is disabled ($MSR[DR] = 0$), the Data Cache Block Set to Zero (**dcbz**) instruction allocates a cache block in the cache and may not verify that the physical address is valid. If a cache block is created for an invalid physical address, a machine check condition may result when an attempt is made to write that cache block back to memory. The cache block could be written back as a result of the execution of an instruction that causes a cache miss and the invalid addressed cache block is the target for replacement or a Data Cache Block Store (**dcbst**) instruction.

Note that any cache control instruction that generates an effective address that corresponds to a direct-store segment ($SR[T] = 1$) is treated as a no-op.

Table 3-28 lists the cache instructions that are accessible to user-level programs.

Note that incoherency may occur if a write-through store is followed by a **dcbz** instruction that is, in turn, followed by a snoop, all to the same cache block. This occurs when the logical address for the **dcbz** and the write-through store are different, but aliased to the same physical page.

To avoid potential adverse effects, **dcbz** should not address write-through memory that can be accessed through multiple logical addresses. Explicit store instructions that write all zeros should be used instead.

Note that broadcasting a sequence of **dcbz** instructions may cause snoop accesses to be retried indefinitely, which may cause the snoop originator to time out or the snooped transaction to not complete. This can be avoided by disabling the broadcasting of **dcbz** by marking the memory space being addressed by the **dcbz** instruction as not global in the BAT or PTE.

Note that incoherency may occur if the following sequence of accesses hits the same cache block: a write-through, a **dcbz** instruction, a snoop. This occurs when the logical address for the **dcbz** and the write-through store are different, but aliased to the same physical page.

3.2.5.4 External Control Instructions

The **eciwx** instruction provides an alternative way to map special devices. The MMU translation of the EA is not used to select the special device, as it is used in loads and stores. Rather, it is used as an address operand that is passed to the device over the address bus. Four other signals (the burst and size signals on the 60x bus) are used to select the device; these four signals output the 4-bit resource ID (RID) field in the EAR register. The **eciwx** instruction also loads a word from the data bus that is output by the special device. Executing these instructions when $MSR[DR] = 0$ causes a programming error, and the physical address on the bus is undefined. Executing these instructions to a direct-store segment causes a DSI exception. The external control instructions are listed in Table 3-29.

Table 3-29. External Control Instructions

Name	Mnemonic	Operand Syntax
External Control In Word Indexed	eciwx	rD,rA,rB
External Control Out Word Indexed	ecowx	rS,rA,rB

3.2.6 PowerPC OEA Instructions

The OEA includes the structure of the memory management model, supervisor-level registers, and exception model.

3.2.6.1 System Linkage Instructions

This section describes the system linkage instructions (see Table 3-30). The **sc** instruction is a user-level instruction that permits a user program to call on the system to perform a service and causes the processor to take an exception. The Return from Interrupt (**rfi**) instruction is a supervisor-level instruction that is useful for returning from an exception handler.

The Return from Critical Interrupt (**rfci**) instruction is a supervisor-level instruction that is only implemented in the G2_LE processor core. The **rfci** instruction is useful for returning from a critical interrupt exception handler. This new instruction is described in Section 3.2.8, “Implementation-Specific Instructions.”

Table 3-30. System Linkage Instructions

Name	Mnemonic	Operand Syntax
Return from Interrupt	rfi	—
System Call	sc	—

3.2.6.2 Processor Control Instructions—OEA

Processor control instructions are used to read from and write to the condition register (CR), machine state register (MSR), and special-purpose registers (SPRs), and to read from the time base register (TBU or TBL).

3.2.6.2.1 Move To/From Machine State Register Instructions

Table 3-31 lists the instructions provided by the core for reading from or writing to the MSR.

Table 3-31. Move To/From Machine State Register Instructions

Name	Mnemonic	Operand Syntax
Move from Machine State Register	mfmsr	rD
Move to Machine State Register	mtmsr	rS

3.2.6.2.2 Move To/From Special-Purpose Register Instructions

Simplified mnemonics are provided for the **mtspr** and **mfspir** instructions so they can be coded with the SPR name as part of the mnemonic rather than as a numeric operand. See Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, for simplified mnemonic examples. The **mtspr** and **mfspir** instructions are shown in Table 3-32.

Table 3-32. Move To/From Special-Purpose Register Instructions

Name	Mnemonic	Operand Syntax
Move from Special-Purpose Register	mf spr	rD,SPR
Move to Special-Purpose Register	mt spr	SPR,rS

For **mt spr** and **mf spr** instructions, the SPR number coded in assembly language does not appear directly as a 10-bit binary number in the instruction. The number coded is split into two 5-bit halves that are reversed in the instruction encoding, with the high-order 5 bits appearing in bits 16–20 of the instruction encoding and the low-order 5 bits in bits 11–15.

If the SPR field contains any value other than one of the values shown in Table 3-33, either the program exception handler is invoked or the results are boundedly undefined.

Table 3-33. Implementation-Specific SPR Encodings (mf spr)

SPR ¹			Register Name	Access
Decimal	spr[5–9]	spr[0–4]		
58	00001	11010	CSRR0 ²	Supervisor
59	00001	11011	CSRR1 ²	Supervisor
276	01000	10100	SPRG4 ²	Supervisor
277	01000	10101	SPRG5 ²	Supervisor
278	01000	10110	SPRG6 ²	Supervisor
279	01000	10111	SPRG7 ²	Supervisor
286	01000	11110	SVR ²	Supervisor
309	01001	10101	IBCR ²	Supervisor
310	01001	10110	DBCR ²	Supervisor
311	01001	10111	MBAR ²	Supervisor
317	01001	11101	DABR2 ²	Supervisor
560	10001	10000	IBAT4U ²	Supervisor
561	10001	10001	IBAT4L ²	Supervisor
562	10001	10010	IBAT5U ²	Supervisor
563	10001	10011	IBAT5L ²	Supervisor
564	10001	10100	IBAT6U ²	Supervisor
565	10001	10101	IBAT6L ²	Supervisor
566	10001	10110	IBAT7U ²	Supervisor
567	10001	10111	IBAT7L ²	Supervisor
568	10001	11000	DBAT4U ²	Supervisor
569	10001	11001	DBAT4L ²	Supervisor
570	10001	11010	DBAT5U ²	Supervisor
571	10001	11011	DBAT5L ²	Supervisor

Table 3-33. Implementation-Specific SPR Encodings (mfspr) (continued)

SPR ¹			Register Name	Access
Decimal	spr[5–9]	spr[0–4]		
572	10001	11100	DBAT6U ²	Supervisor
573	10001	11101	DBAT6L	Supervisor
574	10001	11110	DBAT7U ²	Supervisor
575	10001	11111	DBAT7L ²	Supervisor
976	11110	10000	DMISS	Supervisor
977	11110	10001	DCMP	Supervisor
978	11110	10010	HASH1	Supervisor
979	11110	10011	HASH2	Supervisor
980	11110	10100	IMISS	Supervisor
981	11110	10101	ICMP	Supervisor
982	11110	10110	RPA	Supervisor
1008	11111	10000	HID0	Supervisor
1009	11111	10001	HID1	Supervisor
1010	11111	10010	IABR	Supervisor
1011	11111	10011	HID2	Supervisor
1013	11111	10101	DABR ²	Supervisor
1018	11111	11010	IABR2 ²	Supervisor

¹ Note that the order of the two 5-bit halves of the SPR number is reversed compared with actual instruction coding.

For **mtspr** and **mfspr** instructions, the SPR number coded in assembly language does not appear directly as a 10-bit binary number in the instruction. The number coded is split into two 5-bit halves that are reversed in the instruction, with the high-order 5 bits appearing in bits 16–20 of the instruction and the low-order 5 bits in bits 11–15.

² These registers are implementation-specific for G2_LE core only.

Implementation Note—The core ignores the extended opcode differences between **mftb** and **mfspr** by ignoring TB[25] and treating both instructions identically.

3.2.6.3 Memory Control Instructions—OEA

This section describes memory control instructions, which include the following types:

- Cache management instructions
- Segment register manipulation instructions
- Translation lookaside buffer management instructions

3.2.6.3.1 Supervisor-Level Cache Management Instruction

The supervisor-level cache management instruction in the PowerPC architecture, **dcbi**, should not be used on the G2 core. If it is used it can cause a data storage interrupt. The user-level **dcbf** instruction, described in Section 3.2.5.3, “Memory Control Instructions—VEA” and Section 4.8, “Cache Control Instructions,” should be used when the program needs to invalidate cache blocks. Note that the **dcbf** instruction causes modified blocks to be flushed to system memory if they are the target of a **dcbf** instruction, whereas, by definition in the PowerPC architecture, the **dcbi** instruction only invalidates modified blocks.

3.2.6.3.2 Segment Register Manipulation Instructions

The instructions listed in Table 3-34 provide access to the segment registers for the G2 core. These instructions operate completely independent of the MSR[IR] and MSR[DR] bit settings. Refer to “Synchronization Requirements for Special Registers and TLBs” in Chapter 2, “Register Set,” in the *Programming Environments Manual*, for serialization requirements and other recommended precautions to observe when manipulating the segment registers.

Table 3-34. Segment Register Manipulation Instructions

Name	Mnemonic	Operand Syntax
Move from Segment Register	mfsr	rD,SR
Move from Segment Register Indirect	mfsrin	rD,rB
Move to Segment Register	mtsr	SR,rS
Move to Segment Register Indirect	mtsrin	rS,rB

3.2.6.3.3 Translation Lookaside Buffer Management Instructions

The address translation mechanism is defined in terms of segment descriptors and page table entries (PTEs) used by the processors to locate the effective-to-physical address mapping for a particular access. The PTEs reside in page tables in memory. As defined for 32-bit implementations by the PowerPC architecture, segment descriptors reside in 16 on-chip segment registers.

Implementation Note—The G2 core provides the ability to invalidate a TLB entry. The TLB Invalidate Entry (**tlbie**) instruction invalidates the TLB entry indexed by the EA, and operates on both the instruction and data TLBs simultaneously invalidating four TLB entries (both sets in each TLB). The index corresponds to bits 15–19 of the EA. To invalidate all entries within both TLBs, 32 **tlbie** instructions should be issued, incrementing this field by one each time.

The core provides two implementation-specific instructions (**tlbld** and **tlbli**) that are used by software table search operations following TLB misses to load TLB entries on-chip.

For more information on **tlbld** and **tlbli** refer to Section 3.2.8, “Implementation-Specific Instructions.”

Note that the **tlbia** instruction is not implemented on the core.

Refer to Chapter 6, “Memory Management,” for more information about the TLB operations for the G2 core. Table 3-35 lists the TLB instructions.

Table 3-35. Translation Lookaside Buffer Management Instructions

Name	Mnemonic	Operand Syntax
Load Data TLB Entry	tlbld	rB
Load Instruction TLB Entry	tlbli	rB
TLB Invalidate Entry	tlbie	rB
TLB Synchronize	tlbsync	—

Because the presence and exact semantics of the translation lookaside buffer management instructions is implementation-dependent, system software should incorporate uses of the instructions into subroutines to maximize compatibility with programs written for other processors.

For more information on the PowerPC instruction set, refer to Chapter 4, “Addressing Modes and Instruction Set Summary,” and Chapter 8, “Instruction Set,” in the *Programming Environments Manual*.

3.2.7 Recommended Simplified Mnemonics

To simplify assembly language programs, a set of simplified mnemonics is provided for some of the most frequently used operations (such as no-op, load immediate, load address, move register, and complement register). PowerPC compliant assemblers provide the simplified mnemonics listed in “Recommended Simplified Mnemonics” in Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*, and listed with some of the instruction descriptions in this chapter. Programs written to be portable across the various assemblers for the PowerPC architecture should not assume the existence of mnemonics not described in this document.

For a complete list of simplified mnemonics, see Appendix F, “Simplified Mnemonics,” in the *Programming Environments Manual*.

3.2.8 Implementation-Specific Instructions

This section provides a detailed look at the two G2 and one G2_LE core implementation-specific instructions—**tlbld**, **tlbli**, and **rfdi**, respectively.

tlbld

Load Data TLB Entry

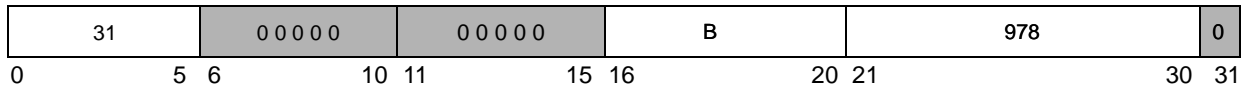
tlbld

Integer Unit

tlbld

rB

Reserved



$EA \leftarrow (rB)$

TLB entry created from DCMP and RPA

DTLB entry selected by $EA[15-19]$ and $SRR1[WAY] \leftarrow$ created TLB entry

The EA is the contents of **rB**. The **tlbld** instruction loads the contents of the data PTE compare (DCMP) and required physical address (RPA) registers into the first word of the selected data TLB entry. The specific DTLB entry to be loaded is selected by the EA and $SRR1[WAY]$ bit.

The **tlbld** instruction should only be executed when address translation is disabled ($MSR[IR] = 0$ and $MSR[DR] = 0$).

Note that it is possible to execute the **tlbld** instruction when address translation is enabled; however, extreme caution should be used in doing so. If data address translation is set ($MSR[DR] = 1$) **tlbld** must be preceded by a **sync** instruction and succeeded by a context synchronizing instruction.

Also, note that care should be taken to avoid modification of the instruction TLB entries that translate current instruction prefetch addresses.

This is a supervisor-level instruction; it is also a G2 core-specific instruction, and not part of the PowerPC instruction set.

Other registers altered:

- None

tlbli

Load Instruction TLB Entry

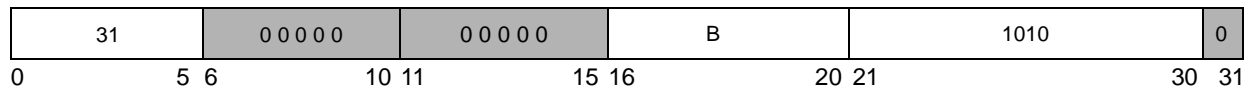
tlbli

Integer Unit

tlbld

rB

 Reserved



$EA \leftarrow (rB)$

TLB entry created from ICMP and RPA

ITLB entry selected by $EA[15-19]$ and $SRR1[WAY] \leftarrow$ created TLB entry

The EA is the contents of **rB**. The **tlbli** instruction loads the contents of the instruction PTE compare (ICMP) and required physical address (RPA) registers into the first word of the selected instruction TLB entry. The specific ITLB entry to be loaded is selected by the EA and $SRR1[WAY]$ bit.

The **tlbli** instruction should only be executed when address translation is disabled ($MSR[IR] = 0$ and $MSR[DR] = 0$).

Note that it is possible to execute the **tlbl**d instruction when address translation is enabled; however, extreme caution should be used in doing so. If instruction address translation is set ($MSR[IR] = 1$), **tlbli** must be followed by a context synchronizing instruction such as **isync** or **rfi**.

Also, note that care should be taken to avoid modification of the instruction TLB entries that translate current instruction prefetch addresses.

This is a supervisor-level instruction; it is also a G2 core-specific instruction, and not part of the PowerPC instruction set.


Other registers altered:

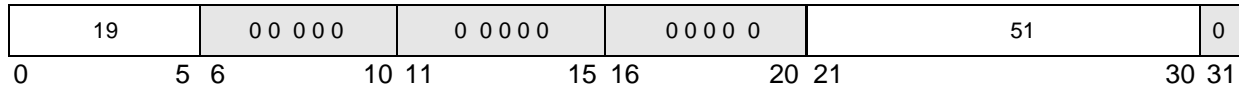
- None

rftci

Return from Critical Interrupt

rftci

 Reserved



MSR[16-27, 30-31] ← CSRR1[16-27, 30-31]
 NIA ← iea CSRR0[0-29] || 0b00

Bits CSRR1[16-27, 30-31] are placed into the corresponding bits of the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address CSRR0[0-29] || 0b00. If the new MSR value enables one or more pending exceptions, the exception associated with the highest priority pending exception is generated; in this case the value placed into CSRR0 by the exception processing mechanism is the address of the instruction that would have been executed next had the exception not occurred. Note that an implementation may define additional MSR bits, and in this case, may also cause them to be saved to CSRR1 from MSR on an exception and restored to MSR from CSRR1 on an **rftci**.

This is a supervisor-level, context synchronizing instruction. This instruction is defined only for 32-bit implementations.

Other registers altered:

- MSR

Chapter 4

Instruction and Data Cache Operation

The G2 core provides two 16-Kbyte, four-way set-associative caches to allow the registers and execution units rapid access to instructions and data. Both the instruction and data caches are tightly coupled to the G2 core bus interface unit (BIU) to allow efficient access to the system memory controller and other bus masters. The G2 core load/store unit (LSU) is also directly coupled to the data cache to allow the efficient movement of data to and from the general-purpose and floating-point registers.

This chapter describes the organization of the cache, cache coherency protocols, cache control instructions, and various cache operations. It describes the interaction between the caches, the load/store unit, the instruction unit, and the memory subsystem. It also describes the cache way-locking features provided in the G2 core.

Note that in this chapter, the term multiprocessor is used in the context of maintaining cache coherency. These multiprocessor devices could be actual processors or other devices that can access system memory, maintain their own caches, and function as bus masters requiring cache coherency.

4.1 Overview

Both the instruction and data caches have 32-byte blocks, and data cache blocks can be snooped or cast out when the cache block is reloaded. The data cache is designed to adhere to a write-back policy, but the G2 core allows control of cacheability, write-back policy, and memory coherency at the page and block level. Both caches use a least recently used (LRU) replacement policy. Burst fill operations to the caches result from cache misses, or in the case of the data cache, cache block write-back operations to memory. Note that in the PowerPC architecture, the term ‘cache block,’ or simply ‘block’ when used in the context of cache implementations, refers to the unit of memory at which coherency is maintained. For the G2 core, the block size is equivalent to the eight-word cache line. This value may be different for other implementations that support the PowerPC architecture.

The data cache is configured as 128 sets of 4 blocks. Each block consists of 32 bytes, 2 state bits, and an address tag. The two state bits implement the three-state MEI (modified/exclusive/invalid) protocol, a coherent subset of the standard four-state MESI protocol. Cache coherency is enforced by on-chip bus snooping logic. Since the G2 core data cache

tags are single-ported, a simultaneous load or store and snoop access represent a resource contention. The snoop access is given first access to the tags. Load or store operations can be performed to the cache on the clock cycle immediately following a snoop access if the snoop misses. Snoop hits may block the data cache for two or more cycles, depending on whether a copy back to main memory is required.

The instruction cache also consists of 128 sets of 4 blocks, and each block consists of 32 bytes, an address tag, and a valid bit. The instruction cache is only written as a result of a block fill operation on a cache miss. In the G2 core, the instruction cache is blocked only until the critical load completes. The G2 core supports instruction fetching from other instruction cache lines following the forwarding of the critical-first-double-word of a cache line load operation. Successive instruction fetches from the cache line being loaded are forwarded, and accesses to other instruction cache lines can proceed during the cache line load operation. The instruction cache is not snooped, and cache coherency must be maintained by software. A fast hardware invalidation capability is provided to support cache maintenance.

The load/store unit provides the data transfer interface between the data cache and the GPRs and FPRs. The LSU provides all logic required to calculate effective addresses, handle data alignment to and from the data cache, and provides sequencing for load and store string and multiple operations. As shown in Figure 1-1, the caches provide a 64-bit interface to the instruction fetcher and LSU. Write operations to the data cache can be performed on a byte, half-word, word, or double-word basis.

The G2 core bus interface unit receives requests for bus operations from the instruction and data caches, and executes the operations according to the G2 core bus protocol. The BIU provides address queues, prioritization, and bus control logic. The BIU also captures snoop addresses for data cache, address queue, and memory reservation (**lwarx** and **stwcx**, instruction) operations. The BIU also contains a touch load address buffer used for address compares during load or store operations. All the data for the corresponding address queues (load and store data queues) is located in the data cache. The data queues are considered temporary storage for the cache and not part of the BIU.

On a cache miss, the G2 core cache blocks are loaded in four beats of 64 bits each when the G2 core is configured with a 64-bit data bus; when the G2 core is configured with a 32-bit bus, cache block loads are performed with eight beats of 32 bits each. The burst load is performed as critical-double-word-first. The data cache is blocked to internal accesses until the load completes; the instruction cache allows sequential fetching during a cache block load. In the G2 core, the critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays. Note that the cache being filled cannot be accessed internally until the fill completes.

When address translation is enabled, the memory access is performed under the control of the page table entry used to translate the effective address. Each page table entry and BAT contains four mode control bits, W, I, M, and G, that specify the storage mode for all

accesses translated using that particular page table entry. The W (write-through) and I (caching-inhibited) bits control how the processor executing the access uses its own cache. The M (memory coherence) bit specifies whether the processor executing the access must use the MEI (modified, exclusive, or invalid) cache coherence protocol to ensure all copies of the addressed memory location are kept consistent. The G (guarded memory) bit controls whether out-of-order data and instruction fetching is permitted.

The G2 core maintains data cache coherency in hardware by coordinating activity between the data cache, memory system, and bus interface logic. As bus operations are performed on the bus by other bus masters, the G2 core bus snooping logic monitors the addresses that are referenced. These addresses are compared with the addresses resident in the data cache. If there is a snoop hit, the G2 core bus snooping logic responds to the bus interface with the appropriate snoop status (for example, a `core_artry_out`). Additional snoop action may be forwarded to the cache as a result of a snoop hit in some cases (a cache push of modified data or cache block invalidation).

The G2 core supports a fully-coherent 4-Gbyte physical memory address space. Bus snooping is used to drive the MEI three-state cache-coherency protocol that ensures the coherency of global memory with respect to the processor's cache. See Section 4.7.1, "MEI State Definitions."

This chapter describes the organization of the G2 core on-chip instruction and data caches, the MEI cache coherence protocol, cache control instructions, various cache operations, and the interaction between the cache, LSU, and BIU. G2 core specific information is noted where applicable.

4.2 Instruction Cache Organization and Control

The instruction fetcher accesses the instruction cache frequently in order to sustain the high throughput provided by the six-entry instruction queue.

4.2.1 Instruction Cache Organization

The instruction cache organization is shown in Figure 4-1. Each cache block contains eight contiguous words from memory that are loaded from an eight-word boundary (that is, bits A27–A31 of the effective addresses are zero); thus, a cache block never crosses a page boundary. Misaligned accesses across a page boundary can incur a performance penalty.

Note that address bits A20–A26 provide an index to select a set. Bits A27–A31 select a byte within a block. The tags consists of bits PA0–PA19. Address translation occurs in parallel, such that higher-order bits (the tag bits in the cache) are physical. Note that the replacement algorithm is strictly an LRU algorithm; that is, the least-recently used block is filled with new instructions on a cache miss.

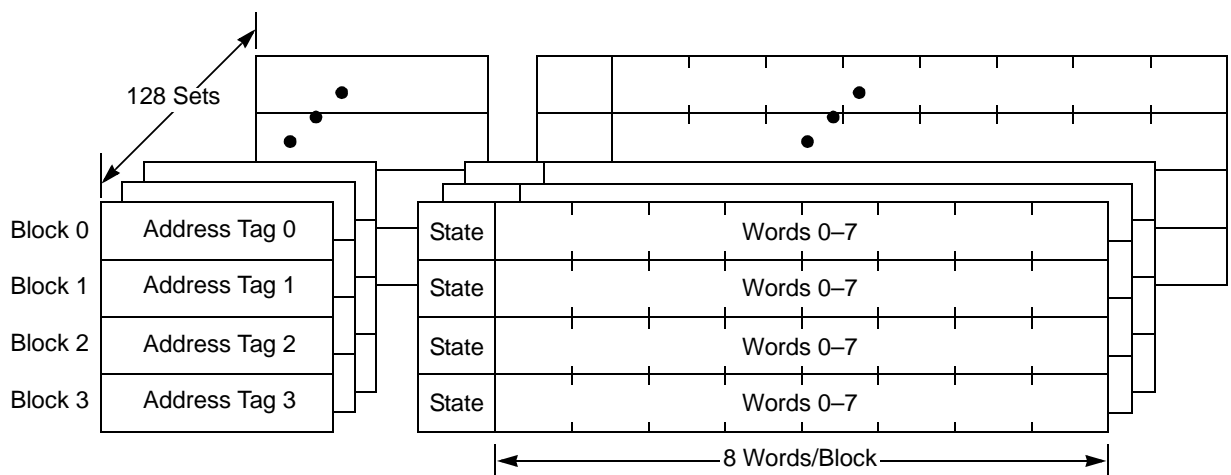


Figure 4-1. Instruction Cache Organization

4.2.2 Instruction Cache Fill Operations

The G2 core instruction cache blocks are loaded in four 64-bit beats, with the critical-double-word loaded first. The instruction cache allows sequential fetching during a cache block load. On a cache miss, the critical and following double words read from memory are simultaneously written to the instruction cache and forwarded to the dispatch queue, thus minimizing stalls due to cache fill latency. There is no snooping of the instruction cache. In the G2 core, the critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.

4.2.3 Instruction Cache Control

In addition to instruction cache control instructions, the G2 core provides several HID0 bits to control invalidating, disabling, and locking the instruction cache. The WIMG bits in the page tables and the IBATs also affect the cacheability of pages and whether the pages are considered guarded.

4.2.3.1 Instruction Cache Invalidation

Although the G2 core instruction cache is automatically invalidated during a power-on or hard reset, asserting `core_sreset` does not invalidate the instruction cache. Software can invalidate the contents of the instruction cache using the instruction cache flash invalidate control bit, `HID0[ICFI]`. Flash invalidation of the instruction cache is accomplished by setting ICFI bit (invalidates the cache) and subsequently clearing the ICFI bit (enables normal operation) in two consecutive `mtspr[HID0]` instructions.

4.2.3.2 Instruction Cache Disabling

The instruction cache may be disabled through the use of the instruction cache enable control bit, HID0[ICE]. When the instruction cache is in the disabled state, the cache tag state bits are ignored and all accesses are propagated to the bus as single-beat transactions. The ICE bit is cleared during a power-on reset, causing the instruction cache to be disabled. To prevent the cache from being enabled or disabled in the middle of a data access, an **isync** instruction should be issued before changing the value of ICE.

4.2.3.3 Instruction Cache Locking

The contents of instruction cache may be locked through the use of HID0[ILOCK]. A locked instruction cache supplies instructions normally on a cache hit, but cache misses are treated as cache-inhibited accesses. The cache-inhibited (`core_ci`) signal is asserted if a cache access misses into a locked cache. The setting of the ILOCK bit must be preceded by an **isync** instruction to prevent the instruction cache from being locked during an instruction access.

Note that the G2 core also provides instruction cache way-locking in addition to entire instruction cache locking as described in Section 4.12, “Cache Locking.”

4.3 Data Cache Organization and Control

The LSU transfers data between the data cache and the GPRs and FPRs and provides buffers for load and store bus operations. The data cache also provides storage for the cache tags required for memory coherency and performs the cache block replacement LRU function.

4.3.1 Data Cache Organization

The organization of the data cache is shown in Figure 4-2. Each cache block contains eight contiguous words from memory that are loaded from an eight-word boundary (that is, bits A27–A31 of the effective addresses are zero); thus, a cache block never crosses a page boundary. Misaligned accesses across a page boundary can incur a performance penalty.

Note that bits A20–A26 provide an index to select a set. Bits A27–A31 select a byte within a block. The tags consists of bits PA0–PA19. Address translation occurs in parallel, such that higher-order bits (the tag bits in the cache) are physical. Note that the replacement algorithm is strictly an LRU algorithm; that is, the least-recently used block is filled with new data on a cache miss.

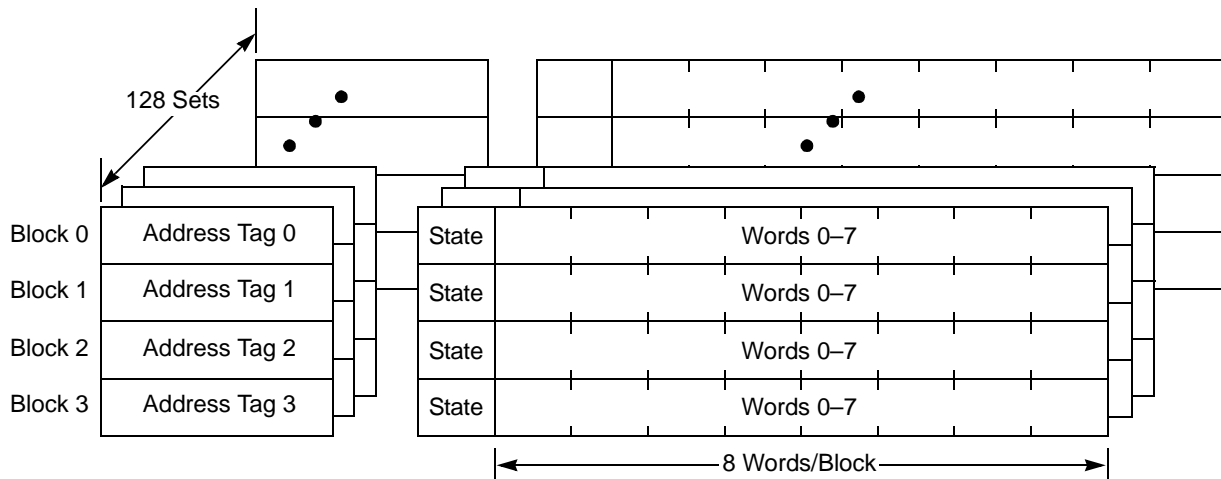


Figure 4-2. Data Cache Organization

4.3.2 Data Cache Fill Operations

When the G2 core is configured with a 64-bit data bus, cache blocks are loaded in four beats of 64 bits each. When the G2 core is configured with a 32-bit bus, cache block loads are performed with eight beats of 32 bits each. The burst load is performed as critical-double-word-first. The data cache is blocked to internal accesses until the load completes. In the G2 core, the critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.

4.3.3 Data Cache Control

The G2 core provides several means of data cache control through the use of the WIMG bits in the page tables, control bits in the HID0 register, and user- and supervisor-level cache control instructions. While memory page level cache control is provided by the WIMG bits, the on-chip data cache can be invalidated, disabled, locked, or broadcast by the control bits in the HID0 register described in this section. (Note that user- and supervisor-level are referred to as problem and privileged state, respectively, in the architecture specification.)

4.3.3.1 Data Cache Invalidation

While the data cache is automatically invalidated when the G2 core is powered up and during a hard reset, assertion of the soft reset signal does not cause data cache invalidation. Software may invalidate the contents of the data cache using the data cache flash invalidate (DCFI) control bit in the HID0 register. Flash invalidation of the data cache is accomplished by setting the DCFI bit (invalidates the cache) and subsequently clearing the DCFI bit (enables normal operation) in two consecutive store operations. If DCFI is not cleared the cache state will remain invalid.

4.3.3.2 Data Cache Disabling

The data cache may be disabled through the use of the data cache enable (DCE) control bit in the HID0 register. When the data cache is in the disabled state, the cache tag state bits are ignored, and all accesses are propagated to the bus as single-beat transactions. The DCE bit is cleared on power-up, causing the data cache to be disabled. To prevent the cache from being enabled or disabled in the middle of a data access, a **sync** instruction should be issued before changing the value of DCE.

Note that while snooping is not performed when the data cache is disabled, cache operations (caused by the **dcbz**, **dcbf**, **dcbst**, and **dcbi** instructions) are not affected by disabling the cache, causing potential coherency errors. An example of this would be a **dcbf** instruction that hits a modified cache block in the disabled cache, causing a copy back to memory of potentially stale data.

NOTE

The **dcbi** instruction should never be used on the G2 core.

Regardless of the state of HID0[DCE], load and store operations are assumed to be weakly ordered. Thus, the LSU can perform load operations that occur later in the program ahead of store operations, even when the data cache is disabled. However, strongly ordered load and store operations can be enforced through the setting of the I bit (of the page WIMG bits) when address translation is enabled. Note that when address translation is disabled, the default WIMG bits cause the I bit to be cleared (accesses are assumed to be cacheable), and thus, the accesses are weakly ordered. Refer to Section 4.6.2, “Caching-Inhibited Attribute (I),” for a description of the operation of the I bit and Section 6.2, “Real Addressing Mode,” for a description of the WIMG bits when address translation is disabled.

4.3.3.3 Data Cache Locking

The contents of the data cache may be locked through the HID0[DLOCK]. A locked data cache supplies data normally on a cache hit, but cache misses are treated as cache-inhibited accesses. The cache-inhibited ($\overline{\text{core_ci}}$) signal is asserted if a cache access misses into a locked cache. The setting of DLOCK must be preceded by a **sync** instruction to prevent the cache from being locked during an access.

Note that the G2 core also provides instruction cache way-locking in addition to entire data cache locking as described in Section 4.12, “Cache Locking.”

4.3.3.4 Data Cache Operations and Address Broadcasts

Executing a **dcbz** instruction generates an address-only broadcast on the bus. Additionally, if HID0[ABE] is set on a G2 core processor, the execution of the **dcbf**, **dcbi**, and **dcbst** instructions also causes an address-only broadcast. The ability of the G2 core to optionally perform address-only broadcasts when executing the **dcbi**, **dcbf**, and **dcbst** instructions

allows the coherency management of an external copy-back L2 cache. Note that these cache control instruction broadcasts are not snooped by the G2 core.

4.3.4 Data Cache Touch Load Support

Touch load operations allow an instruction stream to prefetch data from memory prior to a cache miss. The G2 core supports touch load operations through a temporary cache block buffer located between the BIU and the data cache. The cache block buffer is essentially a floating cache block that is loaded by the BIU on a touch load operation, and is then read by a load instruction that requests that data. After a touch load completes on the bus, the BIU continues to compare the touch load address with subsequent load requests from the data cache. If the load address matches the touch load address in the BIU, the data is forwarded to the data cache from the touch load buffer, the read from memory is canceled, and the touch load address buffer is invalidated.

To avoid the storage of stale data in the touch load buffer, touch load requests that are mapped as write-through or caching-inhibited by the MMU are treated as no-ops by the BIU. Also, subsequent load instructions after a touch load that are mapped as write-through or caching-inhibited do not hit in the touch load buffer, and cause the touch load buffer to be invalidated on a matching address.

While the G2 core provides only a single cache block buffer, other microprocessor implementations may provide buffering for more than one cache block. Programs written for other implementations may issue several **dcbt** or **dcbtst** instructions sequentially, reducing the performance if executed on the G2 core. To improve performance in these situations, HID0[NOOPTI] (bit 31) can be set. This causes the **dcbt** and **dcbtst** instructions to be treated as no-ops, cause no bus activity, and incur only one processor clock cycle of execution latency. NOOPTI is cleared at a power-on reset, enabling the use of the **dcbt** and **dcbtst** instructions.

4.4 Basic Data Cache Operations

This section describes the three types of operations that can occur to the data cache, and how these operations are implemented in the G2 core.

4.4.1 Data Cache Fill

A cache block is filled after a read miss or write miss (read-with-intent-to-modify) occurs in the cache. The cache block that corresponds to the missed address is updated by a burst transfer of the data from system memory. Note that if a read miss occurs in a system with multiple bus masters, and the data is modified in another cache, the modified data is first written to external memory before the cache fill occurs.

4.4.2 Data Cache Cast-Out Operation

The G2 core uses an LRU replacement algorithm to determine which of the four possible cache locations should be used for a cache update on a cache miss. Adding a new block to the cache causes any modified data associated with the least-recently used element to be written back, or cast out, to system memory to maintain memory coherence.

4.4.3 Cache Block Push Operation

When a cache block in the G2 core is snooped and hit by another bus master and the data is modified, the cache block must be written to memory and made available to the snooping device. The cache block that is hit, is pushed out onto the bus. The G2 core supports two kinds of push operations—normal push operations and enveloped high-priority push operations, described in Section 4.7.9, “Enveloped High-Priority Cache Block Push Operation.”

4.5 Data Cache Transactions on Bus

The G2 core transfers data to and from the data cache in single-beat transactions of two words, or in four-beat transactions of eight words which fill a cache block.

4.5.1 Single-Beat Transactions

Single-beat bus transactions can transfer from 1 to 8 bytes to or from the G2 core. Single-beat transactions can be caused by cache write-through accesses, caching-inhibited accesses (I bit of the WIMG bits for the page is set), or accesses when the cache is disabled (HID0[DCE] bit is cleared), and can be misaligned.

4.5.2 Burst Transactions

Burst transactions on the G2 core always transfer eight words of data at a time, and are aligned to a double-word boundary. The G2 core transfer burst (`core_tbst`) output signal indicates to the system whether the current transaction is a single-beat transaction or four-beat burst transfer. Burst transactions have an assumed address order. For cacheable read operations or cacheable, non-write-through write operations that miss the cache, the G2 core presents the double-word aligned address associated with the load or store instruction that initiated the transaction.

As shown in Figure 4-3, this quad word contains the address of the load or store that missed the cache. This minimizes latency by allowing the critical code or data to be forwarded to the processor before the rest of the block is filled. For all other burst operations, however, the entire block is transferred in order (oct-word aligned). Critical-double-word-first fetching on a cache miss applies to both the data and instruction cache.

G2 Core Cache Address
Bits 27:28

0 0	0 1	1 0	1 1
A	B	C	D

If the address requested is in double-word A, the address placed on the bus is that of double-word A, and the four data beats are ordered in the following manner:

Beat				
	0	1	2	3
	A	B	C	D

If the address requested is in double-word C, the address placed on the bus will be that of double-word C, and the four data beats are ordered in the following manner:

Beat				
	0	1	2	3
	C	D	A	B

Figure 4-3. Double-Word Address Ordering—Critical-Double-Word-First

4.5.3 Access to Direct-Store Segments

The G2 core does not provide support for access to direct-store segments. Operations attempting to access a direct-store segment will invoke a DSI exception. See Section 5.5.3, “DSI Exception (0x00300).”

4.6 Memory Management/Cache Access Mode Bits—W, I, M, and G

Some memory characteristics can be set on either a block or page basis by using the WIMG bits in the BAT registers or page table entry (PTE), respectively. The WIMG attributes control the following functionality:

- Write-through (W bit)
- Caching-inhibited (I bit)
- Memory coherency (M bit)
- Guarded memory (G bit)

These bits allow both uniprocessor and multiprocessor system designs to exploit numerous system-level performance optimizations.

Careless specification and use of these bits may create situations where coherency paradoxes are observed by the processor. In particular, this can happen when the state of these bits is changed without appropriate precautions being taken (for example, when

flushing the pages that correspond to the changed bits from the caches of all processors in the system is required, or when the address translations of aliased physical addresses (referred to as real addresses in the architecture specification) specify different values for any of the WIMG bits). The G2 core considers any of these cases to be a programming error that may compromise the coherency of memory. These paradoxes can occur within a single processor or across several devices, as described in Section 4.7.4.1, “Coherency in Single-Processor Systems.”

The WIMG attributes are programmed by the operating system for each page and block. The W and I attributes control how the processor performing an access uses its own cache. The M attribute ensures that coherency is maintained for all copies of the addressed memory location. The G attribute prevents out-of-order loading and prefetching from the addressed memory location.

When an access requires coherency, the processor performing the access must inform the coherency mechanisms throughout the system that the access requires memory coherency. The M attribute determines the kind of access performed on the bus (global or local).

The WIMG attributes occupy 4 bits in the BAT registers for block address translation and in the PTEs for page address translation. The WIMG bits are programmed as follows:

- The operating system uses the **mtspr** instruction to program the WIMG bits in the BAT registers for block address translation. The IBAT register pairs do not have a G bit and all accesses that use the IBAT register pairs are considered not guarded.
- The operating system writes the WIMG bits for each page into the PTEs in system memory as it sets up the page tables.

Note that for accesses performed with direct address translation ($MSR[IR] = 0$ or $MSR[DR] = 0$ for instruction or data access, respectively), the WIMG bits are automatically generated as 0b0011 (the data is write-back, caching is enabled, memory coherency is enforced, and memory is guarded).

4.6.1 Write-Through Attribute (W)

When an access is designated as write-through ($W = 1$), if the data is in the cache, a store operation updates the cached copy of the data. In addition, the update is written to the external memory location (as described below).

While the PowerPC architecture permits multiple store instructions to be combined for write-through accesses except when the store instructions are separated by a **sync** or **eieio** instruction, the G2 core does not implement this ‘combined store’ capability. Note that a store operation that uses the write-through attribute may cause any part of valid data in the cache to be written back to main memory.

The definition of the external memory location to be written to, in addition to the on-chip cache, depends on the implementation of the memory system and can be illustrated by the following examples:

- RAM—The store is sent to the RAM controller to be written into the target RAM.
- I/O device—The store is sent to the memory-mapped I/O control hardware to be written to the target register or memory location.

In systems with multilevel caching, the store must be written to at least a depth in the memory hierarchy that is seen by all processors and devices.

Accesses that correspond to $W = 0$ are considered write-back. For this case, although the store operation is performed to the cache, it is only made to external memory when a copy-back operation is required. Use of the write-back mode ($W = 0$) can improve overall performance for areas of the memory space that are seldom referenced by other masters in the system.

4.6.2 Caching-Inhibited Attribute (I)

If $I = 1$, the memory access is completed by referencing the location in main memory, bypassing the on-chip cache. During the access, the addressed location is not loaded into the cache nor is the location allocated in the cache. It is considered a programming error if a copy of the target location of an access to caching-inhibited memory is resident in the cache. Software must ensure that the location has not been previously loaded into the cache, or, if it has, that it has been flushed from the cache.

The PowerPC architecture permits data accesses from more than one instruction to be combined for cache-inhibited operations, except when the accesses are separated by a **sync** instruction, or by an **ei** instruction when the page or block is also designated as guarded. This ‘combined access’ capability is not implemented on the G2 core. Note that the **ei** is treated as a no-op by the G2 core.

The caching-inhibited (I) bit in the G2 core controls whether load and store operations are strongly or weakly ordered. If an I/O device requires load and store accesses to occur in program order, then the I bit for the page must be set.

4.6.3 Memory Coherency Attribute (M)

This attribute is provided to allow improved performance in systems where hardware-enforced coherency is relatively slow, and software is able to enforce the required coherency. When $M = 0$, the processor does not enforce data coherency. When $M = 1$, the processor enforces data coherency and the corresponding access is considered to be a global access.

When the M attribute is set, and the access is performed, the global signal is asserted to indicate that the access is global. Snooping devices affected by the access must then

respond to this global access if their data is modified by asserting `core_artry_in`, and updating the memory location.

Because instruction memory does not have to be consistent with data memory, the G2 core ignores the M attribute for instruction accesses.

4.6.4 Guarded Attribute (G)

When the guarded bit is set, the memory area (block or page) is designated as guarded, meaning that the processor will perform out-of-order accesses to this area of memory, only as follows:

- Out-of-order load operations from guarded memory areas are performed only if the corresponding data is resident in the cache.
- The processor prefetches from guarded areas, but only when required, and only within the memory boundary dictated by the cache block. That is, if an instruction is certain to be required for execution by the program, it is fetched and the remaining instructions in the block may be prefetched, even if the area is guarded.

This setting can be used to protect certain memory areas from read accesses made by the processor that are not dictated directly by the program. If there are areas of memory that are not fully populated (in other words, there are holes in the memory map within this area), this setting can protect the system from undesired accesses caused by out-of-order load operations or instruction prefetches that could lead to the generation of the machine check exception. Also, the guarded bit can be used to prevent out-of-order load operations or prefetches from occurring to certain peripheral devices that produce undesired results when accessed in this way.

4.6.5 W, I, and M Bit Combinations

Table 4-1 summarizes the six combinations of the WIM bits. Note that either a zero or one setting for the G bit is allowed for each of these WIM bit combinations.

Table 4-1. Combinations of W, I, and M Bits

WIM Setting	Meaning
000	Data may be cached. Loads or stores whose target hits in the cache use that entry in the cache. Memory coherency is not enforced by hardware.
001	Data may be cached. Loads or stores whose target hits in the cache use that entry in the cache. Memory coherency is enforced by hardware.
010	Caching is inhibited. The access is performed to external memory, completely bypassing the cache. Memory coherency is not enforced by hardware.

Table 4-1. Combinations of W, I, and M Bits (continued)

WIM Setting	Meaning
011	Caching is inhibited. The access is performed to external memory, completely bypassing the cache. Memory coherency must be enforced by external hardware (processor provides hardware indication that access is global).
100	Data may be cached. Load operations whose target hits in the cache use that entry in the cache. Stores are written to external memory. The target location of the store may be cached and is updated on a hit. Memory coherency is not enforced by hardware.
101	Data may be cached. Load operations whose target hits in the cache use that entry in the cache. Stores are written to external memory. The target location of the store may be cached and is updated on a hit. Memory coherency is enforced by hardware.

4.6.5.1 Out-of-Order Execution and Guarded Memory

Out-of-order execution occurs when the G2 core performs operations in advance in case the result is needed. Typically, these operations are performed by otherwise idle resources; thus if a result is not required, it is ignored and the out-of-order operation incurs no time penalty (typically).

Supervisor-level programs designate memory as guarded on a block or page level. Memory is designated as guarded if it is not be well-behaved with respect to out-of-order operations.

For example, the memory area that contains a memory-mapped I/O device may be designated as guarded if an out-of-order load or instruction fetch performed to such a device might cause the device to perform unexpected or incorrect operations. Another example of memory that should be designated as guarded is the area that corresponds to the device that resides at the highest implemented physical address (as it has no successor and out-of-order sequential operations such as instruction prefetching may result in a machine check exception). In addition, areas that contain holes in the physical memory space may be designated as guarded.

4.6.5.2 Effects of Out-of-Order Data Accesses

Most data operations may be performed out-of-order, as long as the machine appears to follow a simple sequential model. However, the following out-of-order operations do not occur:

- Out-of-order loading from guarded memory ($G = 1$) does not occur. However, when a load or store operation is required by the program, the entire cache block(s) containing the referenced data may be loaded into the cache.
- Out-of-order store operations that alter the state of the target location do not occur.

- No errors except machine check exceptions are reported due to the out-of-order execution of an instruction until it is known that execution of the instruction is required.

Machine check exceptions resulting solely from out-of-order execution (from nonguarded memory) may be reported. When an out-of-order instruction result is abandoned, only one side effect (other than a possible machine check) may occur—the referenced bit (R) in the corresponding page table entry (and TLB entry) can be set due to an out-of-order load operation. See Chapter 5, “Exceptions,” for more information on the machine check exception.

Thus, an out-of-order load or store instruction will not access guarded memory unless one of the following conditions exist:

- The target memory item is resident in an on-chip cache. In this case, the location may be accessed from the cache or main memory.
- The target memory item is cacheable ($I = 0$) and it is guaranteed that the load or store is in the execution path (assuming there are no intervening exceptions). In this case, the entire cache block containing the target may be loaded into the cache.

4.6.5.3 Effects of Out-of-Order Instruction Fetches

To avoid instruction fetch delay, the processor typically fetches instructions ahead of those currently being executed. Such instruction prefetching is said to be out-of-order in that prefetched instructions may not be executed due to intervening branches or exceptions.

During instruction prefetching, no errors except machine check exceptions are reported due to the out-of-order fetching of an instruction until it is known that execution of the instruction is required.

Machine check exceptions resulting solely from out-of-order execution (from nonguarded memory) may be reported. When an out-of-order instruction result is abandoned, only one side effect (other than a possible machine check) may occur—the referenced bit (R) in the corresponding page table entry (and TLB entry) can be set due to an out-of-order load operation. See Chapter 5, “Exceptions,” for more information on the machine check exception.

Instruction fetching from guarded memory is not permitted.

4.7 Cache Coherency—MEI Protocol

The primary objective of a coherent memory system is to provide the same image of memory to all devices using the system. Coherency allows synchronization and cooperative use of shared resources. Otherwise, multiple copies of a memory location, some containing stale values, could exist in a system resulting in errors when the stale values are used. Each potential bus master must follow rules for managing the state of its cache.

The G2 core cache coherency protocol is a coherent subset of the standard MESI four-state cache protocol that omits the shared state. Since data cannot be shared, the G2 core signals all cache block fills as if they were write misses (read-with-intent-to-modify), flushing the corresponding copies of the data in all caches external to the G2 core prior to the G2 core cache block fill operation. Following the cache block load, the G2 core is the exclusive owner of the data and may write to it without a bus broadcast transaction.

To maintain this coherency, all global reads observed on the bus by the G2 core are snooped as if they are writes, causing the G2 core to write a modified cache block back to memory and invalidate the cache block, or simply invalidate the cache block if it is unmodified. The exception to this rule occurs when a snooped transaction is a caching-inhibited read (either burst or single-beat, where `core_tt[0:4] = 0x1010`; see Table 8-6 for clarification), in which case the G2 core does not invalidate the snooped cache block. If the cache block is modified, the block is written back to memory, and the cache block is marked exclusive unmodified. If the cache block is marked exclusive unmodified when snooped, no bus action is taken, and the cache block remains in the exclusive unmodified state. This treatment of caching-inhibited reads decreases the possibility of data thrashing by allowing noncaching devices to read data without invalidating the entry from the G2 core data cache.

4.7.1 MEI State Definitions

The G2 core data cache characterizes each 32-byte block it contains as being in one of three MEI states. Addresses presented to the cache are indexed into the cache directory with bits A20:A26, and the upper-order 20 bits from the physical address translation (PA0–PA19) are compared against the indexed cache directory tags. If neither of the indexed tags matches, the result is a cache miss. If a tag matches, a cache hit occurred and the directory indicates the state of the cache block through two state bits kept with the tag. The three possible states for a cache block in the cache are the modified state (M), the exclusive state (E), and the invalid state (I). The three MEI states are defined in Table 4-2.

Table 4-2. MEI State Definitions

MEI State	Definition
Modified (M)	The addressed cache block is valid only in the cache. The cache block is modified with respect to system memory—that is, the modified data in the cache block has not been written back to memory.
Exclusive (E)	The addressed block is in this cache only. The data in this cache block is consistent with system memory.
Invalid (I)	This state indicates that the addressed cache block is not resident in the cache.

4.7.2 MEI State Diagram

The G2 core provides dedicated hardware to provide memory coherency by snooping bus transactions. The address retry capability of the G2 core enforces the MEI protocol, as shown in Figure 4-4. Figure 4-4 assumes that the WIM bits for the page or block are set to 001; that is, write-back, caching-not-inhibited, and memory coherency enforced.

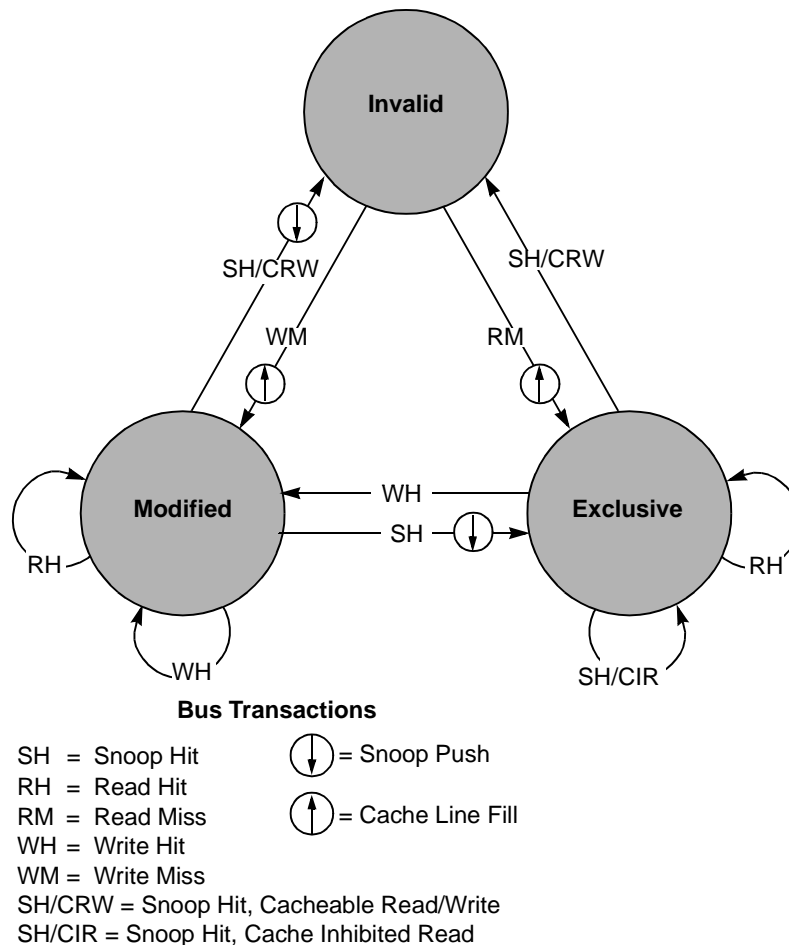


Figure 4-4. MEI Cache Coherency Protocol—State Diagram (WIM = 001)

Section 4.11, “MEI State Transactions,” provides a detailed list of MEI transitions for various operations and WIM bit settings.

4.7.3 MEI Hardware Considerations

While the G2 core provides the hardware required to monitor bus traffic for coherency, the G2 core data cache tags are single ported, and a simultaneous load or store and snoop access represent a resource conflict. In general, the snoop access has highest priority and is given first access to the tags. The load or store access will then occur on the clock following the snoop. The snoop is not given priority into the tags when the snoop coincides with a tag write (for example, validation after a cache block load). In these situations, the snoop is retried and must re-arbitrate before the lookup is possible.

Occasionally, cache snoops cannot be serviced and must be retried. These retries occur if the cache is busy with a burst read or write when the snoop operation takes place.

Note that it is possible for a snoop to hit a modified cache block that is already in the process of being written to the copy-back buffer for replacement purposes. If this happens, the G2

core retries the snoop, and raises the priority of the cast-out operation to allow it to go to the bus before the cache block fill.

The global ($\overline{\text{core_gbl}}$) signal, asserted as part of the address attribute field during a bus transaction, enables the snooping hardware of the G2 core. Address bus masters assert $\overline{\text{core_gbl}}$ to indicate that the current transaction is a global access (that is, an access to memory shared by more than one device). If $\overline{\text{core_gbl}}$ is not asserted for the transaction, that transaction is not snooped by the G2 core. Note that the $\overline{\text{core_gbl}}$ signal is not asserted for instruction fetches, and that $\overline{\text{core_gbl}}$ is asserted for all data read or write operations when using direct address translation. (Note that direct address translation is referred to as the real addressing mode, not the direct-store segment, in the architecture specification.)

Normally, $\overline{\text{core_gbl}}$ reflects the M-bit value specified for the memory reference in the corresponding translation descriptor(s). Care must be taken to minimize the number of pages marked as global, because the retry protocol enforces coherency and can use considerable bus bandwidth if a lot of data is shared. Therefore, available bus bandwidth can decrease as more traffic is marked global.

The G2 core snoops a transaction if the transfer start ($\overline{\text{core_ts}}$) and $\overline{\text{core_gbl}}$ signals are asserted together in the same bus clock (this is a qualified snooping condition). No snoop update to the G2 core cache occurs if the snooped transaction is not marked global. Also, because cache block cast-outs and snoop pushes do not require snooping, the $\overline{\text{core_gbl}}$ signal is not asserted for these operations.

When the G2 core detects a qualified snoop condition, the address associated with the $\overline{\text{core_ts}}$ signal is compared with the cache tags. Snooping finishes if no hit is detected. If, however, the address hits in the cache, the G2 core reacts according to the MEI protocol shown in Figure 4-4.

To facilitate external monitoring of the internal cache tags, the cache set entry signals ($\text{core_cse}[0:1]$) represent in binary the cache set being replaced on read operations (including read-with-intent-to-modify operations). The $\text{core_cse}[0:1]$ signals do not apply for write operations to memory, or during noncacheable or touch load operations. Note that these signals are valid only for G2 core burst operations. Table 4-3 shows the $\text{core_cse}[0:1]$ (cache set entry) encodings.

Table 4-3. $\text{core_cse}[0:1]$ Signal Encoding

$\text{core_cse}[0:1]$	Cache Set Element
00	Set 0
01	Set 1
10	Set 2
11	Set 3

4.7.4 Coherency Precautions

The G2 core supports a three-state coherency protocol that supports the modified, exclusive, and invalid (MEI) cache states. This protocol is a compatible subset of the MESI four-state protocol and operates coherently in systems that contain four-state caches. In addition, the G2 core does not broadcast cache operations caused by cache instructions. They are intended for the management of the local cache but not for other caches in the system.

4.7.4.1 Coherency in Single-Processor Systems

The following situations concerning coherency can be encountered within a single-processor system:

- Load or store to a caching-inhibited page (WIM = 0bx1x) and a cache hit occurs. Caching is inhibited for this page (I = 1)—Load or store operations to a caching-inhibited page that hit in the cache cause boundedly undefined results.
- Store to a page marked write-through (WIM = 0b10x) and a cache read hit to a modified cache block.

This page is marked as write-through (W = 1)—The G2 core pushes the modified cache block to memory and the block remains marked modified (M).

Note that when WIM bits are changed, it is critical that the cache contents reflect the new WIM bit settings. For example, if a block or page that had allowed caching becomes caching-inhibited, software should ensure that the appropriate cache blocks are flushed to memory and invalidated.

4.7.5 Load and Store Coherency Summary

Table 4-4 provides a summary of memory coherency actions performed by the G2 core on load operations. Noncacheable cases are not part of this table.

Table 4-4. Memory Coherency Actions on Load Operations

Cache State	Bus Operation	$\overline{\text{core_artry}}$	Action
M	None	Don't care	Read from cache
E	None	Don't care	Read from cache
I	Read	Negated	Load data and mark E
I	Read	Asserted	Retry read operation

Table 4-5 provides an overview of memory coherency actions on store operations. This table does not include noncacheable or write-through cases. The read-with-intent-to-modify (RWITM) examples involve selecting a replacement class and casting-out modified data that may have resided in that replacement class.

Table 4-5. Memory Coherency Actions on Store Operations

Cache State	Bus Operation	$\overline{\text{core_artry}}$	Action
M	None	Don't care	Modify cache
E	None	Don't care	Modify cache, mark M
I	RWITM	Negated	Load data, modify it, mark M
I	RWITM	Asserted	Retry the RWITM

4.7.6 Atomic Memory References

The Load Word and Reserve Indexed (**lwarx**) and Store Word Conditional Indexed (**stwcx.**) instructions provide an atomic update function for a single, aligned word of memory. While an **lwarx** instruction will normally be paired with an **stwcx.** instruction with the same effective address, an **stwcx.** instruction to any address will cancel the reservation. For detailed information on these instructions, refer to Chapter 3, “Instruction Set Model,” in this book and Chapter 8, “Instruction Set,” in the *Programming Environments Manual*.

4.7.7 Cache Reaction to Specific Bus Operations

There are several bus transaction types defined for the G2 core bus. The G2 core must snoop these transactions and perform the appropriate action to maintain memory coherency as shown in Table 4-6. A processor may assert `core_artry_out` for any bus transaction due to internal conflicts that prevent the appropriate snooping. The transactions in Table 4-6 correspond to the transfer type signals `core_tt[0:4]`, described in Section 8.3.4.1, “Transfer Type.”

Table 4-6. Response to Bus Transactions

Snooped Transaction	G2 Core Response
Clean block	No action is taken
Flush block	No action is taken
Write-with-flush Write-with-flush-atomic	Write-with-flush and write-with-flush-atomic operations occur after the processor issues a store or stwcx. instruction, respectively. <ul style="list-style-type: none"> If the addressed block is in the exclusive state, the address snoop forces the state of the addressed block to invalid. If the addressed block is in the modified state, the address snoop causes $\overline{\text{core_artry_out}}$ to be asserted and initiates a push of the modified block out of the cache and changes the state of the block to invalid. The execution of an stwcx. instruction cancels the reservation associated with any address.
Kill block	The kill block operation is an address-only bus transaction initiated when a dcbz instruction is executed; when snooped by the G2 core, the addressed cache block is invalidated if in the E state, or flushed to memory and invalidated if in the M state, and any associated reservation is canceled.

Table 4-6. Response to Bus Transactions (continued)

Snooped Transaction	G2 Core Response
Write-with-kill	In a write-with-kill operation, the processor snoops the cache for a copy of the addressed block. If one is found, an additional snoop action is initiated internally and the cache block is forced to the I state, killing modified data that may have been in the block. Any reservation associated with the block is also canceled.
Read Read-atomic	<p>The read operation is used by most single-beat and burst read operations on the bus. All burst reads observed on the bus are snooped as if they were writes, causing the addressed cache block to be flushed. A read on the bus with the <code>core_gbl</code> signal asserted causes the following responses:</p> <ul style="list-style-type: none"> • If the addressed block in the cache is invalid, the G2 core takes no action. • If the addressed block in the cache is in the exclusive state, the block is invalidated. • If the addressed block in the cache is in the modified state, the block is flushed to memory and the block is invalidated. • If the snooped transaction is a caching-inhibited read and the block in the cache is in the exclusive state, the snoop causes no bus activity and the block remains in the exclusive state. If the block is in the cache in the modified state, the G2 core initiates a push of the modified block out to memory and marks the cache block as exclusive. <p>Read-atomic operations appear on the bus in response to <code>lwarx</code> instructions and generate the same snooping responses as read operations.</p>
Read-with-intent-to-modify (RWITM) RWITM-atomic	<p>A RWITM operation is issued to acquire exclusive use of a memory location for the purpose of modifying it.</p> <ul style="list-style-type: none"> • If the addressed block is invalid, the G2 core takes no action. • If the addressed block in the cache is in the exclusive state, the G2 core initiates an additional snoop action to change the state of the cache block to invalid. • If the addressed block in the cache is in the modified state, the block is flushed to memory and the block is invalidated. <p>The RWITM-atomic operations appear on the bus in response to <code>stwcx.</code> instructions and are snooped like RWITM instructions.</p>
sync	No action is taken
TLB invalidate	No action is taken

4.7.8 Operations Causing `core_artry` Assertion

The following scenarios cause the G2 core to assert the `core_artry_out` signal:

- Snoop hits to a block in the M state (flush or clean)

This case is a normal snoop hit and will result in `core_artry_out` being asserted if the snooped transaction was a flush or clean request. If the snooped transaction was a kill request, `core_artry_out` will not be asserted.
- Snoop attempt during the last `core_ta` of a cache line fill

In No-`core_drtry` mode, during the cycle that the last `core_ta` is asserted to the G2 core on a cache line fill, the tag is being written to its new state by the G2 core and is not accessible. This will result in a collision being signaled by asserting `core_artry_out`. With `core_drtry` enabled, the cache tags are inaccessible to a snoop operation one cycle after the last `core_ta`.

- Snoop hit after the first $\overline{\text{core_ta}}$ of a burst load operation
After the first $\overline{\text{core_ta}}$ of a burst load operation, the data tags are committed to being written; snoop operations cannot be serviced until the load completes, thereby causing the assertion of $\overline{\text{core_artry_out}}$.
- Snoop hits to line in the cast-out buffer
The G2 core cast-out buffer is kept coherent with main memory, and snoop operations that hit in the cast-out buffer will cause the assertion of $\overline{\text{core_artry_out}}$.
- Snoop attempt during cycles that **dcbz** instruction or load or store operation is updating the tag
During the execution of a **dcbz** instruction or during a load or store operation that requires a cache line cast out, the cache tags will be inaccessible during the first and last cycle of the operation.
- Snoop attempt during the cycle that a **dcbf** or **dcbst** instruction is updating the tag
If the EA of a **dcbf** or **dcbst** instruction hits in the cache, the tag will be changed to its new state. During that clock, the tag is not accessible and snoop transactions during that cycle will cause the assertion of $\overline{\text{core_artry_out}}$.

4.7.9 Enveloped High-Priority Cache Block Push Operation

In cases where the G2 core has completed the address tenure of a read operation, and then detects a snoop hit to a modified cache block by another bus master, the G2 core provides a high-priority push operation. If the address snooped is the same as the address of the data to be returned by the read operation, $\overline{\text{core_artry_out}}$ is asserted one or more times until the data tenure of the read operation is completed. The cache block push transaction can be enveloped within the address and data tenures of a read operation. This feature prevents deadlocks in system organizations that support multiple memory-mapped buses.

More specifically, the G2 core internally detects the scenario where a load request is outstanding and the processor has pipelined a write operation on top of the load. Normally, when the data bus is granted to the G2 core, the resulting data bus tenure is used for the load operation. The enveloped high-priority cache block push feature defines a bus signal, the data bus write only qualifier ($\overline{\text{core_dbwo}}$), which, when asserted with a qualified data bus grant, indicates that the resulting data tenure should be used for the store operation instead. This signal is described in Section 9.10, “Using $\overline{\text{core_dbwo}}$ (Data Bus Write Only).” Note that the enveloped copy-back operation is an internally pipelined bus operation.

4.8 Cache Control Instructions

Software must use the appropriate cache management instructions to ensure that caches are kept consistent when data is modified by the processor. When a processor alters a memory location that may be contained in an instruction cache, software must ensure that updates

to memory are visible to the instruction fetching mechanism. Although the instructions to enforce coherency vary among implementations and, hence, operating systems should provide a system service for this function, the following sequence is typical:

1. **dcbst** (update memory)
2. **sync** (wait for update)
3. **icbi** (invalidate copy in cache)
4. **isync** (invalidate copy in own instruction buffer)

These operations are necessary because the processor does not maintain instruction memory coherent with data memory. Software is responsible for enforcing coherency of instruction caches and data memory. Since instruction fetching may bypass the data cache, changes made to items in the data cache may not be reflected in memory until after the instruction fetch completes.

The PowerPC architecture defines instructions for controlling both the instruction and data caches when they exist. The G2 core interprets the cache control instructions (**icbi**, **dcbi**, **dcbt**, **dcbz**, and **dcbst**) as if they pertain only to the G2 core caches. They are not intended for use in managing other caches in the system.

The **dcbz** instruction causes an address-only broadcast on the bus if the contents of the block are from a page marked global through the WIMG bits. This broadcast is performed for coherency reasons; the **dcbz** instruction is the only cache control instruction that can allocate and take new ownership of a line. Note that if the **HID0[ABE]** bit is set on a G2 core processor, the execution of the **dcbf**, **dcbi**, and **dcbst** instructions will also cause an address-only broadcast. The **dcbz** instruction is also the only cache operation that is snooped by the G2 core. The cache instructions are intended primarily for the management of the on-chip cache, and do not perform address-only broadcasts for the maintenance of other caches in the system. The ability of the G2 core to optionally perform address-only broadcasts when executing the **dcbi**, **dcbf**, and the **dcbst** instructions allows the coherency management of an external copy-back L2 cache. Note that the **dcbi** instruction should never be used on the G2 core.

The other instructions do not broadcast either for the purpose of invalidating or flushing other caches in the system or for managing system resources. Any bus activity caused by these instructions is the direct result of performing the operation in the G2 core cache. Note that a data access exception is generated if the effective address of a **dcbi**, **dcbst**, **dcbf**, or **dcbz** instruction cannot be translated due to the lack of a TLB entry. (Note that exceptions are referred to as interrupts in the architecture specification.)

Note that in the PowerPC architecture, the term ‘cache block’ or ‘block,’ when used in the context of cache implementations, refers to the unit of memory at which coherency is maintained. For the G2 core, this is the eight-word cache line. This value may be different for other implementations that support the PowerPC architecture. In-depth descriptions of

coding these instructions is provided in Chapter 3, “Addressing Modes and Instruction Set Summary,” and Chapter 10, “Instruction Set,” in the *Programming Environments Manual*.

4.8.1 Data Cache Block Invalidate (dcbi) Instruction

If the block containing the byte addressed by the EA is in the data cache, the cache block is invalidated regardless of whether the block is in the exclusive or modified state. If `HID0[ABE]` is set on a G2 core when a **dcbi** instruction is executed, the G2 core will perform an address-only bus transaction. The **dcbi** instruction can only be executed when the G2 core is in the supervisor state.

4.8.2 Data Cache Block Touch (dcbt) Instruction

This instruction provides a method for improving performance through the use of software-initiated prefetch hints. The G2 core performs the fetch when the address hits in the TLB or BAT registers, and when it is a permitted load access from the addressed page. The operation is treated similarly to a byte load operation with respect to coherency.

If the address translation does not hit in the TLB or BAT mechanism, or if it does not have load access permission, the instruction is treated as a no-op.

If the cache is locked or disabled, or if the access is to a page that is marked as guarded, the **dcbt** instruction is treated as a no-op.

If the access is directed to a write-through or caching-inhibited page, the instruction is treated as a no-op.

The **dcbt** instruction never affects the referenced or changed bits in the hashed page table.

A successful **dcbt** instruction affects the state of the TLB and cache LRU bits as defined by the LRU algorithm.

The touch load buffer will be marked invalid if the contents of the touch buffer have been moved to the cache, if any data cache management instruction has been executed, if a **dcbz** instruction is executed that matches the address of the cache block in the touch buffer, or if another **dcbt** instruction is executed.

4.8.3 Data Cache Block Touch for Store (dcbtst) Instruction

The **dcbtst** instruction, like the data cache block touch instruction (**dcbt**), allows software to prefetch a cache block in anticipation of a store operation (read-with-intent-to-modify).

4.8.4 Data Cache Block Clear to Zero (dcbz) Instruction

If the block containing the byte addressed by the EA is in the data cache, all bytes are cleared.

If the block containing the byte addressed by the EA is not in the data cache and the corresponding page is caching-allowed, the block is established in the data cache without fetching the block from main memory, and all bytes of the block are cleared. If the contents of the cache block are from a page marked global through the WIM bits, an address-only bus transaction is run.

If the page containing the byte addressed by the EA is caching-inhibited or write-through, then the system alignment exception handler is invoked.

The **dcbz** instruction is treated as a store to the addressed byte with respect to address translation and protection.

4.8.5 Data Cache Block Store (dcbst) Instruction

If the block containing the byte addressed by the EA is in coherency-required mode, and a block containing the byte addressed by the EA is in the data cache of any processor and has been modified, the writing of it to main memory is initiated. On a G2 core, if the cache block is unmodified, `HID0[ABE]` is set, and if the contents of the cache block are from a page marked global through the WIM bits, an address-only bus transaction is run.

The function of this instruction is independent of the write-through and caching-inhibited/caching-allowed modes of the block containing the byte addressed by the EA.

This instruction is treated as a load to the addressed byte with respect to address translation and protection.

4.8.6 Data Cache Block Flush (dcbf) Instruction

The action taken depends on the memory mode associated with the target, and on the state of the cache block. The following list describes the action taken for the various cases. These actions are executed regardless of whether the page containing the addressed byte is in caching-inhibited or caching-allowed mode. The following actions occur in both coherency-required (`WIM = 0bxx1`) and coherency-not-required mode (`WIM = 0bxx0`).

The **dcbf** instruction causes the following cache activity:

- Unmodified block—invalidates the block in the processor's cache
- Modified block—copies the block to memory and invalidates data cache block
- Absent block—does nothing

The G2 core treats this instruction as a load to the addressed byte with respect to address translation and protection.

4.8.7 Enforce In-Order Execution of I/O (eieio) Instruction

As defined by the PowerPC architecture, the **eieio** instruction provides an ordering function for the effects of load and store instructions executed by a given processor. Executing **eieio** ensures that all memory accesses previously initiated by the given processor are completed with respect to main memory before any memory accesses subsequently initiated by the processor access main memory. The **eieio** instruction orders loads and stores to caching-inhibited memory only.

The **eieio** instruction is intended for use only in performing memory-mapped I/O operations. It enforces strong ordering of cache-inhibited memory accesses during I/O operations between the processor and I/O devices.

When executed by the G2 core, the **eieio** instruction is treated as a no-op; caching-inhibited load and store operations (inhibited by the WIMG bits for the page) are performed in strict program order.

4.8.8 Instruction Cache Block Invalidate (icbi) Instruction

The execution of an **icbi** instruction causes all four cache sets indexed by the EA to be marked invalid. No cache hit is required, and no MMU translation is performed.

4.8.9 Instruction Synchronize (isync) Instruction

The **isync** instruction waits for all previous instructions to complete and then discards any previously fetched instructions, causing subsequent instructions to be fetched (or refetched) from memory and to execute in the context established by the previous instructions. This instruction has no effect on other processors or on their caches.

4.9 System Bus Interface and Cache Instructions

Table 4-7 provides an overview of the bus operations initiated by cache control instructions. The cache control, TLB management, and synchronization instructions supported by the G2 core may affect or be affected by the operation of the bus. None of the instructions will actively broadcast through address-only transactions on the bus (except for **dcbz**), and no broadcasts by other masters are snooped by the G2 core (except for kills). The operation of the instructions, however, may indirectly cause bus transactions to be performed, or their completion may be linked to the bus. Table 4-7 summarizes how these instructions may operate with respect to the bus.

Table 4-7. Bus Operations Caused by Cache Control Instructions (WIM = 001)

Operation	Cache State	Next Cache State	Bus Operations	Comment
sync	Don't care	No change	None	Waits for memory queues to complete bus activity
icbi	Don't care	I	None	—
dcbi ¹	Don't care	I	None	—
dcbf	I, E	I	None	—
dcbf	M	I	Write-with-kill	Block is pushed
dcbst	I, E	No change	None	—
dcbst	M	E	Write	Block is pushed
dcbz	I	M	Write-with-kill	—
dcbz	E, M	M	Kill block	Writes over modified data
dcbt	I	No change	Read	Fetches cache block is stored in touch load queue
dcbt	E, M	No change	None	—
dcbtst	I	No change	Read-with-intent-to-modify	Fetches cache block is stored in touch load queue
dcbtst	E, M	No change	None	—

¹ The **dcbi** instruction should never be used on the G2 core.

Note that Table 4-7 assumes that the WIM bits are set to 001; that is, since the cache is operating in write-back mode, caching is permitted and coherency is enforced.

Table 4-7 does not include noncacheable or write-through cases, nor does it completely describe the mechanisms for the operations described. For more information, see Section 4.11, “MEI State Transactions.”

For detailed information on the cache control instructions, refer to Chapter 3, “Instruction Set Model,” in this book and Chapter 8, “Instruction Set,” in the *Programming Environments Manual*. The G2 core contains snooping logic to monitor the bus for these commands and the control logic required to keep the cache and the memory queues coherent. For additional details about the specific bus operations performed by the G2 core, see Chapter 9, “Core Interface Operation.”

4.10 Bus Interface

The bus interface buffers bus requests from the instruction and data caches, and executes the requests per the G2 core bus protocol. It includes address register queues, prioritization logic, and bus control logic. The bus interface also captures snoop addresses for snooping in the cache and in the address register queues, snoops for reservations, and holds the touch load address for the cache. All data storage for the address register buffers (load and store

data buffers) are located in the cache section. The data buffers are considered temporary storage for the cache and not part of the bus interface.

The general functions and features of the bus interface are as follows:

- Address register buffers that include:
 - Instruction cache load address buffer
 - Data cache load address buffer
 - Data cache touch load address buffer (associated data block buffer located in cache)
 - Data cache cast-out/store address buffer (associated data line buffer located in cache)
 - Data cache snoop copy-back address buffer (associated data line buffer located in cache)
 - Reservation address buffer for snoop monitoring
- Pipeline collision detection for data cache buffers
- Reservation address snooping for **lwarx/stwcx**. instructions
- One-level address pipelining
- Loadahead of store capability

Figure 4-5 is a conceptual block diagram of the bus interface. The address register queues hold transaction requests that the bus interface may issue on the bus independently of the other requests. The bus interface may have up to two transactions operating on the bus at any given time through the use of address pipelining.

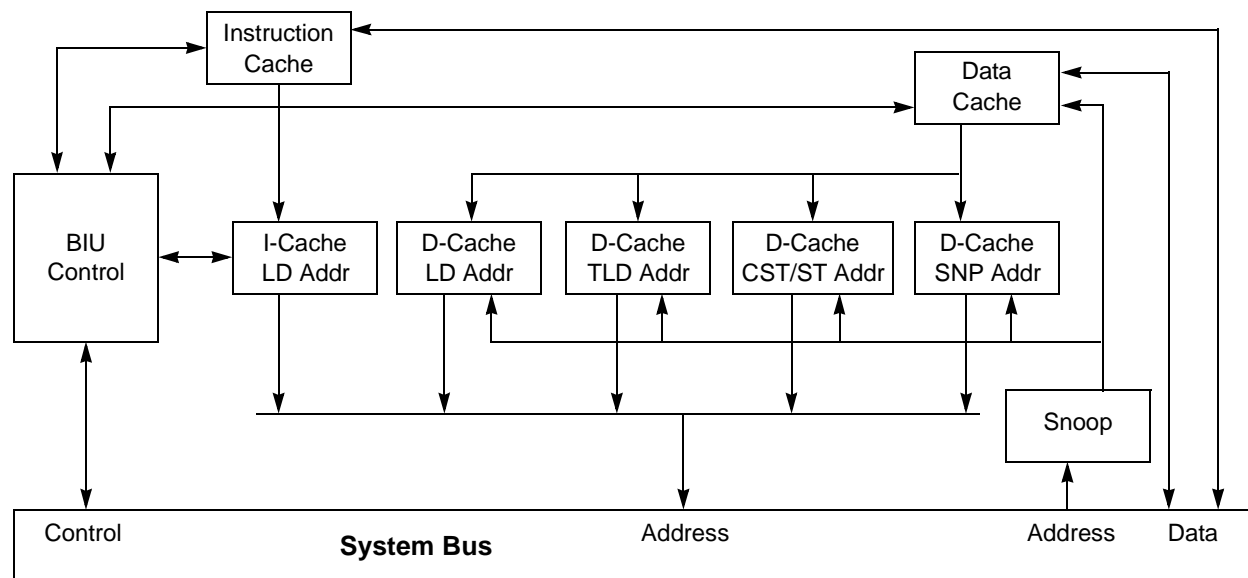


Figure 4-5. Bus Interface Address Buffers

For additional information about the G2 core bus interface and the bus protocols, refer to Chapter 9, “Core Interface Operation.”

4.11 MEI State Transactions

Table 4-8 shows MEI state transitions for various operations. Bus operations are described in Table 4-6.

Table 4-8. MEI State Transitions

Operation	Cache Operation	Bus Sync	WIM	Current State	Next State	Cache Actions	Bus Operation
Load (T = 0)	Read	No	x0x	I	Same	1 Cast out of modified block (as required)	Write-with-kill
						2 Pass four-beat read to memory queue	Read
Load (T = 0)	Read	No	x0x	E,M	Same	Read data from cache	—
Load (T = 0)	Read	No	x1x	I	Same	Pass single-beat read to memory queue	Read
Load (T = 0)	Read	No	x1x	E	I	CRTRY read	—
Load (T = 0)	Read	No	x1x	M	I	CRTRY read (push sector to write queue)	Write-with-kill
lwarx	Read	Acts like other reads but bus operation uses special encoding					
Store (T = 0)	Write	No	00x	I	Same	1 Cast out of modified block (if necessary)	Write-with-kill
						2 Pass RWITM to memory queue	RWITM
Store (T = 0)	Write	No	00x	E,M	M	Write data to cache	—
Store ≠ stwcx. (T = 0)	Write	No	10x	I	Same	Pass single-beat write to memory queue	Write-with-flush
Store ≠ stwcx. (T = 0)	Write	No	10x	E	Same	1 Write data to cache	—
						2 Pass single-beat write to memory queue	Write-with-flush
Store ≠ stwcx. (T = 0)	Write	No	10x	M	Same	1 CRTRY write	—
						2 Push block to write queue	Write-with-kill
Store (T = 0) or stwcx. (WIM = 10x)	Write	No	x1x	I	Same	Pass single-beat write to memory queue	Write-with-flush
Store (T = 0) or stwcx. (WIM = 10x)	Write	No	x1x	E	I	CRTRY write	—

Table 4-8. MEI State Transitions (continued)

Operation	Cache Operation	Bus Sync	WIM	Current State	Next State	Cache Actions	Bus Operation
Store (T = 0) or stwcx. (WIM = 10x)	Write	No	x1x	M	I	1 CRTRY write	—
						2 Push block to write queue	Write-with-kill
stwcx.	Conditional write	If the reserved bit is set, this operation is like other writes except the bus operation uses a special encoding.					
dcbf	Data cache block flush	No	xxx	I,E	Same	1 CRTRY dcbf	—
						2 Pass flush	Flush
				Same	I	3 State change only	—
dcbf	Data cache block flush	No	xxx	M	I	Push block to write queue	Write-with-kill
dcbst	Data cache block store	No	xxx	I,E	Same	1 CRTRY dcbst	—
						2 Pass clean	Clean
				Same	Same	3 No action	—
dcbst	Data cache block store	No	xxx	M	E	Push block to write queue	Write-with-kill
dcbz	Data cache block set to zero	No	x1x	x	x	Alignment trap	—
dcbz	Data cache block set to zero	No	10x	x	x	Alignment trap	—
dcbz	Data cache block set to zero	Yes	00x	I	Same	1 CRTRY dcbz	—
						2 Cast out of modified block	Write-with-kill
						3 Pass kill	Kill
				Same	M	4 Clear block	—
dcbz	Data cache block set to zero	No	00x	E,M	M	Clear block	—
dcbt	Data cache block touch	No	x1x	I	Same	Pass single-beat read to memory queue	Read
dcbt	Data cache block touch	No	x1x	E	I	CRTRY read	—
dcbt	Data cache block touch	No	x1x	M	I	1 CRTRY read	—
						2 Push block to write queue	Write-with-kill

Table 4-8. MEI State Transitions (continued)

Operation	Cache Operation	Bus Sync	WIM	Current State	Next State	Cache Actions	Bus Operation
dcbt	Data cache block touch	No	x0x	I	Same	1 Cast out of modified block (as required)	Write-with-kill
						2 Pass four-beat read to memory queue	Read
dcbt	Data cache block touch	No	x0x	E,M	Same	No action	—
Single-beat read	Reload dump 1	No	xxx	I	Same	Forward data_in	—
Four-beat read (double-word-aligned)	Reload dump	No	xxx	I	E	Write data_in to cache	—
Four-beat write (double-word-aligned)	Reload dump	No	xxx	I	M	Write data_in to cache	—
E→I	Snoop write or kill	No	xxx	E	I	State change only (committed)	—
M→I	Snoop kill	No	xxx	M	I	State change only (committed)	—
Push M→I	Snoop flush	No	xxx	M	I	Conditionally push	Write-with-kill
Push M→E	Snoop clean	No	xxx	M	E	Conditionally push	Write-with-kill
tlbie	TLB invalidate	No	xxx	x	x	1 CRTRY TLBI	—
						2 Pass TLBI	—
						3 No action	—
sync	Synchroni- zation	No	xxx	x	x	1 CRTRY sync	—
						2 Pass sync	—
						3 No action	—

Note: Single-beat writes are not snooped in the write queue.

4.12 Cache Locking

This section describes the entire cache locking and cache way-locking features of the G2 core.

4.12.1 Cache Locking Terminology

Cache locking refers to the ability to prevent some or all of a processor’s instruction or data cache from being overwritten. Cache locking can be set for either an entire cache or for individual ways within the cache as follows:

- Entire cache locking—When an entire cache is locked, data for read hits within the cache are supplied to the requesting unit in the same manner as hits from an unlocked cache. Similarly, writes that hit in the data cache are written to the cache in the same way as write hits to an unlocked cache. However, any access that misses in the cache is treated as a cache-inhibited access. Cache entries that are invalid at the time of locking remain invalid and inaccessible until the cache is unlocked. When the cache has been unlocked, all entries (including invalid entries) are available. Entire cache locking is inefficient if the number of instructions or the size of data to be locked is small compared to the cache size.
- Way-Locking—Locking only a portion of the cache is accomplished by locking ways within the cache. Locking always begins with the first way (way 0) and is sequential, that is, locking ways 0, 1, and 2 is possible, but it is not possible to lock only way 0 and way 2. When using way-locking, at least two ways must be left unlocked. The maximum number of lockable ways is six on the G2 core (way 0–way 5).

Unlike entire cache locking, invalid entries in a locked way are accessible and available for data replacement. As hits to the cache fill invalid entries within a locked way, the entries become valid and locked. This behavior differs from entire cache locking in which invalid entries cannot be allocated. Unlocked ways of the cache behave normally.

Table 4-9 summarizes the G2 core cache organization.

Table 4-9. Cache Organization

Instruction Cache Size	Data Cache Size	Associativity	Block Size	Way Size
16 Kbytes	16 Kbytes	4-way	8 words	4 Kbytes

4.12.2 Cache Locking Register Summary

Table 4-10 through Table 4-12 outline the registers and bits used to perform cache locking on the G2 core. Refer to Section 2.1.2.1, “Hardware Implementation Register 0 (HID0),” for a complete description of the HID0 and MSR registers. Refer to Section 2.1.2.3, “Hardware Implementation Register 2 (HID2),” for a complete description of the HID2 register.

Table 4-10. HID0 Bits Used to Perform Cache Locking

Bits	Name	Description
16	ICE	Instruction cache enable. This bit must be set for instruction cache locking. See Section 4.12.3.1.1, “Enabling the Data Cache.”
17	DCE	Data cache enable. This bit must be set for data cache locking. See Section 4.12.3.1.1, “Enabling the Data Cache.”
18	ILOCK	Instruction cache lock. Set to lock the entire instruction cache. See Section 4.12.3.2.5, “Entire Instruction Cache Locking.”
19	DLOCK	Data cache lock. Set to lock the entire data cache. See Section 4.12.3.1.6, “Entire Data Cache Locking.”
20	ICFI	Instruction cache flash invalidate. Setting and then clearing this bit invalidates the entire instruction cache. See Section 4.12.3.2.7, “Invalidating the Instruction Cache (Even if Locked).”
21	DCFI	Data cache flash invalidate. Setting and then clearing this bit invalidates the entire data cache. See Section 4.12.3.1.4, “Invalidating the Data Cache.”

Table 4-11. HID2 Bits Used to Perform Cache Way-Locking

Bits	Name	Description
16–18	IWLCK	Instruction cache way-lock. These bits are used to lock individual ways in the instruction cache. See Section 4.12.3.2.6, “Instruction Cache Way-Locking.”
24–26	DWLCK	Data cache way-lock. These bits are used to lock individual ways in the data cache. See Section 4.12.3.1.7, “Data Cache Way-Locking.”

Table 4-12. MSR Bits Used to Perform Cache Locking

Bits	Name	Description
16	EE	External interrupt enable. This bit must be cleared during instruction and data cache loading. See Section 4.12.3.1.3, “Disabling Exceptions for Data Cache Locking.”
19	ME	Machine check enable. This bit must be cleared during instruction and data cache loading. See Section 4.12.3.1.3, “Disabling Exceptions for Data Cache Locking.”
26	IR	Instruction address translation. This bit must be set to enable instruction address translation by the MMU. See Section 4.12.3.1.2, “Address Translation for Data Cache Locking.”
27	DR	Data address translation. This bit must be set to enable data address translation by the MMU. See Section 4.12.3.1.2, “Address Translation for Data Cache Locking.”

4.12.3 Performing Cache Locking

This section outlines the basic procedures for locking the data and instruction caches and provides some example code for locking the caches. The procedures for the data cache are described first, followed by the corresponding sections for locking the instruction cache.

The basic procedures for cache locking are:

- Enabling the cache
- Enabling address translation for example code

- Disabling exceptions
- Loading the cache
- Locking the cache (entire cache locking or cache way-locking)

In addition, this section describes how to invalidate the data and instruction caches, even when they are locked.

4.12.3.1 Data Cache Locking

This section describes the procedures for performing data cache locking on the G2 core.

4.12.3.1.1 Enabling the Data Cache

To lock the data cache, the data cache enable bit `HID0[DCE]`, bit 17, must be set. The following assembly code enables the data cache:

```
# Enable the data cache. This corresponds
# to setting DCE bit in HID0 (bit 17)

mfspr    r1, HID0
ori      r1, r1, 0x4000
sync
mtspr    HID0, r1
```

4.12.3.1.2 Address Translation for Data Cache Locking

Two distinct memory areas must be set up to enable cache locking:

- The first area is where the code that performs the locking resides and is executed from
- The second area is where the data to be locked resides

Both areas of memory must be in locations that are translated by the memory management unit (MMU). This translation can be performed either with the page table or the block address translation (BAT) registers.

For the purposes of the cache locking example in this document, the two areas of memory are defined using the BAT registers. The first area is a 1-Mbyte area in the upper region of memory that contains the code performing the cache locking. The second area is a 256-Mbyte block of memory (not all of the 256 Mbytes of memory is locked in the cache; this area is set up as an example) that contains the data to lock. Both memory areas use identity translation (the logical memory address equals the physical memory address).

Table 4-13 summarizes the BAT settings used in this example.

Table 4-13. Example BAT Settings for Cache Locking

Area	Base Address	Memory Size	WIMG Bits	BATU Setting	BATL Setting
First	0xFFFF0_0000	1 Mbyte	0b0100 ¹	0xFFFF0_001F	0xFFFF0_0002 ¹
Second	0x0000_0000	256 Mbyte	0b0000	0x0000_1FFF	0x0000_0002

¹ Cache-inhibited memory is not a requirement for data cache locking. A setting of 0xFFFF0_0002 with a corresponding WIMG of 0b0000 marks the memory area as cacheable.

The block address translation upper (BATU) and block address translation lower (BATL) settings in Table 4-13 can be used for both instruction block address translation (IBAT) and data block address translation (DBAT) registers. After the BAT registers have been set up, the MMU must be enabled. The following assembly code enables both instruction and data memory address translation:

```
# Enable instruction and data memory address translation. This
# corresponds to setting IR and DR in the MSR (bits 26 & 27)

mfmsr    r1
ori      r1, r1, 0x0030
mtmsr    r1
sync
```

4.12.3.1.3 Disabling Exceptions for Data Cache Locking

To ensure that exception handler routines do not execute while the cache is being loaded (which could possibly pollute the cache with undesired contents) all exceptions must be disabled. This is accomplished by clearing the appropriate bits in the machine state register (MSR). See Table 4-14 for the bits within the MSR that must be cleared to ensure that exceptions are disabled.

Table 4-14. MSR Bits for Disabling Exceptions

Bits	Name	Description
16	EE	External interrupt enable
19	ME	Machine check enable
20	FE0 ¹	Floating-point exception mode 0
23	FE1 ¹	Floating-point exception mode 1
24	CE	Critical interrupt enable

¹ The floating-point exception may not need to be disabled because the example code shown in this document that performs cache locking does not execute any floating-point operations.

The following assembly code disables all asynchronous exceptions:

```
# Clear the following bits from the MSR:
#   EE (16)      ME (19)
#   FE0 (20)    FE1 (23)
#   ME (24)
```

```

mfmsr    r1
lis      r2, 0xFFFF
ori      r2, r2, 0x667F
and      r1, r1, r2
mtmsr    r1
sync

```

4.12.3.1.4 Invalidating the Data Cache

If a non-empty data cache has modified data, and the data cannot be discarded, the data cache must be flushed before it can be invalidated. Data cache flushing is accomplished by filling the data cache with known data and performing a flash invalidate or a series of **dcbf** instructions that force a flush and invalidation of the data cache block.

The following code sequence shows how to flush the data cache:

```

# r6 contains a block-aligned address in memory with which to fill
# the data cache. For this example, address 0x0 is used
    li      r6, 0x0

# CTR = number of data blocks to load
# Number of blocks = (16K) / (32 Bytes/block)
#                   = 2^14 / 2^5 = 2^9 = 0x200
    li      r1, 0x200
    mtctr   r1

# Save the total number of blocks in cache to r8
    mr      r8, r1

# Load the entire cache with known data
loop:  lwz   r2, 0(r6)
       addi r6, r6, 32      # Find the next block
       bdnz loop          # Decrement the counter, and
# branch if CTR != 0

# Now, flush the cache with dcbf instructions
    li      r6, 0x0        # Address of first block
    mtctr   r8            # Number of blocks

loop2: dcbf   r0, r6
       addi r6, r6, 32      # Find the next block
       bdnz loop2         # Decrement the counter, and
# branch if CTR != 0

```

If the content of the data cache does not need to be flushed to memory, the cache can be directly invalidated. The entire data cache is invalidated through the data cache flash invalidate bit `HID0[DCFI]`, bit 21. Setting `HID0[DCFI]` and then immediately clearing it

causes the entire data cache to be invalidated. The following assembly code invalidates the entire data cache (does not flush modified entries):

```
# Set and then clear the HID0[DCFI] bit, bit 21
    mfspr    r1, HID0
    mr       r2, r1
    ori      r1, r1, 0x0400
    mtspr    HID0, r1
    mtspr    HID0, r2
    sync
```

4.12.3.1.5 Loading the Data Cache

This section explains loading data into the data cache. The data cache can be loaded in several ways. The example in this document loads the data from memory. The following assembly code loads the data cache:

```
# Assuming interrupts are turned off, cache has been flushed,
# MMU on, and loading from contiguous cacheable memory.
# r6 = Starting address of code to lock
# r20 = Temporary register for loading into
# CTR = Number of cache blocks to lock

loop:  lwz     r20, 0(r6)      # Load data into d-cache
       addi   r6, r6, 32     # Find next block to load
       bdnz  loop          # CTR = CTR-1, branch if CTR != 0
```

4.12.3.1.6 Entire Data Cache Locking

Locking of the entire data cache is controlled by the data cache lock bit (HID0[DLOCK], bit 19). Setting HID0[DLOCK] to 1 locks the entire data cache. To unlock the data, the HID0[DLOCK] must be cleared to 0. Setting the DLOCK bit must be preceded by a **sync** instruction to prevent the data cache from being locked during a data access. The following assembly code locks the entire data cache:

```
# Set the DLOCK bit in HID0 (bit 19)

    mfspr    r1, HID0
    ori      r1, r1, 0x1000
    sync
    mtspr    HID0, r1
```

4.12.3.1.7 Data Cache Way-Locking

Data cache way-locking is controlled by HID2[DWLCK], bits 24–26. Table 4-15 shows the HID2[DWLCK[0–2]] settings for the G2 core embedded processor.

Table 4-15. G2 Core DWLCK[0–2] Encodings

DWLCK[0:2]	Ways Locked
0b000	No ways locked
0b001	Way 0 locked
0b010	Ways 0 and 1 locked
0b011	Ways 0, 1, and 2 locked
0b100	Ways 0, 1, 2, and 3 locked
0b101	Ways 0, 1, 2, 3, and 4 locked
0b110	Ways 0, 1, 2, 3, 4, and 5 locked
0b111	Reserved

The following assembly code locks way 0 of the G2 core data cache:

```
# Lock way 0 of the data cache
# This corresponds to setting dwlck(0-2) 0b001 (bits 24-26)

mfspr    r1, HID2
lis      r2, 0xFFFF
ori      r2, r2, 0xFF1F
and      r1, r1, r2
ori      r1, r1, 0x0020
sync
mtspr    HID2, r1
```

4.12.3.1.8 Invalidating the Data Cache (Even if Locked)

There are two methods to invalidate the instruction or data cache:

- Invalidate the entire cache by setting and then immediately clearing the data cache flash invalidate bit HID0[DCFI], bit 21. Even when a cache is locked, toggling DCFI bit invalidates all of the data cache.
- The data cache block invalidate (**dcbi**) instruction can be used to invalidate individual cache blocks on other devices. However, the **dcbi** instruction should never be used on the G2 core.

4.12.3.2 Instruction Cache Locking

This section describes the procedures for performing instruction cache locking on the G2 core.

4.12.3.2.1 Enabling the Instruction Cache

To lock the instruction cache, the instruction cache enable bit HID0[ICE], bit 16 must be set.

```
# Enable the data cache. This corresponds
# to setting DCE bit in HID0 (bit 17)
```

```
    mfspr    r1, HID0
    ori      r1, r1, 0x8000
    sync
    mtspr    HID0, r1
```

4.12.3.2.2 Address Translation for Instruction Cache Locking

Two distinct memory areas must be set up to enable cache locking:

- The first area is where the code that performs the locking resides and is executed from
- The second area is where the instructions to be locked reside

Both areas of memory must be in locations that are translated by the memory management unit (MMU). This translation can be performed either with the page table or the block address translation (BAT) registers.

For the purposes of the cache locking example in this document, two areas of memory are defined using the BAT registers. The first area is a 1-Mbyte area in the upper region of memory that contains the code performing the cache locking. This area of memory must be cache-inhibited for instruction cache locking. The second area is a 256-Mbyte block of memory (not all of the 256 Mbytes of memory is locked in the cache; this area is set up as an example) that contains the instructions to lock. Both memory areas use identity translation (the logical memory address equals the physical memory address). Table 4-16 summarizes the BAT settings used in this example.

Table 4-16. Example BAT Settings for Cache Locking

Area	Base Address	Memory Size	WIMG Bits	BATU Setting	BATL Setting
First	0xFFFF0_0000	1 Mbyte	0b0100 ¹	0xFFFF0_001F	0xFFFF0_0022 ¹
Second	0x0000_0000	256 Mbytes	0b0000	0x0000_1FFF	0x0000_0002

¹ 0xFFFF0_0022 defines a cache-inhibited memory area used for instruction cache locking, and corresponds to a WIMG of 0b0100. Cache-inhibited memory is not a requirement for data cache locking. A setting of 0xFFFF0_0002 with a corresponding WIMG of 0b0000 marks the memory area as cacheable.

The block address translation upper (BATU) and block address translation lower (BATL) settings in Table 4-16 can be used for both instruction block address translation (IBAT) and data block address translation (DBAT) registers. After the BAT registers have been set up, the MMU must be enabled.

The following assembly code enables both instruction and data memory address translation:

```
# Enable instruction and data memory address translation. This
# corresponds to setting IR and DR in the MSR (bits 26 & 27)
```

```
mfmsr    r1
ori      r1, r1, 0x0030
mtmsr    r1
sync
```

4.12.3.2.3 Disabling Exceptions for Instruction Cache Locking

To ensure that exception handler routines do not execute while the cache is being loaded (which could possibly pollute the cache with undesired contents) all exceptions must be disabled. This is accomplished by clearing the appropriate bits in the machine state register (MSR). See Table 4-17 for the bits within the MSR that must be cleared to ensure that exceptions are disabled.

Table 4-17. MSR Bits for Disabling Exceptions

Bit	Name	Description
16	EE	External interrupt enable
19	ME	Machine check enable
20	FE0 ¹	Floating-point exception mode 0
23	FE1 ¹	Floating-point exception mode 1
24	CE	Critical interrupt enable

¹ The floating-point exception may not need to be disabled because the example code shown in this document that performs cache locking does not execute any floating-point operations.

The following assembly code disables all asynchronous exceptions:

```
# Clear the following bits from the MSR:
#   EE (16)      ME (19)
#   FE0 (20)    FE1 (23)
#   ME (24)

mfmsr    r1
lis      r2, 0xFFFF
ori      r2, r2, 0x667F
and      r1, r1, r2
mtmsr    r1
sync
```

4.12.3.2.4 Preloading Instructions into the Instruction Cache

To optimize performance, processors that implement the PowerPC architecture automatically prefetch instructions into the instruction cache. This feature can be used to preload explicit instructions into the cache even when it is known that their execution will be canceled. Although the execution of the instructions is canceled, the instructions remain valid in the instruction cache.

Because instructions are intentionally executed speculatively, care must be taken to ensure that all I/O memory is marked guarded. Otherwise, speculative loads and stores to I/O space could potentially cause data loss. See the *Programming Environments Manual* for a full discussion of guarded memory.

The code that prefetches must be in cache-inhibited memory as in the following example:

```
# Assuming exceptions are disabled, cache has been flushed,
# the MMU is on, and we are executing in a cache-inhibited
# location in memory
# LR and r6 = Starting address of code to lock
# CTR = Number of cache blocks to lock
# r2 = non-zero numerator and denominator
# 'loop' must begin on an 8-byte boundary to ensure that
#   the divw and beqlr+ are fetched on the same cycle.

.orig    0xFFFF04000

loop:    divw.    r2, r2, r2        # LONG divide w/ non-zero result
         beqlr+   # Cause the prefetch to happen

         addi    r6, r6, 32        # Find next block to prefetch
         mtlr    r6                # set the next block
         bdnz-   loop             # Decrement the counter and
         # branch if CTR != 0
```

In the above example, both the **divw** and **beqlr+** instructions are fetched at the same time (this assumes a 64-bit 60x data bus; the preloading code does not work for a 32-bit data bus) due to their placement on a double-word boundary. The divide instruction was chosen because it takes many cycles to execute. During execution of the divide, the processor starts fetching instructions speculatively at the target destination of the branch instruction. The speculation occurs because the branch is statically predicted as taken. This speculative fetching causes the cache block that is pointed to by the link register (LR) to be loaded into the cache. Because the **divw** instruction always produces a non-zero result, the **beqlr+** is not taken and execution of all speculatively fetched instructions is canceled. However, the instructions remain valid in the cache.

If the destination instruction stream contains an unconditional branch to another memory location, it is possible to also prefetch the destination of the unconditional branch instruction. This does not cause a problem if the destination of the unconditional branch is also inside the area of memory that needs to be preloaded. But if the destination of the unconditional branch is not in the area of memory to be loaded, then care must be taken to ensure that the branch destination is to an area of memory that is cache inhibited. Otherwise, unintentional instructions may be locked in the cache and the desired instructions may not be in their expected way within the cache.

4.12.3.2.7 Invalidating the Instruction Cache (Even if Locked)

There are two methods to invalidate the instruction cache. In the first way, invalidate the entire cache by setting and then immediately clearing the instruction cache flash invalidate bit (HID0[ICFI], bit 20). Even when a cache is locked, toggling the ICFI bit invalidates all of the instruction cache. The following assembly code invalidates the entire instruction cache:

```
# Set and then clear the HID0[ICFI] bit, bit 20
```

```
    mfspr    r1, HID0
    mr       r2, r1
    ori      r1, r1, 0x0800

    mtspr    HID0, r1
    mtspr    HID0, r2
    sync
```

In the second method, the instruction cache block invalidate (**icbi**) instruction can be used to invalidate individual cache blocks. The **icbi** instruction invalidates blocks in an entirely locked instruction cache. The **icbi** instruction also may invalidate way-locked blocks within the instruction cache.

Chapter 5

Exceptions

The PowerPC exception mechanism allows the processor to change to supervisor state as a result of external signals, errors, or unusual conditions arising in the execution of instructions, and differ from the arithmetic exceptions defined by the IEEE for floating-point operations. When exceptions (referred to as interrupts in the architecture specification) occur, information about the state of the processor is saved to certain registers and the processor begins execution at an address (exception vector) predetermined for each exception. Processing of exceptions occurs in supervisor mode.

Although multiple exception conditions can map to a single exception vector, a more specific condition may be determined by examining a register associated with the exception—for example, the DSISR or FPSCR. Additionally, certain exception conditions can be explicitly enabled or disabled by software.

The PowerPC architecture requires that exceptions be handled in program order; therefore, although a particular implementation may recognize exception conditions out of order, they are handled strictly in order with respect to the instruction stream. When an instruction-caused exception is recognized, any unexecuted instructions that appear earlier in the instruction stream, including any that have not yet entered the execute state, are required to complete before the exception is taken. Any exceptions caused by those instructions are handled first. Likewise, exceptions that are asynchronous and precise are recognized when they occur, but are not handled until the instruction currently in the completion stage successfully completes execution or generates an exception, and the completed store queue is emptied (see Section 7.1, “Terminology and Conventions,” for the definition). An instruction is said to have completed when the results of that instruction’s execution have been committed to the registers defined by the architecture (for example, the GPRs or FPRs, rather than rename buffers). If a single instruction encounters multiple exception conditions, those exceptions are taken and handled sequentially. Likewise, exceptions that are asynchronous are recognized when they occur, but are not handled until the next instruction to complete in program order successfully completes. Throughout this chapter, the phrase ‘next instruction’ implies the next instruction to complete in program order.

Note that exceptions can occur while an exception handler routine is executing, and multiple exceptions can become nested. It is up to the exception handler to save the states to allow control to ultimately return to the original excepting program.

Unless a catastrophic condition causes a system reset or machine check exception, only one exception is handled at a time. If, for example, a single instruction encounters multiple exception conditions, those conditions are handled sequentially. After the exception handler handles an exception, the instruction execution continues until the next exception condition is encountered. However, in many cases there is no attempt to re-execute the instruction. This method of recognizing and handling exception conditions sequentially guarantees that exceptions are recoverable.

To prevent loss of state information, exception handlers should save the information stored in SRR0 and SRR1 soon after the exception is taken. This prevents loss of information due to a system reset or machine check exception or to an instruction-caused exception in the exception handler before disabling external interrupts.

In this chapter, the following terminology is used to describe the various stages of exception processing:

Recognition	Exception recognition occurs when the condition that can cause an exception is identified by the processor.
Taken	An exception is said to be taken when control of instruction execution is passed to the exception handler; that is, the context is saved and the instruction at the appropriate vector offset is fetched and the exception handler routing is executed in supervisor mode.
Handling	Exception handling is performed by the software linked to the appropriate vector offset. Exception handling is performed at the supervisor-level.

5.1 Exception Classes

The PowerPC architecture supports four types of exceptions:

- Synchronous, precise—These are caused by instructions. All instruction-caused exceptions are handled precisely; that is, the machine state at the time the exception occurs is known and can be completely restored. This means that (excluding the trap and system call exceptions) the address of the faulting instruction is provided to the exception handler and that neither the faulting instruction nor subsequent instructions in the code stream will complete execution before the exception is taken. Once the exception is processed, execution resumes at the address of the faulting instruction (or at an alternate address provided by the exception handler). When an exception is taken due to a trap or system call instruction, execution resumes at an address provided by the handler.
- Synchronous, imprecise—The PowerPC architecture defines two imprecise floating-point exception modes: recoverable and nonrecoverable. Even though the G2 core provides a means to enable the imprecise modes, it implements these modes

identically to the precise mode (that is, all enabled floating-point exceptions are always precise on the G2 core).

- **Asynchronous, maskable**—The external (`core_int`), system management interrupt (`core_smi`), and decremter exceptions are maskable asynchronous exceptions. The critical interrupt (`core_cint`) exception of the G2_LE core is also a maskable asynchronous exception. When these exceptions occur, their handling is postponed until the next instruction completes execution and until any exceptions associated with that instruction complete execution. If there are no instructions in the execution units, the exception is taken immediately upon determination of the correct restart address (for loading SRR0).
- **Asynchronous, nonmaskable**—There are two nonmaskable asynchronous exceptions: system reset and the machine check exception. These exceptions may not be recoverable, or may provide a limited degree of recoverability. All exceptions report recoverability through the MSR[RI] bit.

The G2 core exception classes are shown in Table 5-1.

Table 5-1. Exception Classifications

Synchronous/Asynchronous	Precise/Imprecise	Exception Type
Asynchronous, nonmaskable	Imprecise	Machine check System reset
Asynchronous, maskable	Precise	External interrupt Decrementer System management interrupt Critical interrupt (G2_LE core only)
Synchronous	Precise	Instruction-caused exceptions

Table 5-1 defines exception categories that are handled uniquely by the G2 core. Note that Table 5-1 includes no synchronous imprecise exceptions. While the PowerPC architecture supports imprecise handling of floating-point exceptions, the G2 core implements floating-point exception modes as precise exceptions. Although the PowerPC architecture specifies that the recognition of the machine check exception is nonmaskable, on the G2 core the stimuli that cause this exception are maskable. For example, the machine check exception is caused by the assertion of `core_tea`, `core_ape`, `core_dpe`, or `core_mcp`. However, `core_mcp`, `core_ape`, and `core_dpe` can be disabled by bits 0, 2, and 3, respectively, in `HID0`. Therefore, the machine check caused by asserting `core_tea` is the only truly nonmaskable machine check exception.

The G2 core exceptions, and conditions that cause them, are listed in Table 5-2.

Table 5-2. Exceptions and Conditions

Exception Type	Vector Offset (hex)	Causing Conditions
Reserved	00000	—
System reset	00100	A system reset is caused by the assertion of either <code>core_sreset</code> or <code>core_hreset</code> .
Machine check	00200	A machine check is caused by the assertion of the <code>core_tea</code> signal during a data bus transaction, assertion of <code>core_mcp</code> , or an address or data parity error.
DSI	00300	<p>The cause of a DSI exception can be determined by the bit settings in the DSISR, listed as follows:</p> <ol style="list-style-type: none"> 1 Set if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared. 4 Set if a memory access is not permitted by the page or DBAT protection mechanism; otherwise cleared. 5 Set by an <code>eciwX</code> or <code>ecowX</code> instruction if the access is to an address that is marked as write-through, or execution of a load/store instruction that accesses a direct-store segment. 6 Set for a store operation and cleared for a load operation 9 G2_LE core only. Set a data address breakpoint exception occurs when the data [0–28] in the DABR or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> • Write breakpoints enabled when DABR[30] is set • Read breakpoints enabled when DABR[31] is set 11 Set if <code>eciwX</code> or <code>ecowX</code> is used and EAR[E] is cleared
ISI	00400	<p>An ISI exception is caused when an instruction fetch cannot be performed for any of the following reasons:</p> <ul style="list-style-type: none"> • The effective (logical) address cannot be translated. That is, there is a page fault for this portion of the translation, so an ISI exception must be taken to load the PTE (and possibly the page) into memory. • The fetch access is to a direct-store segment (indicated by SRR1[3] set) • The fetch access violates memory protection (indicated by SRR1[4] set). If the key bits (Ks and Kp) in the segment register and the PP bits in the PTE are set to prohibit read access, instructions cannot be fetched from this location.
External interrupt	00500	An external interrupt is caused when MSR[EE] = 1 and the <code>core_int</code> signal is asserted.
Alignment	00600	<p>An alignment exception is caused when the core cannot perform a memory access for any of the reasons described below:</p> <ul style="list-style-type: none"> • The operand of a floating-point load or store instruction is not word-aligned. • The operand of <code>lmw</code>, <code>stmw</code>, <code>lwarx</code>, and <code>stwcx</code> instructions are not aligned. • The execution of a floating-point load or store instruction to a direct-store segment. • The operand of a load, store, load multiple, store multiple, load string, or store string instruction crosses a segment boundary into a direct-store segment, or crosses a protection boundary. • Execution of a misaligned <code>eciwX</code> or <code>ecowX</code> instruction. • The instruction is <code>lmw</code>, <code>stmw</code>, <code>lswi</code>, <code>lswx</code>, <code>stswi</code>, <code>stswx</code>, and the G2 core is in little-endian mode. This applies to both modified little-endian and true little-endian mode for G2_LE core. • The operand of <code>dcbz</code> is in memory that is write-through-required or caching-inhibited.

Table 5-2. Exceptions and Conditions (continued)

Exception Type	Vector Offset (hex)	Causing Conditions
Program	00700	<p>A program exception is caused by one of the following exception conditions, which correspond to bit settings in SRR1 and arise during execution of an instruction.</p> <p>Floating-point enabled exception—A floating-point enabled exception condition is generated when the following condition is met:</p> <p style="text-align: center;">(MSR[FE0] MSR[FE1]) & FPSCR[FEX] is 1.</p> <ul style="list-style-type: none"> • FPSCR[FEX] is set by the execution of a floating-point instruction that causes an enabled exception or by the execution of one of the ‘move to FPSCR’ instructions that results in both an exception condition bit and its corresponding enable bit being set in the FPSCR. • Illegal instruction—An illegal instruction program exception is generated when execution of an instruction is attempted with an illegal opcode or illegal combination of opcode and extended opcode fields (including PowerPC instructions not implemented in the core), or when execution of an optional instruction not provided in the core is attempted (these do not include those optional instructions that are treated as no-ops). • Privileged instruction—A privileged instruction type program exception is generated when the execution of a privileged instruction is attempted and the MSR register user privilege bit, MSR[PR], is set. In the G2 core, this exception is generated for mtspr or mfspir with an invalid SPR field if SPR[0] = 1 and MSR[PR] = 1. This may not be true for all cores that implement the PowerPC architecture. • Trap—A trap type program exception is generated when any of the conditions specified in a trap instruction is met.
Floating-point unavailable	00800	A floating-point unavailable exception is caused by an attempt to execute a floating-point instruction (including floating-point load, store, and move instructions) when the floating-point available bit is cleared (MSR[FP] = 0).
Decrementer	00900	The decrementer exception occurs when DEC[31] changes from 0 to 1. This exception is enabled with MSR[EE].
Critical interrupt	00A00	A critical interrupt exception is taken when the core_cint signal is asserted and MSR[CE] = 1 (G2_LE only).
Reserved	00B00–00BFF	—
System call	00C00	A system call exception occurs when a System Call (sc) instruction is executed.
Trace	00D00	A trace exception is taken when MSR[SE] = 1 or when the currently completing instruction is a branch and MSR[BE] = 1.
Reserved	00E00	The G2 core does not generate an exception to this vector. Other devices may use this vector for floating-point assist exceptions.
Reserved	00E10–00FFF	—
Instruction translation miss	01000	An instruction translation miss exception is caused when the effective address for an instruction fetch cannot be translated by the ITLB.
Data load translation miss	01100	A data load translation miss exception is caused when the effective address for a data load operation cannot be translated by the DTLB.
Data store translation miss	01200	A data store translation miss exception is caused when the effective address for a data store operation cannot be translated by the DTLB, or where a DTLB hit occurs, and the change bit in the PTE must be set due to a data store operation.

Table 5-2. Exceptions and Conditions (continued)

Exception Type	Vector Offset (hex)	Causing Conditions
Instruction address breakpoint	01300	An instruction address breakpoint exception occurs when the address (bits 0–29) in the IABR matches the next instruction to complete in the completion unit, and IABR[30] is set. Note that the G2_LE core also implements IABR2, which functions identically to IABR.
System management interrupt	01400	A system management interrupt is caused when MSR[EE] = 1 and the $\overline{\text{core_smi}}$ input signal is asserted.
Reserved	01500–02FFF	—

Exceptions are roughly prioritized by exception class, as follows:

1. Nonmaskable, asynchronous exceptions have priority over all other exceptions—system reset and machine check exceptions (although the machine check exception condition can be disabled so the condition causes the processor to go directly into the checkstop state). These exceptions cannot be delayed, and do not wait for the completion of any precise exception handling.
2. Synchronous, precise exceptions are caused by instructions and are taken in strict program order.
3. Maskable asynchronous exceptions (for example, external interrupt and decremter exceptions) are delayed until higher priority exceptions are taken.

System reset and machine check exceptions may occur at any time and are not delayed even if an exception is being handled. As a result, state information for the interrupted exception may be lost; therefore, these exceptions are typically nonrecoverable.

All other exceptions have lower priority than system reset and machine check exceptions, and the exception may not be taken immediately when it is recognized.

5.1.1 Exception Priorities

The exceptions are listed in Table 5-3 in order of highest to lowest priority.

Table 5-3. Exception Priorities

Exception Category	Priority	Exception	Cause
Asynchronous	0	System reset	$\overline{\text{core_hreset}}$ or power-on reset
	1	Machine check	$\overline{\text{core_tea}}$, $\overline{\text{core_mcp}}$, $\overline{\text{core_ape}}$, or $\overline{\text{core_dpe}}$
	2	System reset	$\overline{\text{core_sreset}}$
	3	Critical interrupt	$\overline{\text{core_cint}}$ (G2_LE-only)
	4	System management interrupt	$\overline{\text{core_smi}}$

Table 5-3. Exception Priorities (continued)

Exception Category	Priority	Exception	Cause
Asynchronous (continued)	5	External interrupt	core_int
	6	Decrementer exception	Decrementer passed through 0x00000000
Instruction fetch	0	ITLB miss	Instruction TLB miss
	1	Instruction access	Instruction access exception
Instruction dispatch/execution	0	IABR	Instruction address breakpoint exception
	1	Program	Program exception due to the following: <ul style="list-style-type: none"> • Illegal instruction • Privileged instruction • Trap
	2	System call	System call exception
	3	Floating-point unavailable	Floating-point unavailable exception
	4	Program	Program exception due to a floating-point enabled exception
	5	Alignment	Alignment exception due to the following: <ul style="list-style-type: none"> • Floating-point not word-aligned • lmw, stmw, lwarx, or stwcx. not word-aligned • ecwix or ecowx operands not aligned • Multiple or string access with little-endian bit set
	6	Data access	Data access exception due to a BAT page protection violation
	7	Data access	Data access exception due to the following: <ul style="list-style-type: none"> • eciwx, ecowx, lwarx, or stwcx. to direct-store segment (bit 5 of DSISR) • Crossing from memory segment to direct-store segment (bit 0 of DSISR) • Crossing from direct-store segment to memory segment • Any access to direct-store, SR[T] = 1 • eciwx or ecowx with EAR[E] = 0 (bit 11 of DSISR)
	8	DTLB miss	Data TLB miss exception due to: <ul style="list-style-type: none"> • Store miss • Load miss
9	Alignment	Alignment exception due to a dcbz to a write-through or caching-inhibited page	
	10	Data access	Data access exception due to TLB page protection violation
	11	DTLB miss	Data TLB miss exception due to a change bit not set on a store operation
Post-instruction execution	0	Trace	Trace exception due to the following: <ul style="list-style-type: none"> • MSR[SE] = 1 • MSR[BE] = 1 for branches

Exception priorities are described in detail in “Exception Priorities,” in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

5.1.2 Summary of Front-End Exception Handling

The following list of interrupt categories describes how the G2 core handles exceptions up to the point of signaling the appropriate exception to occur. Note that a recoverable state is reached if the completed store queue is empty (drained, not canceled) and any instruction that is next in program order and has been signaled to complete has completed. If MSR[RI] is clear, the core is in a nonrecoverable state by default. Also, completion of an instruction is defined as performing all architectural register writes associated with that instruction, and then removing that instruction from the completion buffer queue.

- Asynchronous nonmaskable nonrecoverable—(system reset caused by the assertion of either `core_hreset` or internally during power-on reset (POR)). These exceptions have highest priority and are taken immediately regardless of other pending exceptions or recoverability. A nonpredicted address is guaranteed.
- Asynchronous maskable nonrecoverable—(machine check). A machine check exception takes priority over any other pending exception except a nonrecoverable system reset caused by the assertion of either `core_hreset` or internally during POR. A machine check exception is taken immediately regardless of recoverability. A machine check exception can occur only if the machine check enable bit, MSR[ME], is set. If MSR[ME] is cleared, the processor goes directly into checkstop state when a machine check exception condition occurs. A nonpredicted address is guaranteed.
- Asynchronous nonmaskable recoverable—(system reset caused by the assertion of `core_sreset`). This interrupt takes priority over any other pending exceptions except nonrecoverable exceptions listed above. This exception is taken immediately when a recoverable state is reached.
- Asynchronous maskable recoverable—(system management interrupt, critical interrupt (G2_LE only), external interrupt, decremter exception). Before handling this type of exception, the next instruction in program order must complete or except. If this action causes another type of exception, that exception is taken and the asynchronous maskable recoverable exception remains pending. Once an instruction can complete without causing an exception, further instruction completion is halted while the exception not taken remains pending. The exception is taken when a recoverable state is reached.
- Instruction fetch—(ITLB, ISI). When this type of exception is detected, dispatch is halted and the current instruction stream is allowed to drain. If completing any instructions in this stream causes an exception, that exception is taken and the instruction fetch exception is forgotten. Otherwise, as soon as the machine is empty and a recoverable state is reached, the instruction fetch exception is taken.
- Instruction dispatch/execution—(program, DSI, alignment, emulation trap, system call, DTLB miss on load or store, IABR). This type of exception is determined at dispatch or execution of an instruction. The exception remains pending until all instructions in program order before the exception-causing instruction are

completed. The exception is then taken without completing the exception-causing instruction. If any other exception condition is created in completing these previous instructions in the machine, that exception takes priority over the pending instruction dispatch/execution exception, which will then be forgotten.

- Post-instruction execution—(trace). This type of exception is generated following execution and completion of an instruction while a trace mode is enabled. If executing the instruction produces conditions for another type of interrupt, that exception is taken and the post-instruction execution exception is forgotten for that instruction.

5.2 Exception Processing

When an exception is taken, the processor uses the save/restore registers, SRR0 and SRR1, to save the contents of the machine state register for user-level mode and to identify where instruction execution should resume after the exception is handled.

5.2.1 Exception Processing Registers

The G2 core implements the SRR0 and SRR1 registers that are used for saving processor state on an exception. The G2_LE core also uses these registers; additionally, the G2_LE core implements CSRR0 and CSRR1 to specifically save state for critical interrupt exceptions.

5.2.1.1 SRR0 and SRR1 Bit Settings

When an exception occurs, SRR0 is set to point to the instruction at which instruction processing should resume when the exception handler returns control to the interrupted process. All instructions in the program flow preceding this one will have completed and no subsequent instruction will have completed. This may be the address of the instruction that caused the exception or the next one (as in the case of a system call exception). The instruction addressed can be determined from the exception type and status bits. This address is used to resume instruction processing in the interrupted process, typically when an **rfi** instruction is executed. The SRR0 register is shown in Figure 5-1.

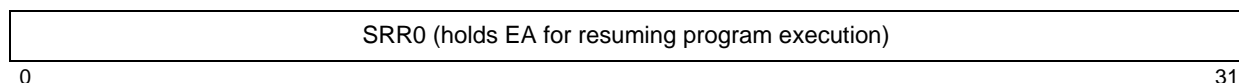
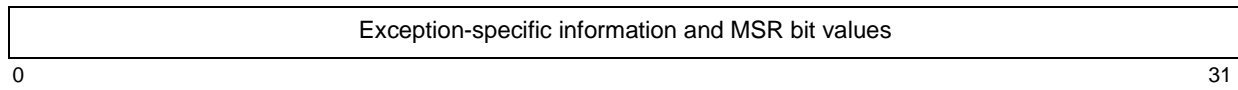


Figure 5-1. Machine Status Save/Restore Register 0 (SSR0)

The save/restore register 1 (SRR1) is used to save machine status (the contents of the MSR) on exceptions and to restore those values when **rfi** is executed. SRR1 is shown in Figure 5-2.



0

31

Figure 5-2. Machine Status Save/Restore Register 1 (SSR1)

Typically, when an exception occurs, bits 0–15 of SRR1 are loaded with exception-specific information and bits 16–31 of MSR are placed into the corresponding bit positions of SRR1. The G2 core loads SRR1 with specific bits for handling machine check exceptions, as shown in Table 5-4.

Table 5-4. SRR1 Bit Settings for Machine Check Exceptions

Bits	Name	Description
0	MSR[0]	Copy of MSR bit 0
1–4	—	Reserved
5–9	MSR[5–9]	Copy of MSR bits 5–9
10–11	—	Reserved
12	MCP	Machine check
13	TEA	TEA error
14	DPE	Data parity error
15	APE	Address parity error
16–31	MSR[16–31]	Copy of MSR bits 16–31

The G2 core loads SRR1 with specific bits for handling the three TLB miss exceptions, as shown in Table 5-5.

Table 5-5. SRR1 Bit Settings for Software Table Search Operations

Bits	Name	Description
0–3	CRF0	Copy of condition register field 0 (CR0)
4	—	Reserved
5–9	MSR[5–9]	Copy of MSR bits 5–9
10–11	—	Reserved
12	KEY	TLB miss protection key
13	I/D	Instruction/data TLB miss 0 DTLB miss 1 ITLB miss
14	WAY	Bit 14 indicates which TLB associativity set should be replaced 0 Set 0 1 Set 1
15	S/L	Store/load protection instruction 0 Load miss 1 Store miss
16–31	MSR[16–31]	Copy of MSR bits 16–31

Note that in some implementations, every instruction fetch when $MSR[IR] = 1$ and every instruction execution requiring address translation when $MSR[DR] = 1$ may modify SRR1.

5.2.1.2 CSRR0 and CSRR1 Bit Settings—G2_LE Only

The G2_LE core also implements the CSRR0 and CSRR1 to save state for critical interrupt exceptions only. Note that the values saved in CSRR0 are the same as those saved in SRR0 for all other exceptions, and the values saved in CSRR1 are the same as those saved in SRR1 for all other exceptions. However, CSRR0 and CSRR1 have unique SPR numbers, as described in Chapter 2, “Register Model.”

Figure 5-3 shows the format of CSRR0.



Figure 5-3. Critical Interrupt Save/Restore Register 0 (CSRR0)

When a critical interrupt exception occurs, CSRR0 is set to point to an instruction such that all prior instructions have completed execution and no subsequent instruction has begun execution. When an **rfci** instruction is executed, the contents of CSRR0 are copied to the next instruction address (NIA)—the 32-bit address of the next instruction to be executed. Figure 5-4 shows the format of CSRR1.

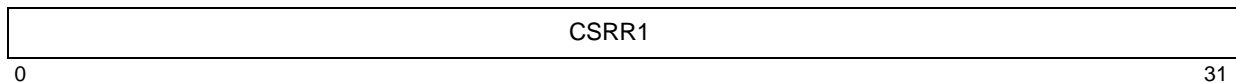


Figure 5-4. Critical Interrupt Save/Restore Register 1 (CSRR1)

When an exception occurs, CSRR1[0–15] are loaded with all zeros and the values of MSR[16–31] are placed in corresponding CSRR1 bit positions. When **rfci** executes, MSR[16–31] are loaded from CSRR1[16–31].

CSRR1[0–15] are defined as reserved. An implementation may define one or more of these bits, and may also cause them to be saved from MSR when an exception is taken, and restored to MSR from CSRR1 when an **rfci** is executed.

5.2.1.3 SPRG4–SPRG7 (G2_LE Only)

The G2_LE core provides four additional SPRG (SPRG4–SPRG7) registers for general operating system use, such as performing a fast state save or for supporting multiprocessor implementations. However, SPRG4–SPRG7 have unique SPR numbers, as described in Chapter 2, “Register Model.” The formats of SPRG4–SPRG7 are shown in Figure 5-5.

Table 5-7. MSR Bit Settings (continued)

Bit	Name	Description
14	TGPR	Temporary GPR remapping (implementation-specific) 0 Normal operation 1 TGPR mode. GPR0–GPR3 are remapped to TGPR0–TGPR3 for use by TLB miss routines The contents of GPR0–GPR3 remain unchanged while MSR[TGPR] = 1. Attempts to use GPR4–GPR31 with MSR[TGPR] = 1 yield undefined results. Temporarily replaces TGPR0–TGPR3 with GPR0–GPR3 for use by TLB miss routines. The TGPR bit is set when either an instruction TLB miss, data read miss, or data write miss exception is taken. The TGPR bit is cleared by an rfi instruction.
15	ILE	Exception little-endian mode. When an exception occurs, this bit is copied into MSR[LE] to select the endian mode for the context established by the exception.
16	EE	External interrupt enable 0 The processor ignores external interrupts, system management interrupts, and decrementer interrupts. 1 The processor is enabled to take an external interrupt, system management interrupt, or decrementer interrupt.
17	PR	Privilege level 0 The processor can execute both user- and supervisor-level instructions. 1 The processor can only execute user-level instructions.
18	FP	Floating-point available 0 The processor prevents dispatch of floating-point instructions, including floating-point loads, stores, and moves. 1 The processor can execute floating-point instructions, and can take floating-point enabled exception type program exceptions.
19	ME	Machine check enable 0 Machine check exceptions are disabled 1 Machine check exceptions are enabled
20	FE0	Floating-point exception mode 0 (see Table 5-8)
21	SE	Single-step trace enable 0 The processor executes instructions normally 1 The processor generates a trace exception on the successful completion of the next instruction
22	BE	Branch trace enable 0 The processor executes branch instructions normally 1 The processor generates a trace exception upon the successful completion of a branch instruction
23	FE1	Floating-point exception mode 1 (see Table 5-8)
24	CE	Critical interrupt exception enable (G2_LE core-only) 0 Critical interrupts disabled 1 Critical interrupts enabled; critical interrupt exception and rftci instruction enabled. The critical interrupt is an asynchronous implementation-specific exception. The critical interrupt exception vector offset is 0x00A00. The Return From Critical Interrupt (rftci) instruction is implemented to return from these exception handlers. Also, CSRR0 and CSRR1, are used to save and restore the processor state for critical interrupts.
25	IP	Exception prefix. The setting of this bit specifies whether an exception vector offset is prepended with Fs or 0s. In the following description, <i>nnnn</i> is the offset of the exception. See Table 5-2. 0 Exceptions are vectored to the physical address 0x000n_nnnn 1 Exceptions are vectored to the physical address 0xFFFFn_nnnn

Table 5-7. MSR Bit Settings (continued)

Bit	Name	Description
26	IR	Instruction address translation 0 Instruction address translation is disabled 1 Instruction address translation is enabled See Chapter 6, “Memory Management.”
27	DR	Data address translation 0 Data address translation is disabled 1 Data address translation is enabled See Chapter 6, “Memory Management.”
28–29	—	Reserved. Full function.
30	RI	Recoverable exception (for system reset and machine check exceptions) 0 Exception is not recoverable 1 Exception is recoverable
31	LE	Little-endian mode enable 0 The processor runs in big-endian mode 1 The processor runs in little-endian mode. For the G2_LE core, see Section 1.1.2.1, “True Little-Endian Mode,” for a definition of whether the core is operating in true little-endian mode or modified little-endian mode.

The IEEE floating-point exception mode bits (FE0 and FE1) together define whether floating-point exceptions are handled precisely, imprecisely, or if they are taken at all. The possible settings and default conditions for the G2 core are shown in Table 5-8. For further details, see Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

Table 5-8. IEEE Floating-Point Exception Mode Bits

FE0	FE1	Mode
0	0	Floating-point exceptions disabled
0	1	Floating-point imprecise nonrecoverable ¹
1	0	Floating-point imprecise recoverable ¹
1	1	Floating-point precise mode

¹ Not implemented in the G2 core.

MSR bits are guaranteed to be written to SRR1 when the first instruction of the exception handler is encountered.

5.2.2 Enabling and Disabling Exceptions

When a condition exists that may cause an exception to be generated, it must be determined whether the exception is enabled for that condition as follows:

- IEEE floating-point enabled exceptions (a type of program exception) are ignored when both MSR[FE0] and MSR[FE1] are cleared. If either of these bits are set, all IEEE enabled floating-point exceptions are taken and cause a program exception.

- Asynchronous, maskable exceptions (that is, the external, system management, and decremter interrupts) are enabled by setting the MSR[EE] bit. When MSR[EE] = 0, recognition of these exception conditions is delayed. MSR[EE] is cleared automatically when an exception is taken, to delay recognition of conditions causing those exceptions.
- A machine check exception can occur only if the machine check enable bit, MSR[ME], is set. If MSR[ME] is cleared, the processor goes directly into checkstop state when a machine check exception condition occurs. Individual machine check exceptions can be enabled and disabled through bits in the HID0 register, as described in Table 2-5.
- The G2_LE core enables the critical interrupt with the MSR[CE] bit.
- System reset exceptions cannot be masked.

5.2.3 Steps for Exception Processing

After it is determined that the exception can be taken (by confirming that any instruction-caused exceptions occurring earlier in the instruction stream have been handled, and by confirming that the exception is enabled for the exception condition), the processor does the following:

1. The machine status save/restore register 0 (SRR0) is loaded with an instruction address that depends on the type of exception. See the individual exception description for details about how this register is used for specific exceptions.
2. SRR1[1–4, 10–15] are loaded with information specific to the exception type.
3. SRR1[5–9, 16–31] are loaded with a copy of the corresponding bits of the MSR.
4. The MSR is set as described in Table 5-7. The new values take effect beginning with the fetching of the first instruction of the exception-handler routine located at the exception vector address.

Note that MSR[IR] and MSR[DR] are cleared for all exception types; therefore, address translation is disabled for both instruction fetches and data accesses beginning with the first instruction of the exception-handler routine.

5. Instruction fetch and execution resumes, using the new MSR value, at a location specific to the exception type. The location is determined by adding the exception's vector (see Table 5-2) to the base address determined by MSR[IP]. If IP is cleared, exceptions are vectored to the physical address 0x000n_nnnn. If IP is set, exceptions are vectored to the physical address 0xFFFFn_nnnn. For a machine check exception that occurs when MSR[ME] = 0 (machine check exceptions are disabled), the processor enters the checkstop state (the machine stops executing instructions). See Section 5.5.2, “Machine Check Exception (0x00200).”

Note that the same steps occur when a critical interrupt occurs (and is enabled) for the G2_LE core, except that CSRR0 is set instead of SRR0 and CSRR1 is set instead of SRR1.

5.2.4 Setting MSR[RI]

The operating system should handle MSR[RI] as follows:

- In the machine check and system reset exceptions—If SRR1[RI] is cleared, the exception is not recoverable. If it is set, the exception is recoverable with respect to the processor.
- In each exception handler—When enough state information has been saved that a machine check or system reset exception can reconstruct the previous state, set MSR[RI].
- In each exception handler—Clear MSR[RI], set the SRR0 and SRR1 (or CSRR0 and CSRR1) registers appropriately, and then execute **rfi** (or **rfci**).
- Note that the RI bit being set indicates that, with respect to the processor, enough processor state data is valid for the processor to continue, but it does not guarantee that the interrupted process can resume.

5.2.5 Returning From an Exception Handler with rfi

The Return From Interrupt (**rfi**) instruction performs context synchronization by allowing previously issued instructions to complete before returning to the interrupted process. In general, execution of the **rfi** instruction ensures the following:

- All previous instructions have completed to a point where they can no longer cause an exception. If a previous instruction causes a direct-store interface error exception, the results must be determined before this instruction is executed.
- Previous instructions complete execution in the context (privilege, protection, and address translation) under which they were issued.
- The **rfi** instruction copies SRR1 bits back into the MSR.
- The instructions following this instruction execute in the context established by this instruction.

For a complete description of context synchronization, refer to Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

5.2.6 Returning From an Interrupt with rfci

The Return From Critical Interrupt (**rfci**) is a G2_LE core-only supervisor level instruction that performs context synchronization by allowing previously issued instructions to complete before returning to the interrupted process. The **rfci** instruction performs the same

functions as **rfi**, except that it uses CSRR0 and CSRR1 to restore the processor state. Thus, execution of the **rfci** instruction ensures the following:

- CSRR1[0, 5–9, 16–31] are placed into the corresponding bits of the MSR. If the new MSR value does not enable any pending exceptions, the next instruction is fetched from the address defined by CSRR0[0–29] || 0b00.
- If the new MSR value enables one or more pending exceptions, the exception associated with the highest priority pending exception is generated. In this case, the exception processing mechanism places in SRR0 the address of the instruction which would have executed next had the exception not occurred.

5.3 Process Switching

The operating system should execute one of the following when processes are switched:

- The **sync** instruction, which orders the effects of instruction execution. All instructions previously initiated appear to have completed before the **sync** instruction completes, and no subsequent instructions appear to be initiated until the **sync** instruction completes. For an example showing the use of a **sync** instruction, see Chapter 2, “Register Set,” of the *Programming Environments Manual*.
- The **isync** instruction, which waits for all previous instructions to complete and then discards any fetched instructions, causing subsequent instructions to be fetched (or refetched) from memory and to execute in the context (privilege, translation, protection, etc.) established by the previous instructions.
- The **stwcx.** instruction, to clear any outstanding reservations, which ensures that an **lwarx** instruction in the old process is not paired with an **stwcx.** instruction in the new process.

The operating system should set the MSR[RI] bit as described in Section 5.2.4, “Setting MSR[RI].”

5.4 Exception Latencies

Latencies for taking various exceptions depend on the state of the machine when the exception conditions occur. This latency may be as short as one cycle, in which case an exception is signaled in the cycle following the appearance of the exception condition. The latencies are as follows:

- Hard reset and machine check—In most cases, a hard reset or machine check exception will have a single-cycle latency. A two- to three-cycle delay may occur only when a predicted instruction is next to complete, and the branch guess that forced this instruction to be predicted was resolved to be incorrect.
- Soft reset—The latency of a soft reset exception is affected by recoverability. The time to reach a recoverable state may depend on the time needed to complete or

except an instruction at the point of completion, the time needed to drain the completed store queue (see Section 7.1, “Terminology and Conventions,” for the definition), and the time waiting for a correct empty state so that a valid MSR[IP] may be saved. For lower-priority externally-generated interrupts, a delay may be incurred waiting for another interrupt generated while reaching a recoverable state to be serviced.

Further delays are possible for other types of exceptions depending on the number and type of instructions that must be completed before those exceptions may be serviced. See Section 5.1.2, “Summary of Front-End Exception Handling,” to determine possible maximum latencies for different exceptions.

5.5 Exception Definitions

Table 5-9 shows all the types of exceptions that can occur with the G2 core and the MSR bit settings when the processor transitions to supervisor mode. The state of these bits prior to the exception is typically stored in SRR1 (or CSRR1 for critical interrupts on the G2_LE core). Note that MSR[CE] is cleared for the following exceptions in system reset, machine check, and critical interrupt.

Table 5-9. MSR Setting Due to Exception

Exception Type	MSR Bit																
	POW	TGPR	ILE	EE	PR	FP	ME	FE0	SE	BE	FE1	CE ¹	IP	IR	DR	RI	LE
System reset	0	0	—	0	0	0	—	0	0	0	0	0	1	0	0	0	ILE
Machine check	0	0	—	0	0	0	0	0	0	0	0	0	—	0	0	0	ILE
DSI	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
ISI	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
External	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
Alignment	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
Program	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
Floating-point unavailable	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
Decrementer	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
Critical Interrupt	0	0	—	0	0	0	—	0	0	0	0	0	—	0	0	0	ILE
System call	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
Trace exception	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
ITLB miss	0	1	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
DTLB miss on load	0	1	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
DTLB miss on store	0	1	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE

Table 5-9. MSR Setting Due to Exception (continued)

Exception Type	MSR Bit																
	POW	TGPR	ILE	EE	PR	FP	ME	FE0	SE	BE	FE1	CE ¹	IP	IR	DR	RI	LE
Instruction address breakpoint	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE
System management interrupt	0	0	—	0	0	0	—	0	0	0	0	—	—	0	0	0	ILE

Note:

0 Bit is cleared.

1 Bit is set.

ILE Bit is copied from the ILE bit in the MSR.

— Bit is not altered.

Reserved bits are read as if written as 0.

¹ G2_LE core only.

5.5.1 Reset Exceptions (0x00100)

The system reset exception is a nonmaskable, asynchronous exception signaled to the G2 core either through the assertion of the reset signals (`core_sreset` or `core_hreset`) or internally during the power-on reset (POR) process. The assertion of the soft reset signal, `core_sreset`, as described in Section 8.3.10.2, “Soft Reset (`core_sreset`)—Input,” causes the system reset exception to be taken and the physical base address of the handler is determined by the MSR[IP] bit.

The assertion of the hard reset signal, `core_hreset`, as described in Section 8.3.10.1, “Hard Reset (`core_hreset`)—Input,” causes the system reset exception to be taken.

Note that there are some byte ordering precautions necessary when coming out of reset in big-endian mode and switching to little-endian mode. The following sections describe the differences between a hard and soft reset and the byte ordering implications for reset exception handling.

5.5.1.1 Hard Reset and Power-On Reset

As described in Section 5.1.2, “Summary of Front-End Exception Handling,” the hard reset exception is a nonrecoverable, nonmaskable asynchronous exception. When `core_hreset` is asserted or at power-on reset (POR), the G2 core immediately branches to the address determined by the state of the `core_msrip` signal, as described in Table 5-10, without attempting to reach a recoverable state.

Table 5-10. Hard Reset MSR Value and Exception Vector

core_msrip	MSR[0–31]	Fetch Instructions from Handler at System Reset Vector
asserted	0x0000_0040 (MSR[IP] = 1)	0xFFF0_0100
negated	0x0000_0000 (MSR[IP] = 0)	0x0000_0100

A hard reset has the highest priority of any exception, and is always nonrecoverable. Table 5-11 shows the state of the machine just before it fetches the first instruction of the system reset handler after a hard reset.

Table 5-11. Settings Caused by Hard Reset

Register	Setting	Register	Setting
GPRs	Unknown	PVR	See Table 2-3
FPRs	Unknown	HID0	0000_0000
FPSCR	00000000	HID1	0000_0000
CR	All 0s	HID2	0000_0000 or 0800_0000
SRs	Unknown	DMISS and IMISS	All 0s
MSR	0000_0040 or 0000_0000 or 0001_0041 or 0001_0001	DCMP and ICMP	All 0s
XER	0000_0000	RPA	All 0s
TBU	0000_0000	IABR	All 0s
TBL	0000_0000	DSISR	0000_0000
LR	0000_0000	DAR	0000_0000
CTR	0000_0000	DEC	FFFF_FFFF
SDR1	0000_0000	HASH1	0000_0000
SRR0 (and CSRR0)	0000_0000	HASH2	0000_0000
SRR1 (and CSRR1)	0000_0000	TLBs	Unknown
SPRGs	0000_0000	Cache	All cache blocks invalidated
Tag directory	All 0s. (However, LRU bits are initialized so each side of the cache has a unique LRU value.)	BATs	Unknown

The $\overline{\text{core_hreset}}$ signal can be asserted for the following reasons:

- System power-on reset
- System reset from a panel switch

For information on the $\overline{\text{core_hreset}}$ signal, see Section 8.3.10.1, “Hard Reset (core_hreset)—Input.”

The following is also true after a hard reset operation:

- External checkstops are enabled

- The on-chip test interface has given control of the I/Os to the rest of the chip for functional use
- Since the reset exception has data and instruction translation disabled (MSR[DR] and MSR[IR] both cleared), the chip operates in real addressing mode as described in Section 6.2, “Real Addressing Mode.”

5.5.1.2 Soft Reset

As described in Section 5.1.2, “Summary of Front-End Exception Handling,” the soft reset exception is a type of system reset exception that is recoverable, nonmaskable, and asynchronous. When `core_sreset` is asserted, the processor attempts to reach a recoverable state by allowing the next instruction to either complete or cause an exception, blocking the completion of subsequent instructions, and allowing the completed store queue to drain (see Section 7.1, “Terminology and Conventions,” for the definition).

Unlike a hard reset, no registers or latches are initialized; however, the instruction cache is disabled (`HID0[ICE] = 0`). After `core_sreset` is recognized as asserted, the processor begins fetching instructions from the system reset routine at offset `0x0100`. When a soft reset occurs, registers are set as shown in Table 5-12. A soft reset is recoverable provided that attaining the recoverable state does not cause a machine check exception. This interrupt case is third in priority, following hard reset and machine check.

When a soft reset occurs, registers are set as shown in Table 5-12 in addition to the clearing of `HID0[ICE]`.

Table 5-12. Soft Reset Exception—Register Settings

Register	Setting Description			
SRR0	Set to the effective address of the instruction that the processor would have attempted to complete next if no exception conditions were present.			
SRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]. Note that if the processor state is corrupted to the extent that execution cannot be reliably restarted, SRR1[30] is cleared.			
MSR	POW 0	FP 0	FE1 0	RI 0
	TGPR 0	ME —	CE 0	LE Set to value of ILE
	ILE —	FE0 0	IP —	
	EE 0	SE 0	IR 0	
	PR 0	BE 0	DR 0	

5.5.1.3 Byte Ordering Considerations for G2_LE Only

All exception handler routines are executed in the endian mode determined by the setting of the `MSR[ILE]`, `MSR[LE]`, and `HID2[LET]` bits (see Table 1-1 for endian mode indication) when the exception is taken. A special case for exception handlers is the system reset exception handler for both hard and soft reset for the G2_LE core. When the `core_tle` signal is negated at the time `core_hreset` is negated, the system exception handler of the

device enters into the big-endian mode. If MSR[ILE], MSR[LE], and HID2[LET] are subsequently set (during or after the reset routine has completed), a subsequent soft reset causes the system reset exception handler to be entered in true little-endian mode, potentially resulting in illegal instruction execution (if the beginning of the handler is written assuming big-endian code). Note that the reverse occurs for true little-endian mode.

The following assembly language code highlights register settings necessary when in big-endian mode coming out of hard reset and subsequently changing the processor state to true little-endian mode and setting the MSR[ILE], MSR[LE], and HID2[LET] bits. The first eight instructions of the system reset exception handler is written in big-endian format, in order to facilitate the mode switch. The rest of the reset handler is written in true little-endian format for the remaining supervisor or OS code. This reset code assumes that caching is not enabled out of reset. Due to the complexities involved with keeping the memory system coherent, it is strongly recommended not to change endianness at any other time once it is determined at hard reset.

```
.orig 0xFFFF0 0100 # default IP vector

# Begin HRESET_ handler with Big-Endian Mode

xor    r2,r2,r2        initialize register
xor    r1,r1,r1        # initialize register
oris   r2,r2,0x0800    # set bit in r2 for HID2[4]LET
mfspr  HID2,r2        # load HID2 setting LET bit
oris   r1,r1,0x0001    # set bit in r1 for MSR[15]ILE
ori    r1,r1,0x0001    # set bit in r1 for MSR[31]LE
mtmsr  r1              # load MSR setting ILE and LE bits
isync                          # wait for all instructions to complete

# End Big-Endian mode, True Little-Endian enabled
# modify the 8 Big-Endian instructions into valid True Little-Endian instructions
# True Little-Endian Mode
mfspr  SRR1,r1        # load the Machine State with LE enabled
xor    r0,r0,r0        # initialize register
oris   r0,r0,0x0001    # set Starting address at b'0001 0000
mfspr  SRR0,r0        # load the next instruction address

# whatever instructions the supervisor/OS wants.

rfi                                # return from HRESET_ interrupt routine
# End HRESET_ handler in True Little-Endian Mode
```

See Section 3.1.2, “Endian Modes and Byte Ordering,” for more information on the endian modes of the G2 and G2_LE cores.

5.5.2 Machine Check Exception (0x00200)

The G2 core conditionally initiates a machine check exception after detecting the assertion of the core_tea or core_mcp signals on the 60x bus (assuming the machine check is enabled with MSR[ME] = 1). The assertion of one of these signals indicates that a bus error

occurred and the system terminates the current transaction. One clock cycle after the signal is asserted, the data bus signals go to the high-impedance state; however, data entering the GPR or the cache is not invalidated. Note that if `HID0[EMCP]` is cleared, the core ignores the assertion of the `core_mcp` signal.

A machine check exception also occurs when an address or data parity error is detected on the bus and the address or data parity error is enabled in `HID0`. See Section 2.1.2.1, “Hardware Implementation Register 0 (`HID0`),” for more information.

Note that the G2 core makes no attempt to force recoverability on a machine check; however, it does guarantee that the machine check exception is always taken immediately upon request, with a nonpredicted address saved in `SRR0`, regardless of the current machine state. Because pending stores in the store queue (see Figure 7-4) are not canceled when a machine check exception occurs, two consecutive stores that result in the assertion of `core_tea` can cause the processor to checkstop. To prevent a checkstop in this case, a **sync** instruction must be placed between two stores that can result in assertion of `core_tea`.

Software can use the machine check exception in a recoverable mode to probe memory. For this case, a **sync**, load, **sync** instruction sequence is used. If the load access results in a system error (for example, the assertion of `core_tea`), the processor can handle this in a recoverable state. If the **sync** instruction is not used, a second access to the same address as the first load could cause the processor to enter the checkstop state.

If the `MSR[ME]` bit is set, the exception is recognized and handled; otherwise, the G2 core attempts to enter an internal checkstop. Note that the resulting machine check exception has priority over any exceptions caused by the instruction that generated the bus operation.

Machine check exceptions are only enabled when `MSR[ME] = 1`; this is described in Section 5.5.2.1, “Machine Check Exception Enabled (`MSR[ME] = 1`).” If `MSR[ME] = 0` and a machine check occurs, the processor enters the checkstop state; this is described in Section 5.5.2.2, “Checkstop State (`MSR[ME] = 0`).”

5.5.2.1 Machine Check Exception Enabled (`MSR[ME] = 1`)

When a machine check exception is taken, registers are updated as shown in Table 5-13.

When a machine check exception is taken, instruction execution for the handler begins at offset `0x00200` from the physical base address indicated by `MSR[IP]`.

In order to return to the main program, the exception handler should do the following:

1. `SRR0` and `SRR1` should be given the values to be used by the **rfi** instruction
2. Execute **rfi**

Table 5-13. Machine Check Exception—Register Settings

Register	Setting Description
SRR0	Set to the address of the next instruction that would have been completed in the interrupted instruction stream. Neither this instruction nor any others beyond it will have been completed. All preceding instructions will have been completed.
SRR1	0–11 Cleared 12 <u>core_mcp</u> —Machine check signal caused exception 13 <u>core_tea</u> —Transfer error acknowledge signal caused exception 14 <u>core_dpe</u> —Data parity error condition (and signal assertion) caused exception 15 <u>core_ape</u> —Address parity error condition (and signal assertion) caused exception 16–31 Loaded from MSR[16–31]
MSR	POW 0 FP 0 FE1 0 RI 0 TGPR 0 ME — CE 0 LE Set to value of ILE ILE — FE0 0 IP — EE 0 SE 0 IR 0 PR 0 BE 0 DR 0

Note: When a machine check exception is taken, the exception handler should set MSR[ME] as soon as it is practical to handle another core_tea assertion. Otherwise, subsequent core_tea assertions cause the processor to automatically enter the checkstop state.

5.5.2.2 Checkstop State (MSR[ME] = 0)

When the G2 core enters the checkstop state, it asserts the checkstop output signal, core_ckstp_out. The following events cause the G2 core to enter the checkstop state:

- Machine check exception occurs with MSR[ME] cleared
- External checkstop input, core_ckstp_in, is asserted.

When a processor is in the checkstop state, instruction processing is suspended and generally cannot be restarted without resetting the processor. The contents of all latches are frozen within two cycles upon entering the checkstop state so that the state of the processor can be analyzed as an aid in problem determination.

Note that not all processors that implement the PowerPC architecture provide the same level of error checking. The reasons a processor can enter checkstop state are implementation-dependent.

5.5.3 DSI Exception (0x00300)

A DSI exception occurs when no higher priority exception exists and a data memory access cannot be performed. The condition that caused the DSI exception can be determined by reading the DSISR register, a supervisor-level SPR (SPR18) that can be read by using the **mf spr** instruction. Bit settings are provided in Table 5-14. Table 5-14 also indicates the memory element that is saved to the DAR.

Table 5-14. DSI Exception—Register Settings

Register	Setting Description																																								
SRR0	Set to the effective address of the instruction that caused the exception.																																								
SRR1	0–15 Cleared 16–31 Loaded with MSR[16–31]																																								
MSR	<table border="0"> <tr> <td>POW</td><td>0</td> <td>FP</td><td>0</td> <td>FE1</td><td>0</td> <td>RI</td><td>0</td> </tr> <tr> <td>TGPR</td><td>0</td> <td>ME</td><td>—</td> <td>CE</td><td>—</td> <td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>ILE</td><td>—</td> <td>FE0</td><td>0</td> <td>IP</td><td>—</td> <td></td><td></td> </tr> <tr> <td>EE</td><td>0</td> <td>SE</td><td>0</td> <td>IR</td><td>0</td> <td></td><td></td> </tr> <tr> <td>PR</td><td>0</td> <td>BE</td><td>0</td> <td>DR</td><td>0</td> <td></td><td></td> </tr> </table>	POW	0	FP	0	FE1	0	RI	0	TGPR	0	ME	—	CE	—	LE	Set to value of ILE	ILE	—	FE0	0	IP	—			EE	0	SE	0	IR	0			PR	0	BE	0	DR	0		
POW	0	FP	0	FE1	0	RI	0																																		
TGPR	0	ME	—	CE	—	LE	Set to value of ILE																																		
ILE	—	FE0	0	IP	—																																				
EE	0	SE	0	IR	0																																				
PR	0	BE	0	DR	0																																				
DSISR	<table border="0"> <tr> <td>0</td> <td>Set if a load or store instruction results in a direct-store error exception due to a load or store instruction accesses a direct-store segment by setting a T bit.</td> </tr> <tr> <td>1</td> <td>Set by the data TLB miss exception handler if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared.</td> </tr> <tr> <td>2–3</td> <td>Cleared</td> </tr> <tr> <td>4</td> <td>Set if a memory access is not permitted by the page or BAT protection mechanism; otherwise cleared.</td> </tr> <tr> <td>5</td> <td>Set if the lwarx or stwcx. instruction is attempted to direct-store space</td> </tr> <tr> <td>6</td> <td>Set for a store operation and cleared for a load operation</td> </tr> <tr> <td>9</td> <td>G2_LE core only. Set when a data address breakpoint exception when the data (bit 29) in the DABR1 or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> • Write breakpoints enabled when DABR[30] is set • Read breakpoints enabled when DABR[31] is set </td> </tr> <tr> <td>7–31</td> <td>Cleared</td> </tr> </table>	0	Set if a load or store instruction results in a direct-store error exception due to a load or store instruction accesses a direct-store segment by setting a T bit.	1	Set by the data TLB miss exception handler if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared.	2–3	Cleared	4	Set if a memory access is not permitted by the page or BAT protection mechanism; otherwise cleared.	5	Set if the lwarx or stwcx . instruction is attempted to direct-store space	6	Set for a store operation and cleared for a load operation	9	G2_LE core only. Set when a data address breakpoint exception when the data (bit 29) in the DABR1 or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> • Write breakpoints enabled when DABR[30] is set • Read breakpoints enabled when DABR[31] is set 	7–31	Cleared																								
0	Set if a load or store instruction results in a direct-store error exception due to a load or store instruction accesses a direct-store segment by setting a T bit.																																								
1	Set by the data TLB miss exception handler if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared.																																								
2–3	Cleared																																								
4	Set if a memory access is not permitted by the page or BAT protection mechanism; otherwise cleared.																																								
5	Set if the lwarx or stwcx . instruction is attempted to direct-store space																																								
6	Set for a store operation and cleared for a load operation																																								
9	G2_LE core only. Set when a data address breakpoint exception when the data (bit 29) in the DABR1 or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> • Write breakpoints enabled when DABR[30] is set • Read breakpoints enabled when DABR[31] is set 																																								
7–31	Cleared																																								
DAR	Set to the effective address of a memory element as described in the following list: <ul style="list-style-type: none"> • A byte in the first word accessed in the page that caused the DSI exception, for a byte, half word, or word memory access. • A byte in the first word accessed in the BAT area that caused the DSI exception for a byte, half word, or word access to a BAT area. • A byte in the block that caused the exception for icbi, dcbz, dcbst, dcbf, or dcbi¹ instructions. • The EA that causes a data breakpoint for the G2_LE core. • Any EA in the memory range addressed (for direct-store exceptions). 																																								

¹ The **dcbi** instruction should never be used on the G2 core.

DSI exceptions can occur for any of the following reasons:

- The instruction is not supported for the type of memory addressed
- Any access to a direct-store segment ($SR[T] = 1$)
- The attempted access violates the memory protection defined by $SR[Ks, Kp]$, $PTE[PP]$, or $DBAT_n[PP]$.

Note that the OEA specifies an additional case that may cause a DSI exception—when an effective address for a load, store, or cache operation cannot be translated by the TLBs. On the G2 core, this condition causes a TLB miss exception instead. These scenarios are common among all processors that implement the PowerPC architecture. The following additional scenarios can cause a DSI exception in the G2 core:

- A bus error indicates crossing from a direct-store segment to a memory segment

- The execution of any load/store instruction to a direct-store segment ($SR[T] = 1$)
- A data access crosses from a memory segment ($SR[T] = 0$) into a direct-store segment ($SR[T] = 1$)

Finally, the G2_LE core causes a DSI exception when either the DABR or DABR2 is enabled and the address of an access matches with the value in the CEA field and the breakpoint is enabled for the type of access (read or write) in DABR/DABR2. See Chapter 11, “Debug Features,” and Section 2.1.2.15, “Data Address Breakpoint Register (DABR and DABR2)—G2_LE Only,” for more information.

DSI exceptions can be generated by load/store instructions and cache control instructions (**dcbi**, **dcbz**, **dcbst**, and **dcbf**). Note that the **dcbi** instruction should never be used on the G2 core.

The G2 core supports the crossing of page boundaries. However, if the second page has a translation error or protection violation associated with it, the G2 core takes the DSI exception in the middle of the instruction. In this case, the data address register (DAR) always points to a byte address in the first word of the offending page.

If an **stwcx.** instruction has an effective address for which a normal store operation would cause a DSI exception, the G2 core takes the DSI exception without checking for the reservation.

If the XER indicates that the byte count for an **lswi** or **stswi** instruction is zero, a DSI exception does not occur, regardless of the effective address.

The condition that caused the exception is defined in the DSISR. These conditions also use the data address register (DAR) as shown in Table 5-14.

When a DSI exception is taken, instruction execution for the handler begins at offset 0x00300 from the physical base address indicated by MSR[IP].

The architecture permits certain instructions to be partially executed when they cause a DSI exception. These are as follows:

- Load multiple or load string instructions—some registers in the range of registers to be loaded may have been loaded.
- Store multiple or store string instructions—some bytes of memory in the range addressed may have been updated.

In these cases, the number of registers and amount of memory altered are instruction- and boundary-dependent. However, memory protection is not violated. Furthermore, if some of the data accessed is in direct-store space ($SR[T] = 1$) and the instruction is not supported for direct-store accesses, the locations in direct-store space are not accessed.

For update forms, the update register (**rA**) is not altered.

5.5.4 ISI Exception (0x00400)

The ISI exception is implemented as it is defined by the PowerPC architecture. An ISI exception occurs when no higher priority exception exists and an attempt to fetch the next instruction fails for any of the following reasons:

- If an instruction TLB miss fails to find the desired PTE, then a page fault is synthesized. The ITLB miss handler branches to the ISI exception handler to retrieve the translation from a storage device.
- An attempt is made to fetch an instruction from a direct-store segment while instruction translation is enabled ($MSR[IR] = 1$)
- An attempt is made to fetch an instruction from no-execute memory
- An attempt is made to fetch an instruction from guarded memory when $MSR[IR] = 1$
- The fetch access violates memory protection

Register settings for this exception are described in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

When an ISI exception is taken, instruction execution for the handler begins at offset 0x00400 from the physical base address indicated by $MSR[IP]$.

5.5.5 External Interrupt (0x00500)

An external interrupt is signaled to the G2 core by the assertion of the $\overline{core_int}$ signal as described in Section 8.3.9.1, “External Interrupt ($\overline{core_int}$)—Input.” The interrupt may not be recognized if a higher priority exception occurs simultaneously or if the $MSR[EE]$ bit is cleared when $\overline{core_int}$ is asserted.

After the $\overline{core_int}$ is recognized, the G2 core generates a recoverable halt to instruction completion. The G2 core allows the next instruction in program order to complete, including handling any exceptions that instruction may generate. However, the G2 core blocks subsequent instructions from completing and allows any outstanding stores to occur to system memory. If any other exceptions are encountered in this process, they are taken first and the external interrupt is delayed until a recoverable halt is achieved. At this time, the G2 core saves the state information and takes the external interrupt as defined by the PowerPC architecture.

The register settings for the external interrupt are shown in Table 5-15.

Table 5-15. External Interrupt—Register Settings

Register	Setting																																								
SRR0	Set to the effective address of the instruction that the processor would have attempted to execute next if no interrupt conditions were present.																																								
SRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]																																								
MSR	<table style="width: 100%; border: none;"> <tr> <td>POW</td><td>0</td> <td>FP</td><td>0</td> <td>FE1</td><td>0</td> <td>RI</td><td>0</td> </tr> <tr> <td>TGPR</td><td>0</td> <td>ME</td><td>—</td> <td>CE</td><td>—</td> <td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>ILE</td><td>—</td> <td>FE0</td><td>0</td> <td>IP</td><td>—</td> <td></td><td></td> </tr> <tr> <td>EE</td><td>0</td> <td>SE</td><td>0</td> <td>IR</td><td>0</td> <td></td><td></td> </tr> <tr> <td>PR</td><td>0</td> <td>BE</td><td>0</td> <td>DR</td><td>0</td> <td></td><td></td> </tr> </table>	POW	0	FP	0	FE1	0	RI	0	TGPR	0	ME	—	CE	—	LE	Set to value of ILE	ILE	—	FE0	0	IP	—			EE	0	SE	0	IR	0			PR	0	BE	0	DR	0		
POW	0	FP	0	FE1	0	RI	0																																		
TGPR	0	ME	—	CE	—	LE	Set to value of ILE																																		
ILE	—	FE0	0	IP	—																																				
EE	0	SE	0	IR	0																																				
PR	0	BE	0	DR	0																																				

When an external interrupt is taken, instruction execution for the handler begins at offset 0x00500 from the physical base address indicated by MSR[IP].

The G2 core only recognizes the interrupt condition ($\overline{\text{core_int}}$ asserted) if the MSR[EE] bit is set; it ignores the interrupt condition if the MSR[EE] bit is cleared. To guarantee that the external interrupt is taken, the $\overline{\text{core_int}}$ signal must be held asserted until the G2 core takes the interrupt. If the $\overline{\text{core_int}}$ signal is negated before the interrupt is taken, the G2 core is not guaranteed to take an external interrupt. The interrupt handler must send a command to the device that asserted $\overline{\text{core_int}}$, acknowledging the interrupt and instructing the device to negate $\overline{\text{core_int}}$ before the handler re-enables recognition of external interrupts.

5.5.6 Alignment Exception (0x00600)

This section describes conditions that can cause alignment exceptions in the G2 core. The G2 core implements the alignment exception as it is defined in the PowerPC architecture. For information on bit settings and how exception conditions are detected, refer to the *Programming Environments Manual*. Note that the PowerPC architecture allows individual processors to determine whether an exception is required to handle various alignment conditions.

Similar to DSI exceptions, alignment exceptions use the SRR0 and SRR1 to save the machine state and the DSISR to determine the source of the exception. The G2 core initiates an alignment exception when it detects any of the following conditions:

- The operand of a floating-point load or store operation is not word-aligned
- The operand of an **lmw**, **stmw**, **lwarx**, or **stwcx**. instruction is not word-aligned.
- A multiple or string access is attempted with the MSR[LE] bit set
- The operand of a floating-point load or store operation is to a direct-store segment
- The operand of an elementary, multiple or string load or store crosses a segment boundary with a change to the direct-store attribute (T bit different).

- The operand of a **eciwx** or **ecowx** instruction is not aligned
- The operand of a **dcbz** instruction is in a page that is write-through or caching-inhibited

Note that although the MPC603e processor generates an alignment exception for a misaligned little-endian access (MSR[LE] = 1), the G2 core does not.

The register settings for alignment exceptions are shown in Table 5-15.

The architecture does not support the use of a misaligned EA by **lwarx** or **stwcx.** instructions. If one of these instructions specifies a misaligned EA, the exception handler should not emulate the instruction, but should treat the occurrence as a programming error.

Table 5-16. Alignment Interrupt—Register Settings

Register	Setting
SRR0	Set to the effective address of the instruction that caused the exception
SRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]
MSR	POW 0 FP 0 FE1 0 RI 0 TGPR 0 ME — CE — LE Set to value of ILE ILE — FE0 0 IP — EE 0 SE 0 IR 0 PR 0 BE 0 DR 0
DSISR	0–11 Cleared 12–13 Cleared. (Note that these bits can be set by several 64-bit PowerPC instructions that are not supported in the G2 core.) 14 Cleared 15–16 For instructions that use register indirect with index addressing—set to bits 29–30 of the instruction For instructions that use register indirect with immediate index addressing—cleared 17 For instructions that use register indirect with index addressing—set to bit 25 of the instruction For instructions that use register indirect with immediate index addressing—set to bit 5 of the instruction 18–21 For instructions that use register indirect with index addressing—set to bits 21–24 of the instruction For instructions that use register indirect with immediate index addressing—set to bits 1–4 of the instruction 22–26 Set to bits 6–10 (identifying either the source or destination) of the instruction. Undefined for dcbz. 27–31 Set to bits 11–15 of the instruction (rA). Set to either bits 11–15 of the instruction or to any register number not in the range of registers loaded by a valid form instruction, for lmw , lswi , and lswx instructions. Otherwise undefined.
DAR	Set to the EA of the data access as computed by the instruction causing the alignment exception. When the operand of an lmw , stmw , lwarx , or stwcx. instruction is not word-aligned, that address value + 4 is stored into the DAR.

5.5.6.1 Integer Alignment Exceptions

The G2 core is optimized for load and store operations that are aligned on natural boundaries. Operations that are not naturally aligned may suffer performance degradation,

depending on the type of operation, the boundaries crossed, and the mode that the processor is in during execution. More specifically, these operations may either cause an alignment exception or they may cause the processor to break the memory access into multiple, smaller accesses with respect to the cache and the memory subsystem.

The G2 core can initiate an alignment exception for the access shown in Table 5-17. In this case, the appropriate range check is performed before the instruction begins execution. As a result, if an alignment exception is taken, it is guaranteed that no portion of the instruction has been executed.

Table 5-17. Access Types

MSR[DR]	SR[T]	Access Type
1	0	Page-address translation access

A page-address translation access occurs when MSR[DR] is set, SR[T] is cleared, and there is not a match in the BAT. Note the following points:

- The following is true for all loads and stores except strings/multiples:
 - Byte operands never cause an alignment exception
 - Half-word operands can cause an alignment exception if the EA ends in 0xFFFF
 - Word operands can cause an alignment exception if the EA ends in 0xFFD–FFF
 - Double-word operands cause an alignment exception if the EA ends in 0xFF9–FFF
- The **dcbz** instruction causes an alignment exception if the access is to a page or block with the W (write-through) or I (cache-inhibit) bit set in the TLB or BAT, respectively.

A misaligned memory access that does not cause an alignment exception will not perform as well as an aligned access of the same type. The resulting performance degradation due to misaligned accesses depends on how well each individual access behaves with respect to the memory hierarchy. At a minimum, additional cache access cycles are required that can delay other processor resources from using the cache. More dramatically, for an access to a noncacheable page, each discrete access involves individual processor bus operations that reduce the effective bandwidth of that bus.

Finally, note that when the G2 core is in page address translation mode, there is no special handling for accesses that fall into BAT regions.

5.5.6.2 Load/Store Multiple Alignment Exceptions

Most alignment exceptions store the address as computed by the instruction in the DAR. However, when the operand of an **lmw**, **stmw**, **lwarx**, or **stwcx**. instruction is not word-aligned that address value + 4 is stored into the DAR.

5.5.7 Program Exception (0x00700)

The G2 core implements the program exception as it is defined by the PowerPC architecture (OEA). A program exception occurs when no higher priority exception exists and one or more of the exception conditions defined in the OEA occur.

When a program exception is taken, instruction execution for the handler begins at offset 0x00700 from the physical base address indicated by MSR[IP]. The exception conditions are as follows:

- Floating-point enabled exception—These exceptions correspond to IEEE-defined exception conditions, such as overflows, and divide by zeros that may occur during the execution of a floating-point arithmetic instruction. As a group, these exceptions are enabled by the FE0 and FE1 bits in the MSR. Individual conditions are enabled by specific bits in the FPSCR. For general information about this exception, see the *Programming Environments Manual*. For more information about how these exceptions are implemented in the G2 core, see Section 5.5.7.1, “IEEE Floating-Point Exception Program Exceptions.”
- Illegal instruction—An illegal instruction program exception is generated when execution of an instruction is attempted with an illegal opcode or illegal combination of opcode and extended opcode fields (including PowerPC instructions not implemented in the G2 core). These do not include those optional instructions treated as no-ops.
- Privileged instruction—A privileged instruction type program exception is generated when the execution of a privileged instruction is attempted and the MSR register user privilege bit, MSR[PR], is set. In the G2 core, this exception is generated for **mtspr** or **mfspr** with an invalid SPR field if SPR[0] = 1 and MSR[PR] = 1. This may not be true for all processors that implement the PowerPC architecture.
- Trap—A trap type program exception is generated when any of the conditions specified in a trap instruction is met.

5.5.7.1 IEEE Floating-Point Exception Program Exceptions

Floating-point exceptions are signaled by condition bits set in the floating-point status and control register (FPSCR). They can cause the system floating-point enabled exception handler to be invoked. The G2 core handles all floating-point exceptions precisely. The G2 core implements the FPSCR as it is defined by the PowerPC architecture; for more information about the FPSCR, see the *Programming Environments Manual*.

Floating-point operations that change exception sticky bits in the FPSCR may suffer a performance penalty. When an exception is disabled in the FPSCR and MSR[FE] = 0, updates to the FPSCR exception sticky bits are serialized at the completion stage. This serialization may result in a one- or two-cycle execution delay. The penalty is incurred only

when the exception bit is changed and not on subsequent operations with the same exception. See Chapter 7, “Instruction Timing,” for a full description of completion serialization.

When an exception is enabled in the FPSCR, the instruction traps to the emulation trap exception vector without updating the FPSCR or the target FPR. The emulation trap exception handler is required to complete the instruction. The emulation trap exception handler is invoked regardless of the FE setting in the MSR.

The two IEEE floating-point imprecise modes, defined by the PowerPC architecture when $MSR[FE0] \neq MSR[FE1]$, are treated as precise exceptions (that is, $MSR[FE0] = MSR[FE1] = 1$). This is regardless of the setting of MSR[NI].

For the highest and most predictable floating-point performance, all exceptions should be disabled in the FPSCR and MSR. For more information about the program exception, see the *Programming Environments Manual*.

5.5.7.2 Illegal, Reserved, and Unimplemented Instructions Program Exceptions

In accordance with the PowerPC architecture, the G2 core considers all instructions defined for 64-bit implementations and unimplemented optional instructions, such as **fsqrt**, **eciwx**, and **ecowx** as illegal and takes a program exception when one of these instructions is encountered. Likewise, if a supervisor-level instruction is encountered when the processor is in user-level mode, a privileged instruction-type program exception is taken.

5.5.8 Floating-Point Unavailable Exception (0x00800)

A floating-point unavailable exception occurs when no higher priority exception exists, an attempt is made to execute a floating-point instruction (including floating-point load, store, and move instructions), and the floating-point available bit in the MSR is disabled ($MSR[FP] = 0$). Register settings for this exception are described in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

When a floating-point unavailable exception is taken, instruction execution for the handler begins at offset 0x00800 from the physical base address indicated by MSR[IP].

5.5.9 Decrementer Exception (0x00900)

The G2 core implements the decrementer interrupt exception as it is defined in the PowerPC architecture. A decrementer exception request is made when the decrementer counts down through zero. The request is held until there are no higher priority exceptions and $MSR[EE] = 1$. At this point the decrementer exception is taken. If multiple decrementer exception requests are received before the first can be reported, only one exception is reported. The occurrence of a decrementer exception cancels the request.

Register settings for this exception are described in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

When a decremter exception is taken, instruction execution for the handler begins at offset 0x00900 from the physical base address indicated by MSR[IP].

5.5.10 Critical Interrupt Exception (0x00A00)—G2_LE Only

A critical interrupt is signaled to the G2_LE core by the assertion of the $\overline{\text{core_int}}$ signal as described in Section 8.3.9.2, “Critical Interrupt ($\overline{\text{core_cint}}$)—Input: G2_LE Core-Only.” The interrupt may not be recognized if a higher priority exception occurs simultaneously or if the MSR[CE] bit is cleared when $\overline{\text{core_cint}}$ is asserted.

The following events occur when the G2_LE recognizes the assertion of $\overline{\text{core_cint}}$:

- Multi-cycle instructions not in the completion stage are terminated
- Outstanding load or store instructions that have not been completed are terminated
- Any outstanding page table search activity is terminated
- The effective address for resuming program execution is saved into CSRR0
- The contents of MSR are saved into CSRR1
- The MSR register is loaded with all zeros except the IP, ILE, and ME bits which remain unchanged
- Exception processing starts at offset value 0x00A00 from the physical base address indicated by MSR[IP]

Some types of instructions (for example load multiple/string and floating-point instructions) cause additional interrupt recognition latency. Timing critical applications must consider these instruction execution latencies in calculating worst-case interrupt recognition latency.

Upon returning from a critical interrupt handler routine the core restarts any terminated or uncompleted instructions, including terminated load multiple or load string instructions. Note that these restarted load instructions may cause side-effects on peripheral devices that have auto-decremter or status bit changes caused by the subsequent load accesses.

The register settings for the critical interrupt are shown in Table 5-15.

Table 5-18. Critical Interrupt—Register Settings

Register	Setting																																								
CSRR0	Set to the effective address of the instruction that the processor would have attempted to execute next if no interrupt conditions were present.																																								
CSRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]																																								
MSR	<table> <tr> <td>POW</td><td>0</td> <td>FP</td><td>0</td> <td>FE1</td><td>0</td> <td>RI</td><td>0</td> </tr> <tr> <td>TGPR</td><td>0</td> <td>ME</td><td>—</td> <td>CE</td><td>0</td> <td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>ILE</td><td>—</td> <td>FE0</td><td>0</td> <td>IP</td><td>—</td> <td></td><td></td> </tr> <tr> <td>EE</td><td>0</td> <td>SE</td><td>0</td> <td>IR</td><td>0</td> <td></td><td></td> </tr> <tr> <td>PR</td><td>0</td> <td>BE</td><td>0</td> <td>DR</td><td>0</td> <td></td><td></td> </tr> </table>	POW	0	FP	0	FE1	0	RI	0	TGPR	0	ME	—	CE	0	LE	Set to value of ILE	ILE	—	FE0	0	IP	—			EE	0	SE	0	IR	0			PR	0	BE	0	DR	0		
POW	0	FP	0	FE1	0	RI	0																																		
TGPR	0	ME	—	CE	0	LE	Set to value of ILE																																		
ILE	—	FE0	0	IP	—																																				
EE	0	SE	0	IR	0																																				
PR	0	BE	0	DR	0																																				

The G2_LE core only recognizes the interrupt condition (core_cint asserted) if the MSR[CE] bit is set; it ignores the interrupt condition if the MSR[CE] bit is cleared. To guarantee that the critical interrupt is taken, the core_cint signal must be held asserted until the G2_LE core takes the interrupt. If the core_cint signal is negated before the interrupt is taken, the G2_LE core is not guaranteed to take a critical interrupt. The interrupt handler must send a command to the device that asserted core_cint, acknowledging the interrupt and instructing the device to negate core_cint before the handler re-enables recognition of critical interrupts.

The additional SPRG4–7 registers on the G2_LE core can reduce overall latency for critical interrupts, as fewer GPRs need to be saved upon entering a critical interrupt handler routine. The G2_LE core also implements the **rftci** instruction for specifically returning from critical interrupt routines and restoring the processor state from CSRR0 and CSRR1.

5.5.11 System Call Exception (0x00C00)

The G2 core implements the system call exception as it is defined by the PowerPC architecture. A system call exception request is made when a system call (**sc**) instruction is completed. If no higher priority exception exists, the system call exception is taken, with SRR0 being set to the EA of the instruction following the **sc** instruction. Register settings for this exception are described in Chapter 6, “Exceptions,” in the *Programming Environments Manual*.

When a system call exception is taken, instruction execution for the handler begins at offset 0x00C00 from the physical base address indicated by MSR[IP].

5.5.12 Trace Exception (0x00D00)

The trace exception is taken under one of the following conditions:

- When MSR[SE] is set, a single-step instruction trace exception is taken when no higher priority exception exists and any instruction (other than **rfti**, **rftci**, **mtmsr**, or **isync**) is successfully completed. Note that other processors will take the trace

exception on **isync** instructions (when MSR[SE] is set); the G2 core does not take the trace exception on **isync** instructions. Single-step instruction trace mode is described in Section 5.5.12.1, “Single-Step Instruction Trace Mode.”

- When MSR[BE] is set, the branch trace exception is taken after each branch instruction is completed.
- The G2 core deviates from the architecture by not taking trace exceptions on **isync** instructions. Single-step instruction trace mode is described in Section 5.5.12.2, “Branch Trace Mode.”

Successful completion implies that the instruction caused no other exceptions. A trace exception is never taken for an **sc** or trap instruction that takes a trap exception.

MSR[SE] and MSR[BE] are cleared when the trace exception is taken. In the normal use of this function, MSR[SE] and MSR[BE] are restored when the exception handler returns to the interrupted program using an **rfi** instruction.

Register settings for the trace mode are described in Table 5-19.

Table 5-19. Trace Exception—Register Settings

Register	Setting Description																																								
SRR0	Set to the address of the instruction following the one for which the trace exception was generated.																																								
SRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]																																								
MSR	<table border="0"> <tr> <td>POW</td><td>0</td> <td>FP</td><td>0</td> <td>FE1</td><td>0</td> <td>RI</td><td>0</td> </tr> <tr> <td>TGPR</td><td>0</td> <td>ME</td><td>—</td> <td>CE</td><td>—</td> <td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>ILE</td><td>—</td> <td>FE0</td><td>0</td> <td>IP</td><td>—</td> <td></td><td></td> </tr> <tr> <td>EE</td><td>0</td> <td>SE</td><td>0</td> <td>IR</td><td>0</td> <td></td><td></td> </tr> <tr> <td>PR</td><td>0</td> <td>BE</td><td>0</td> <td>DR</td><td>0</td> <td></td><td></td> </tr> </table>	POW	0	FP	0	FE1	0	RI	0	TGPR	0	ME	—	CE	—	LE	Set to value of ILE	ILE	—	FE0	0	IP	—			EE	0	SE	0	IR	0			PR	0	BE	0	DR	0		
POW	0	FP	0	FE1	0	RI	0																																		
TGPR	0	ME	—	CE	—	LE	Set to value of ILE																																		
ILE	—	FE0	0	IP	—																																				
EE	0	SE	0	IR	0																																				
PR	0	BE	0	DR	0																																				

Note that a trace or instruction address breakpoint exception condition generates a soft stop instead of an exception if soft stop has been enabled by the JTAG/COP logic. If trace and breakpoint conditions occur simultaneously, the breakpoint conditions receive higher priority.

When a trace exception is taken, instruction execution for the handler begins at offset 0x00D00 from the base address indicated by MSR[IP].

5.5.12.1 Single-Step Instruction Trace Mode

The single-step instruction trace mode is enabled by setting MSR[SE]. Encountering the single-step breakpoint causes the following action—trap to address vector 0x00D00.

The single-step trace action traps after an instruction execution and completion.

5.5.12.2 Branch Trace Mode

The branch trace mode is enabled by setting MSR[BE]. Encountering the branch trace breakpoint causes the following action—trap to interrupt vector 0x00D00.

The branch trace action is to trap after the completion of any branch instruction whenever MSR[BE] is set.

5.5.13 Instruction TLB Miss Exception (0x01000)

When the effective address for an instruction load, store, or cache operation cannot be translated by the ITLB, an instruction TLB miss exception is generated. Register settings for the instruction and data TLB miss exceptions are described in Table 5-20.

If the instruction TLB miss exception handler fails to find the desired PTE, then a page fault must be synthesized. The handler must restore the machine state and clear MSR[TGPR] before invoking the ISI exception (0x00400).

Software table search operations are discussed in Chapter 6, “Memory Management.”

When an instruction TLB miss exception is taken, instruction execution for the handler begins at offset 0x01000 from the physical base address indicated by MSR[IP].

5.5.14 Data TLB Miss on Load Exception (0x01100)

When the effective address for a data load or cache operation cannot be translated by the DTLB, a data TLB miss on load exception is generated. Register settings for the instruction and data TLB miss exceptions are described in Table 5-20.

If a data TLB miss exception handler fails to find the desired PTE, then a page fault must be synthesized. The handler must restore the machine state and clear MSR[TGPR] before invoking the DSI exception (0x00300).

Software table search operations are discussed in Chapter 6, “Memory Management.”

When a data TLB miss on load exception is taken, instruction execution for the handler begins at offset 0x01100 from the physical base address indicated by MSR[IP].

Table 5-20. Instruction and Data TLB Miss Exceptions—Register Settings

Register	Setting Description
SRR0	Set to the address of the next instruction to be executed in the program for which the TLB miss exception was generated.
SRR1	0–3 Loaded from condition register CR0 field 4–11 Cleared 12 KEY. Key for TLB miss (SR[Ks] or SR[Kp], depending on whether the access is a user or supervisor access). 13 D/I. Data or instruction access. 0 = Data TLB miss 1 = Instruction TLB miss 14 WAY. Next TLB set to be replaced (set per LRU). 0 = Replace TLB associativity set 0 1 = Replace TLB associativity set 1 15 S/L. Store or load data access. 0 = Data TLB miss on load 1 = Data TLB miss on store (or C = 0) 16–31 Loaded from MSR[16–31]
MSR	POW 0 FP 0 FE1 0 RI 0 TGPR 1 ME — CE — LE Set to value of ILE ILE — FE0 0 IP — EE 0 SE 0 IR 0 PR 0 BE 0 DR 0

5.5.15 Data TLB Miss on Store Exception (0x01200)

When the effective address for a data store or cache operation cannot be translated by the DTLB, a data TLB miss on store exception is generated. The data TLB miss on store exception is also taken when the changed bit (C = 0) for a DTLB entry needs to be updated for a store operation. Register settings for the instruction and data TLB miss exceptions are described in Table 5-20.

If a data TLB miss exception handler fails to find the desired PTE, then a page fault must be synthesized. The handler must restore the machine state and clear MSR[TGPR] before invoking the DSI exception (0x00300).

Software table search operations are discussed in Chapter 6, “Memory Management.”

When a data TLB miss on store exception is taken, instruction execution for the handler begins at offset 0x01200 from the physical base address indicated by MSR[IP].

5.5.16 Instruction Address Breakpoint Exception (0x01300)

The instruction address breakpoint is controlled by the IABR and IABR2 special purpose register. Bits [0–29] of IABR and IABR2 holds an effective address to which each instruction’s address is compared. The exception is enabled by setting bit 30 in the IABR and IABR2. The exception is taken when an instruction breakpoint address matches on the

next instruction to complete. The instruction tagged with the match is not completed before the instruction address breakpoint exception is taken.

The breakpoint action can be trapped to interrupt vector 0x01300 (default).

Note that the G2_LE core also has a second instruction address breakpoint register, IABR2, that functions identically to IABR, and allows for two instruction breakpoints to be enabled.

The bit settings for when an instruction address breakpoint exception is taken are shown in Table 5-21.

Table 5-21. Instruction Address Breakpoint Exception—Register Settings

Register	Setting Description																																								
SRR0	Set to the address of the next instruction to be executed in the program for which the TLB miss exception was generated.																																								
SRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]																																								
MSR	<table> <tr> <td>POW</td><td>0</td> <td>FP</td><td>0</td> <td>FE1</td><td>0</td> <td>RI</td><td>0</td> </tr> <tr> <td>TGPR</td><td>0</td> <td>ME</td><td>—</td> <td>CE</td><td>—</td> <td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>ILE</td><td>—</td> <td>FE0</td><td>0</td> <td>IP</td><td>—</td> <td></td><td></td> </tr> <tr> <td>EE</td><td>0</td> <td>SE</td><td>0</td> <td>IR</td><td>0</td> <td></td><td></td> </tr> <tr> <td>PR</td><td>0</td> <td>BE</td><td>0</td> <td>DR</td><td>0</td> <td></td><td></td> </tr> </table>	POW	0	FP	0	FE1	0	RI	0	TGPR	0	ME	—	CE	—	LE	Set to value of ILE	ILE	—	FE0	0	IP	—			EE	0	SE	0	IR	0			PR	0	BE	0	DR	0		
POW	0	FP	0	FE1	0	RI	0																																		
TGPR	0	ME	—	CE	—	LE	Set to value of ILE																																		
ILE	—	FE0	0	IP	—																																				
EE	0	SE	0	IR	0																																				
PR	0	BE	0	DR	0																																				

The default breakpoint action is to trap before the execution of the matching instruction.

Table 5-22 shows the priority of actions taken when more than one mode is enabled for the same instruction.

Note that a trace or instruction address breakpoint exception condition generates a soft stop instead of an exception if soft stop has been enabled by the JTAG/COP logic. If trace and breakpoint conditions occur simultaneously, the breakpoint conditions receive higher priority.

The G2 core requires that an **mtspr** instruction that updates the IABR be followed by a context-synchronizing instruction. If the **mtspr** instruction enables the instruction address breakpoint exception, the context-synchronizing instruction cannot generate a breakpoint response. The G2 core also cannot block a breakpoint response on the context-synchronizing instruction if the breakpoint was disabled by the **mtspr** instruction. See “Synchronization Requirements for Special Registers and TLBs” in Chapter 2, “Register Set,” in the *Programming Environments Manual*, for more information on this requirement.

Table 5-22. Breakpoint Action for Multiple Modes Enabled for the Same Address

IABR[IE]	MSR[BE]	MSR[SE]	First Action	Next Action	Comments
1	1	0	Instruction address breakpoint	Trace (branch)	Enabling both modes is useful only if both trace and address breakpoint interrupts are needed.
1	0	1	Instruction address breakpoint	Trace (single-step)	Enabling both modes is useful only if different breakpoint actions are required.
0	1	1	Trace (branch)	None	The action for branch trace and single-step trace is the same. Enabling both trace modes is redundant except for hard stop on branches.
1	1	1	Instruction address breakpoint	Trace	Enabling all modes is redundant. This entry is for clarification only.

Section 2.1.2.14, “Instruction Address Breakpoint Registers (IABR and IABR2),” and Chapter 11, “Debug Features,” provide more information about the instruction breakpoint facility.

5.5.17 System Management Interrupt (0x01400)

The system management interrupt behaves like an external interrupt except for the signal asserted and the vector taken. A system management interrupt is signaled to the G2 core by the assertion of the core_smi signal. The interrupt may not be recognized if a higher priority exception occurs simultaneously or if MSR[EE] is cleared when core_smi is asserted. Note that core_smi takes priority over core_int if they are recognized simultaneously.

After the core_smi is detected (and provided that MSR[EE] is set), the G2 core generates a recoverable halt to instruction completion. The G2 core requires the next instruction in program order to complete or except, block completion of any following instructions, and allow the completed store queue to drain (see Section 7.1, “Terminology and Conventions,” for the definition). If any higher priority exceptions are encountered in this process, they are taken first and the system management interrupt is delayed until a recoverable halt is achieved. At this time the G2 core saves state information and takes the system management interrupt.

The register settings for the external interrupt exception are shown in Table 5-23.

Table 5-23. System Management Interrupt—Register Settings

Register	Setting Description																																								
SRR0	Set to the effective address of the instruction that the processor would have attempted to complete next if no interrupt conditions were present.																																								
SRR1	0–15 Cleared 16–31 Loaded from MSR[16–31]																																								
MSR	<table border="0"> <tr> <td>POW</td><td>0</td> <td>FP</td><td>0</td> <td>FE1</td><td>0</td> <td>RI</td><td>0</td> </tr> <tr> <td>TGPR</td><td>0</td> <td>ME</td><td>—</td> <td>CE</td><td>—</td> <td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>ILE</td><td>—</td> <td>FE0</td><td>0</td> <td>IP</td><td>—</td> <td></td><td></td> </tr> <tr> <td>EE</td><td>0</td> <td>SE</td><td>0</td> <td>IR</td><td>0</td> <td></td><td></td> </tr> <tr> <td>PR</td><td>0</td> <td>BE</td><td>0</td> <td>DR</td><td>0</td> <td></td><td></td> </tr> </table>	POW	0	FP	0	FE1	0	RI	0	TGPR	0	ME	—	CE	—	LE	Set to value of ILE	ILE	—	FE0	0	IP	—			EE	0	SE	0	IR	0			PR	0	BE	0	DR	0		
POW	0	FP	0	FE1	0	RI	0																																		
TGPR	0	ME	—	CE	—	LE	Set to value of ILE																																		
ILE	—	FE0	0	IP	—																																				
EE	0	SE	0	IR	0																																				
PR	0	BE	0	DR	0																																				

When a system management interrupt is taken, instruction execution for the handler begins at offset 0x01400 from the physical base address indicated by MSR[IP].

The G2 core recognizes the interrupt condition ($\overline{\text{core_smi}}$ asserted) only if the MSR[EE] bit is set; otherwise, the interrupt condition is ignored. To guarantee that the external interrupt is taken, the $\overline{\text{core_smi}}$ signal must be held active until the G2 core takes the interrupt. If the $\overline{\text{core_smi}}$ signal is negated before the interrupt is taken, the G2 core is not guaranteed to take a system management interrupt. The interrupt handler must send a command to the device that asserted $\overline{\text{core_smi}}$, acknowledging the interrupt and instructing the device to negate $\overline{\text{core_smi}}$.

Chapter 6

Memory Management

This chapter describes the G2 core implementation of the memory management unit (MMU) specifications provided by the PowerPC operating environment architecture (OEA). The MMU implementation of the G2 core is the same as that of the MPC603e microprocessor. However, the G2_LE core implements four additional IBAT entries and four additional DBAT entries.

The primary function of the MMU in a processor of this family is the translation of logical (effective) addresses to physical addresses (referred to as real addresses in the architecture specification) for memory accesses, and I/O accesses (I/O accesses are assumed to be memory-mapped). In addition, the MMU provides access protection on a segment, block, or page basis. This chapter describes the specific hardware used to implement the MMU model of the OEA in the core. Refer to Chapter 7, “Memory Management,” in the *Programming Environments Manual* for a complete description of the conceptual model.

Two general types of accesses generated by processors that implement the PowerPC architecture require address translation—instruction accesses, and data accesses to memory generated by load and store instructions. Generally, the address translation mechanism is defined in terms of segment descriptors and page tables defined by the PowerPC architecture for locating the effective-to-physical address mapping for instruction and data accesses. The segment information translates the effective address to an interim virtual address and the page table information translates the virtual address to a physical address.

The segment descriptors, used to generate the interim virtual addresses, are stored as on-chip segment registers on 32-bit implementations (such as the G2 core). In addition, two translation lookaside buffers (TLBs) are implemented on the core to keep recently-used page address translations on-chip. Although the OEA describes one MMU (conceptually), the G2 core hardware maintains separate TLBs and table search resources for instruction and data accesses that can be accessed independently (and simultaneously). Therefore, the core is described as having two MMUs, one for instruction accesses (IMMU) and one for data accesses (DMMU).

The block address translation (BAT) mechanism is a software-controlled array that stores the available block address translations on-chip. BAT array entries are implemented as pairs of BAT registers that are accessible as supervisor-level special-purpose registers

(SPRs). There are separate instruction and data BAT mechanisms, and in the G2 core, they reside in the instruction and data MMUs, respectively.

The MMUs, together with the exception processing mechanism, provide the necessary support for the operating system to implement a paged virtual memory environment and for enforcing protection of designated memory areas. Exception processing is described in Chapter 5, “Exceptions.” Section 5.2, “Exception Processing,” describes the MSR which controls some of the critical functionality of the MMUs.

6.1 MMU Features

The G2 core completely implements all features required by the memory management specification of the OEA for 32-bit implementations. Thus, it provides 4 Gbytes of effective address space accessible to supervisor and user programs with a 4-Kbyte page size and 256-Mbyte segment size. In addition, the MMUs of 32-bit processors use an interim virtual address (52 bits) and hashed page tables in the generation of 32-bit physical addresses. These processors also have a BAT mechanism for mapping large blocks of memory. Block sizes range from 128 Kbytes to 256 Mbytes and are software-programmable.

Table 6-1 summarizes all G2 core MMU features including the architectural features of PowerPC MMUs (defined by the OEA) for 32-bit processors and the implementation-specific features provided by the core.

Table 6-1. MMU Features Summary

Feature Category	Architecturally Defined/ G2 Core-Specific	Feature
Address ranges	Architecturally defined	2 ³² bytes of effective address
		2 ⁵² bytes of virtual address
		2 ³² bytes of physical address
Page size	Architecturally defined	4 Kbytes
Segment size	Architecturally defined	256 Mbytes
Block address translation	Architecturally defined	Range of 128 Kbytes–256 Mbytes sizes
		Implemented with IBAT and DBAT registers in BAT array
Memory protection	Architecturally defined	Segments selectable as no-execute
		Pages selectable as user/supervisor and read-only
		Blocks selectable as user/supervisor and read-only
Page history	Architecturally defined	Referenced and changed bits defined and maintained
Page address translation	Architecturally defined	Translations stored as PTEs in hashed page tables in memory
		Page table size determined by mask in SDR1 register

Table 6-1. MMU Features Summary (continued)

Feature Category	Architecturally Defined/ G2 Core-Specific	Feature
TLBs	Architecturally defined	Instructions for maintaining optional TLBs (tlbie instruction in G2 core)
	G2 core-specific	64-entry (32-entry byway), two-way set-associative ITLB 64-entry(32-entry byway), two-way set-associative DTLB
Segment descriptors	Architecturally defined	Stored as segment registers on-chip
Page table search support	G2 core-specific	Three MMU exceptions defined: ITLB miss exception, DTLB miss on load exception, and DTLB miss on store (or C = 0) exception; MMU-related bits set in SRR1 for these exceptions.
		IMISS and DMISS registers (missed effective address) HASH1 and HASH2 registers (PTEG addr) ICMP and DCMP registers (for comparing PTEs) RPA register (for loading TLBs)
		tlbli rB instruction for loading ITLB entries tlbid rB instruction for loading DTLB entries
		Shadow registers for GPR0–GPR3 (can use r0–r3 in table search handler without corruption of r0–r3 in context that was previously executing), called TGPR0–TGPR3.

6.1.1 Memory Addressing

A program references memory using the effective (logical) address computed by the processor when it executes a load, store, or cache instruction, and when it fetches the next instruction. The effective address is translated to a physical address according to the procedures described in Chapter 7, “Memory Management,” in the *Programming Environments Manual*, augmented with information in this chapter. The memory subsystem uses the physical address for the access.

For a complete discussion of effective address calculation, see Section 3.2.2.3, “Effective Address Calculation.”

6.1.2 MMU Organization

Figure 6-1 shows the conceptual organization of a PowerPC MMU in a 32-bit implementation; note that it does not describe the specific hardware used to implement the memory management function for a particular processor. Processors may optionally implement on-chip TLBs and may optionally support the automatic search of the page tables for PTEs. In addition, other hardware features (invisible to the system software) not depicted in the figure may be implemented.

Figure 6-2 and Figure 6-3 show the conceptual organization of the G2 core instruction and data MMUs, respectively. The instruction addresses shown in Figure 6-2 are generated by the processor for sequential instruction fetches and addresses that correspond to a change

of program flow. Data addresses shown in Figure 6-3 are generated by load and store instructions and by cache instructions.

As shown in the figures, after an address is generated, the higher-order bits of the effective address, EA0–EA19 (or a smaller set of address bits, EA0–EA n , in the cases of blocks), are translated into physical address bits PA0–PA19. The lower-order address bits, A20–A31, are untranslated and, therefore, identical for both effective and physical addresses. After translating the address, the MMUs pass the resulting 32-bit physical address to the memory subsystem.

In addition to the higher-order address bits, the MMUs automatically keep an indicator of whether each access was generated as an instruction or data access and a supervisor/user indicator that reflects the state of the PR bit of the MSR when the effective address was generated. In addition, for data accesses, there is an indicator of whether the access is for a load or a store operation. This information is then used by the MMUs to appropriately direct the address translation and to enforce the protection hierarchy programmed by the operating system. Section 5.2, “Exception Processing,” describes the MSR, which controls some of the critical functionality of the MMUs.

The figures show how the A20–A26 address bits index into the on-chip instruction and data caches to select a cache set. The remaining physical address bits are then compared with the tag fields (comprised of bits PA0–PA19) of the four selected cache blocks to determine if a cache hit has occurred. In the case of a cache miss, the instruction or data access is then forwarded to the bus interface unit which then initiates a 60x bus access to the memory subsystem.

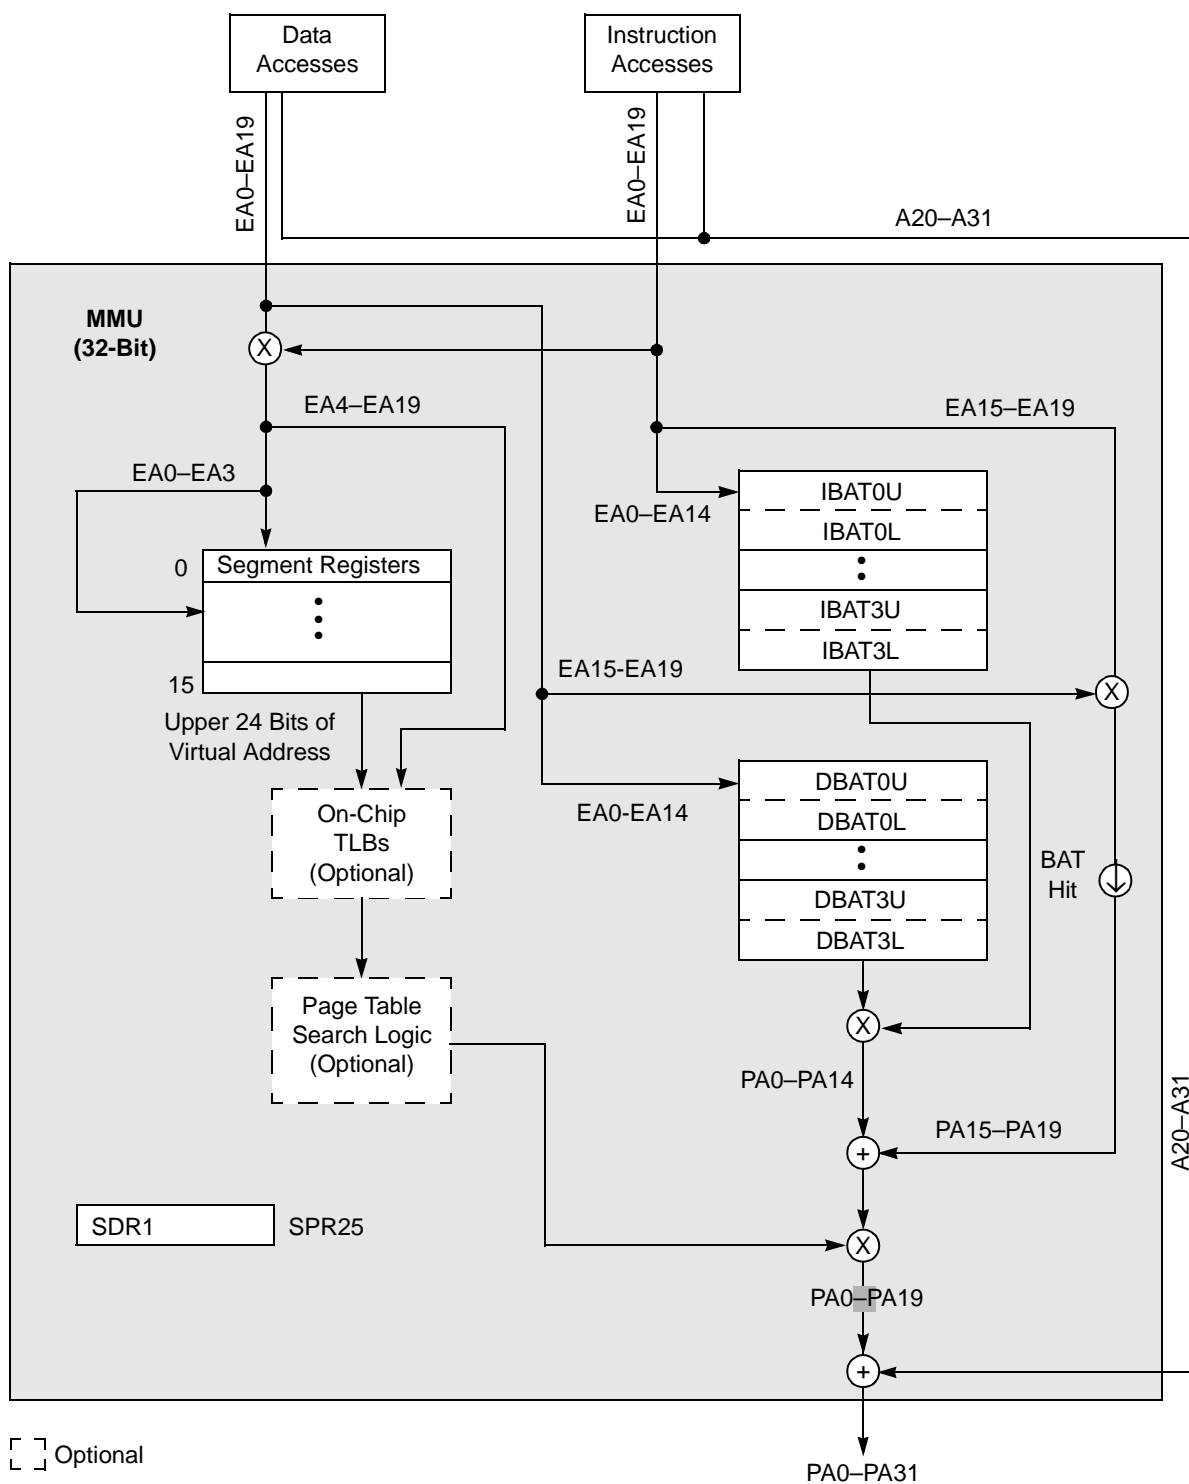


Figure 6-1. MMU Conceptual Block Diagram—32-Bit Implementations

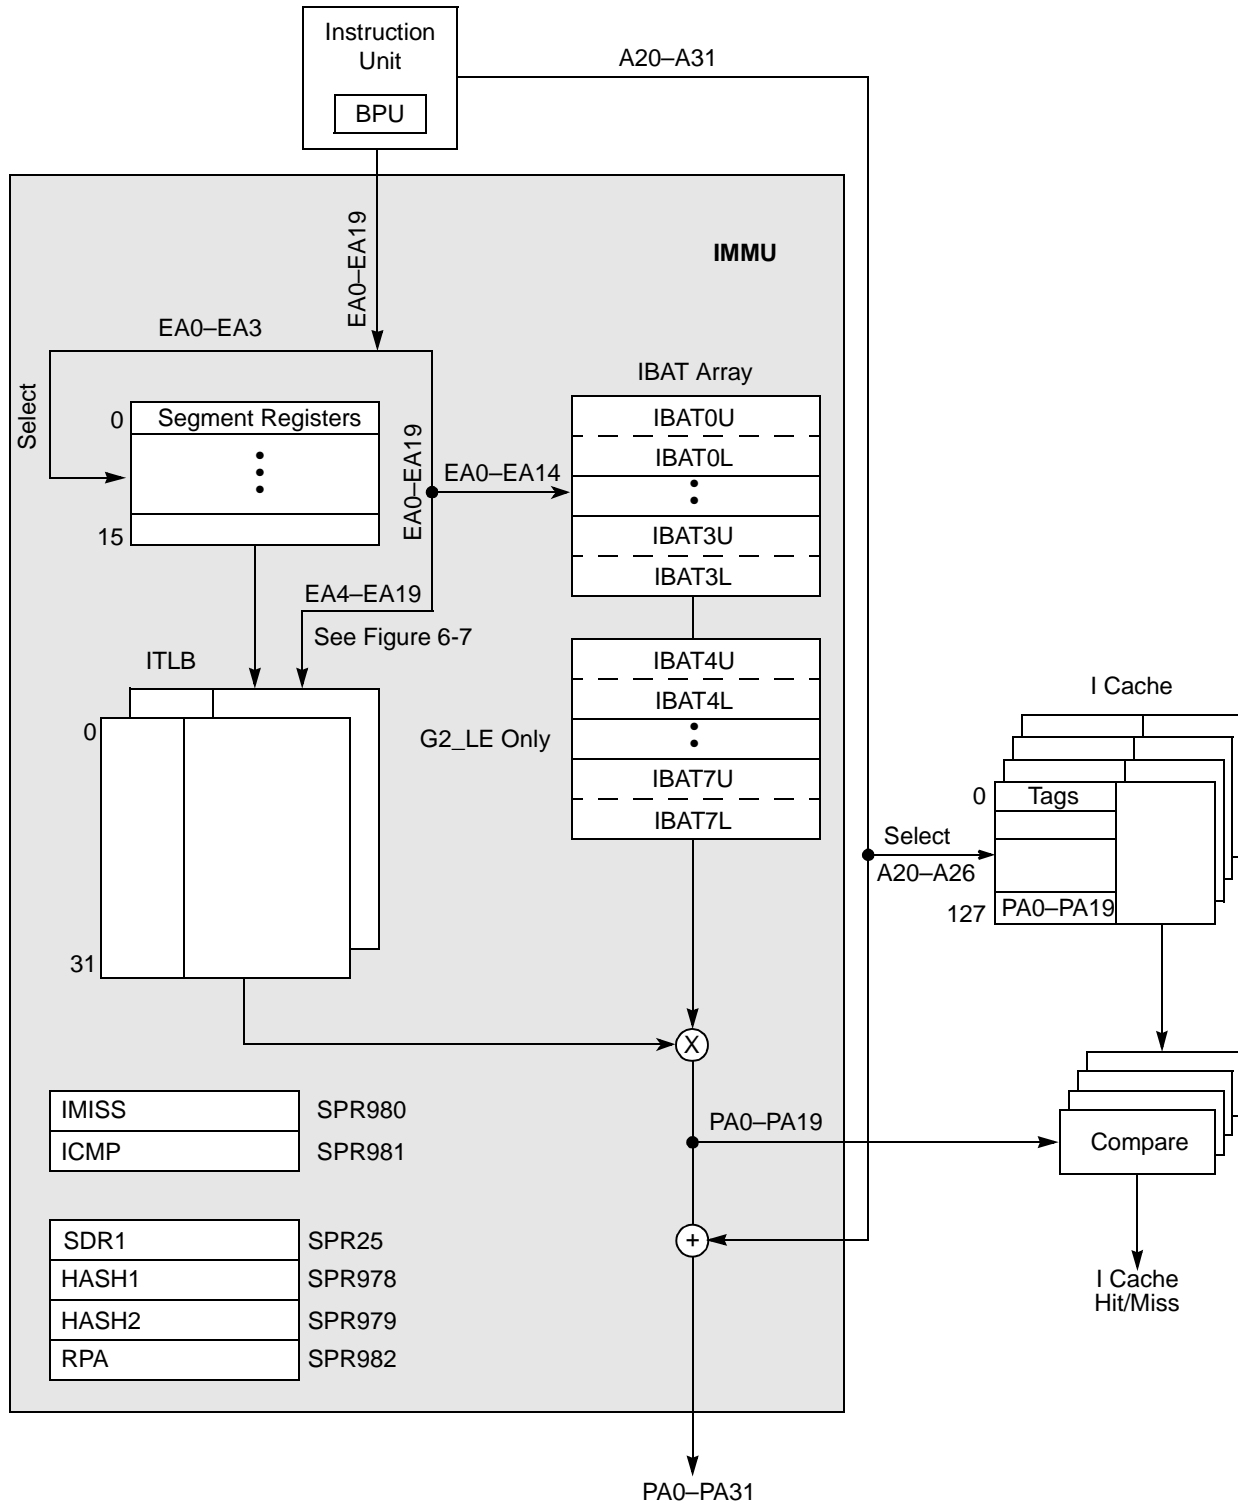


Figure 6-2. G2 Core IMMU Block Diagram

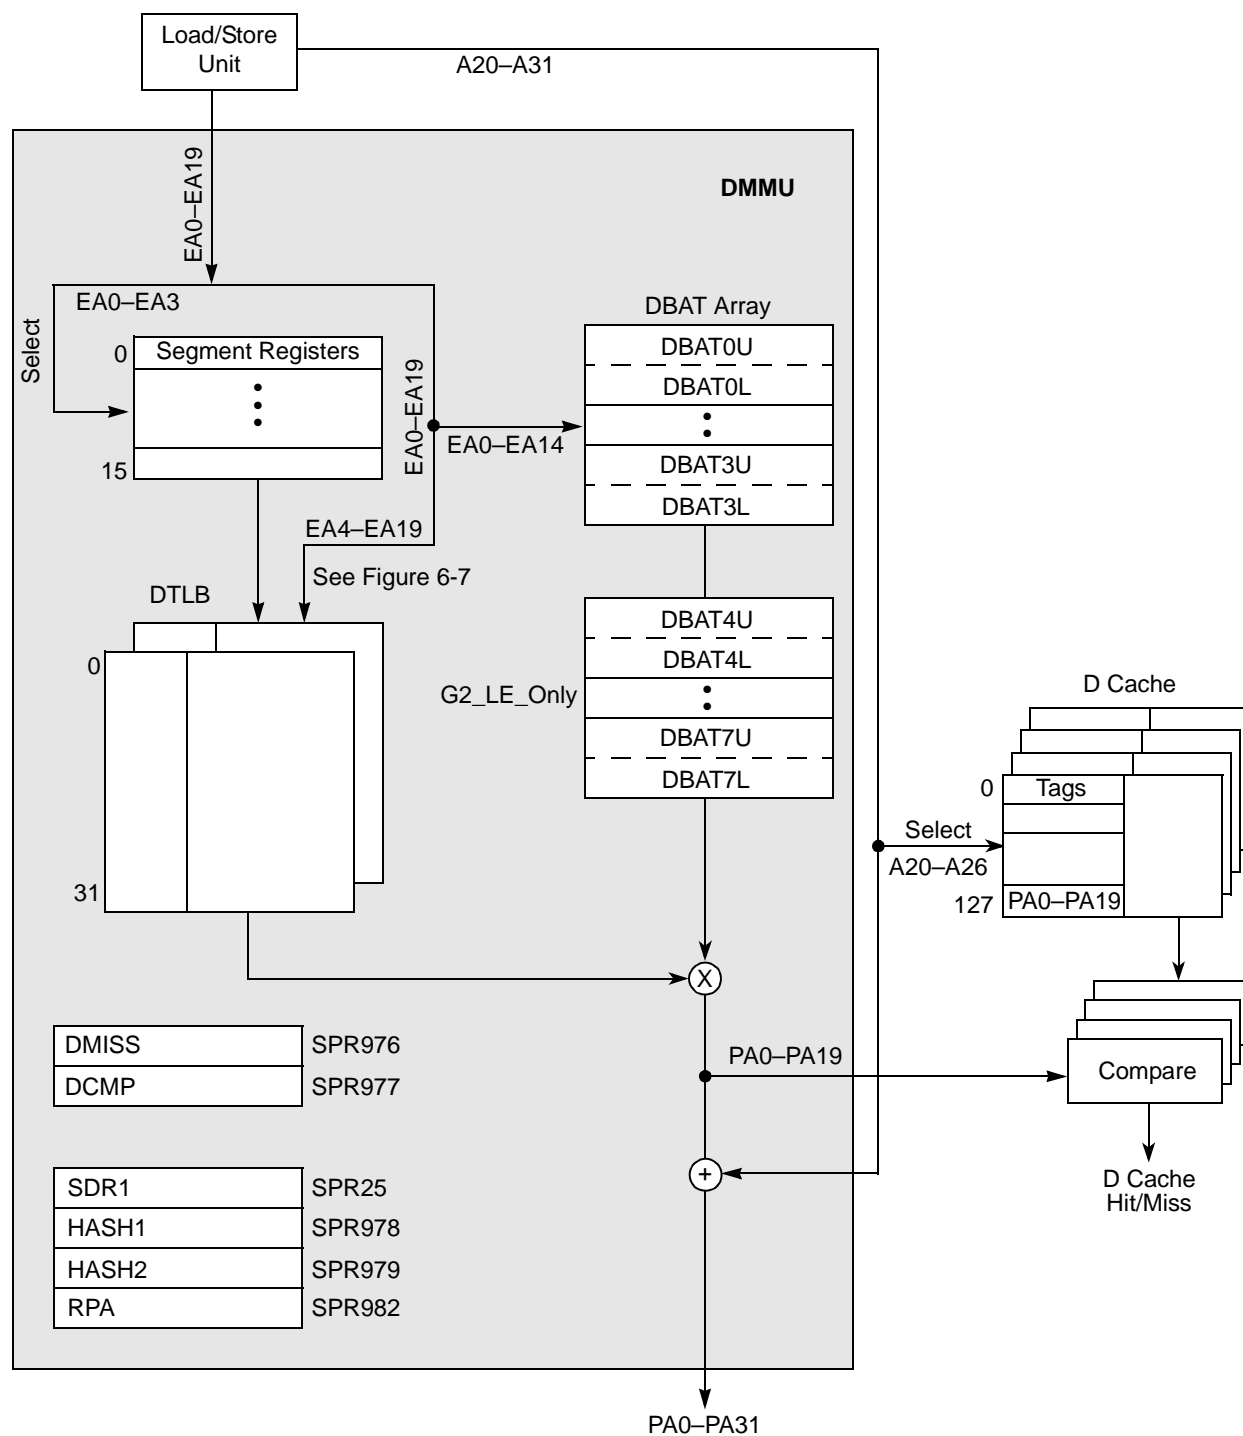


Figure 6-3. G2 Core DMMU Block Diagram

6.1.3 Address Translation Mechanisms

Processors that implement the PowerPC architecture support the following four types of address translation:

- Page address translation—translates the page frame address for a 4-Kbyte page size.
- Block address translation—translates the block number for blocks that range in size from 128 Kbytes to 256 Mbytes.
- Direct-store interface address translation—used to generate direct-store interface accesses on the external bus; not implemented in the G2 core.
- Real addressing mode translation—when address translation is disabled, the physical address is identical to the effective address.

Figure 6-4 shows the three implemented address translation mechanisms provided by the MMUs. The segment descriptors shown in the figure, control the page address translation mechanism. When an access uses page address translation, the appropriate segment descriptor is required. In 32-bit implementations, one of the 16 on-chip segment registers (which contain segment descriptors) is selected by the 4 highest-order effective address bits.

A control bit in the corresponding segment descriptor then determines if the access is to memory (memory-mapped) or to the direct-store interface space (selected when the direct-store translation control bit (T bit) in the corresponding segment descriptor is set). Note that the direct-store interface existed in previous processors only for compatibility with I/O devices that use this interface. When an access is determined to be to the direct-store interface space, the G2 core takes a DSI exception as described in Section 5.5.3, “DSI Exception (0x00300),” if it is a data access. The G2 core takes an ISI exception as described in Section 5.5.4, “ISI Exception (0x00400),” if it is an instruction access.

For memory accesses translated by a segment descriptor, the interim virtual address is generated using the information in the segment descriptor. Page address translation corresponds to the conversion of this virtual address into the 32-bit physical address used by the memory subsystem. In most cases, the physical address for the page resides in an on-chip TLB and is available for quick access. However, if the page address translation misses in an on-chip TLB, the MMU causes a search of the page tables in memory (using the virtual address information and a hashing function) to locate the required physical address. When this occurs, the G2 core vectors to the exception handlers that search the page tables with software.

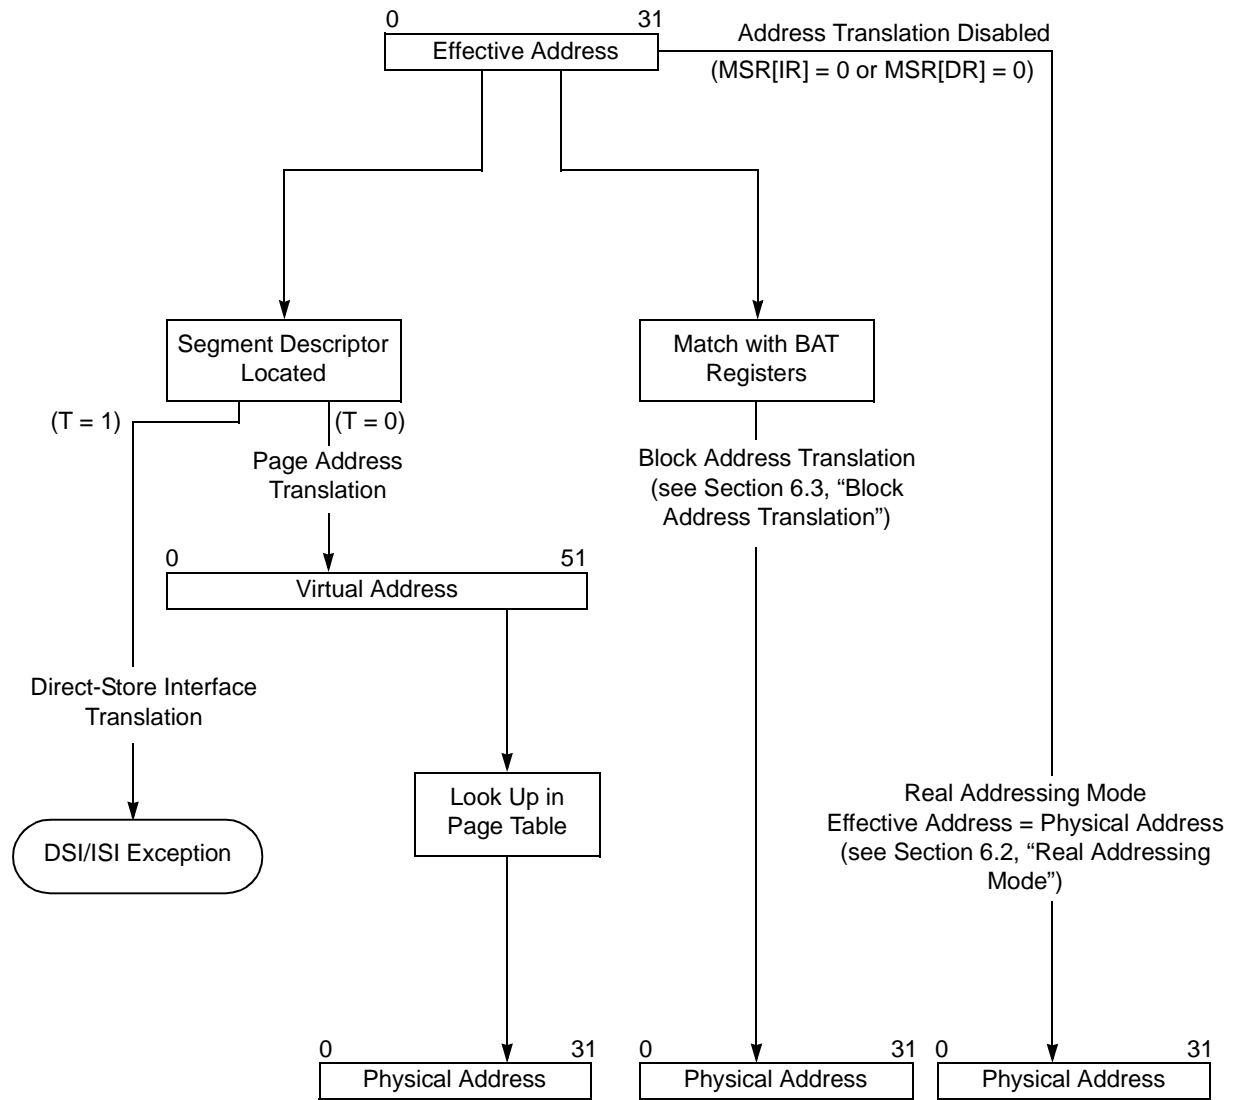


Figure 6-4. Address Translation Types

Block address translation occurs in parallel with page address translation and is similar to page address translation; however, fewer higher-order effective address bits are translated into physical address bits (more lower-order address bits (at least 17) are untranslated to form the offset into a block). Also, instead of segment descriptors and a TLB, block address translations use the on-chip BAT registers as a BAT array. If an effective address matches the corresponding field of a BAT register, the information in the BAT register is used to generate the physical address; in this case, the results of the page translation (occurring in parallel) are ignored (even if the segment corresponds to the direct-store interface space).

Real addressing mode translation occurs when address translation is disabled; in this case, the physical address generated is identical to the effective address. Instruction and data address translation is enabled with the $MSR[IR]$ and $MSR[DR]$ bits, respectively. Thus, when the processor generates an access, and the corresponding address translation enable

bit in MSR is cleared (MSR[IR] for instruction accesses and MSR[DR] for data accesses), the resulting physical address is identical to the effective address and all other translation mechanisms are ignored.

6.1.4 Memory Protection Facilities

In addition to the translation of effective addresses to physical addresses, the MMUs provide access protection of supervisor areas from user access and can designate areas of memory as read-only, as well as, no-execute or guarded. Table 6-2 shows the eight protection options supported by the MMUs for pages.

Table 6-2. Access Protection Options for Pages

Option	User Read		User Write	Supervisor Read		Supervisor Write
	I-Fetch	Data		I-Fetch	Data	
Supervisor-only	—	—	—	√	√	√
Supervisor-only-no-execute	—	—	—	—	√	√
Supervisor-write-only	√	√	—	√	√	√
Supervisor-write-only-no-execute	—	√	—	—	√	√
Both user/supervisor	√	√	√	√	√	√
Both user/supervisor-no-execute	—	√	√	—	√	√
Both read-only	√	√	—	√	√	—
Both read-only-no-execute	—	√	—	—	√	—

Note:

- √ access permitted.
- protection violation.

The operating system programs whether instructions can be fetched from an area of memory by appropriately using the no-execute option provided in the segment descriptor. Each of the remaining options is enforced, based on a combination of information in the segment descriptor and the page table entry. Thus, the supervisor-only option allows only read and write operations generated while the processor is operating in supervisor mode (corresponding to MSR[PR] = 0) to access the page. User accesses that map into a supervisor-only page cause an exception to be taken.

Finally, there is a facility in the VEA and OEA that allows pages or blocks to be designated as guarded, preventing out-of order accesses that may cause undesired side effects. For example, areas of the memory map that are used to control I/O devices can be marked as guarded so that accesses (for example, instruction prefetches) do not occur unless they are explicitly required by the program.

For more information on memory protection, see “Memory Protection Facilities” in Chapter 7, “Memory Management,” in the the *Programming Environments Manual*.

6.1.5 Page History Information

The MMUs of these processors also define referenced (R) and changed (C) bits in the page address translation mechanism that can be used as history information relevant to the page. This information can then be used by the operating system to determine the areas of memory to write back to disk when new pages must be allocated in main memory. While these bits are initially programmed by the operating system into the page table, the architecture specifies that the R and C bits may be maintained either by the processor hardware (automatically) or by some software-assist mechanism that updates these bits when required as needed by the G2 core. The software table search routines used by the G2 core set the R bit when a PTE is accessed; the core causes an exception (to vector to the software table search routines) when the C bit in the corresponding TLB entry requires updating. See Section 6.4.1.3, “Scenarios for Referenced and Changed Bit Recording,” for more details.

6.1.6 General Flow of MMU Address Translation

The following sections describe the general flow used by processors that implement the PowerPC architecture to translate effective addresses to virtual and then physical addresses.

6.1.6.1 Real Addressing Mode and Block Address Translation Selection

When an instruction or data access is generated and the corresponding instruction or data translation is disabled ($MSR[IR] = 0$ or $MSR[DR] = 0$), real addressing mode translation is used (physical address equals effective address) and the access continues to the memory subsystem as described in Section 6.2, “Real Addressing Mode.”

Figure 6-5 shows the flow used by the MMUs in determining whether to select real addressing mode, block address translation, or to use the segment descriptor to select page address translation.

Note that if the BAT array search results in a hit, the access is qualified with the appropriate protection bits. If the access violates the protection mechanism, an exception (ISI or DSI exception) is generated.

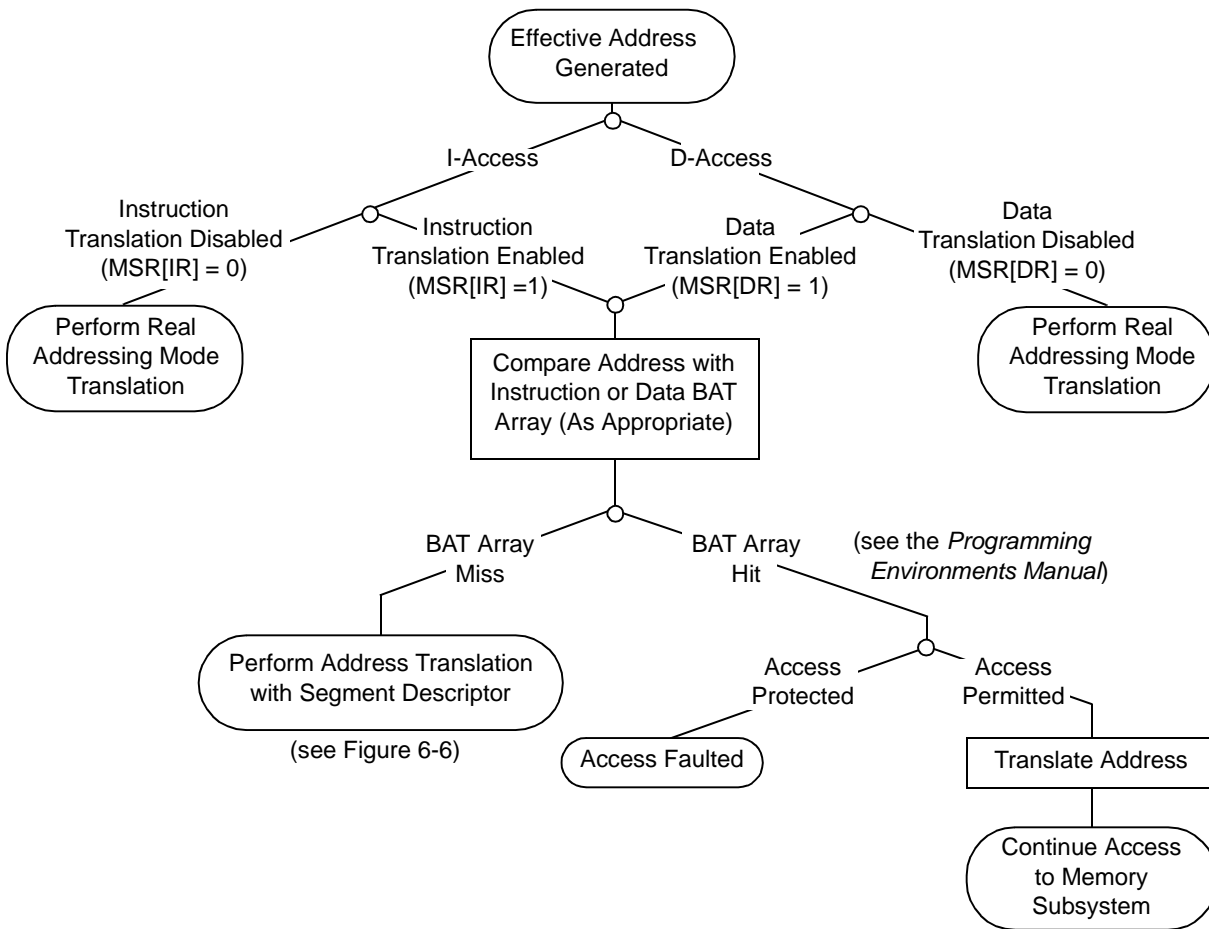
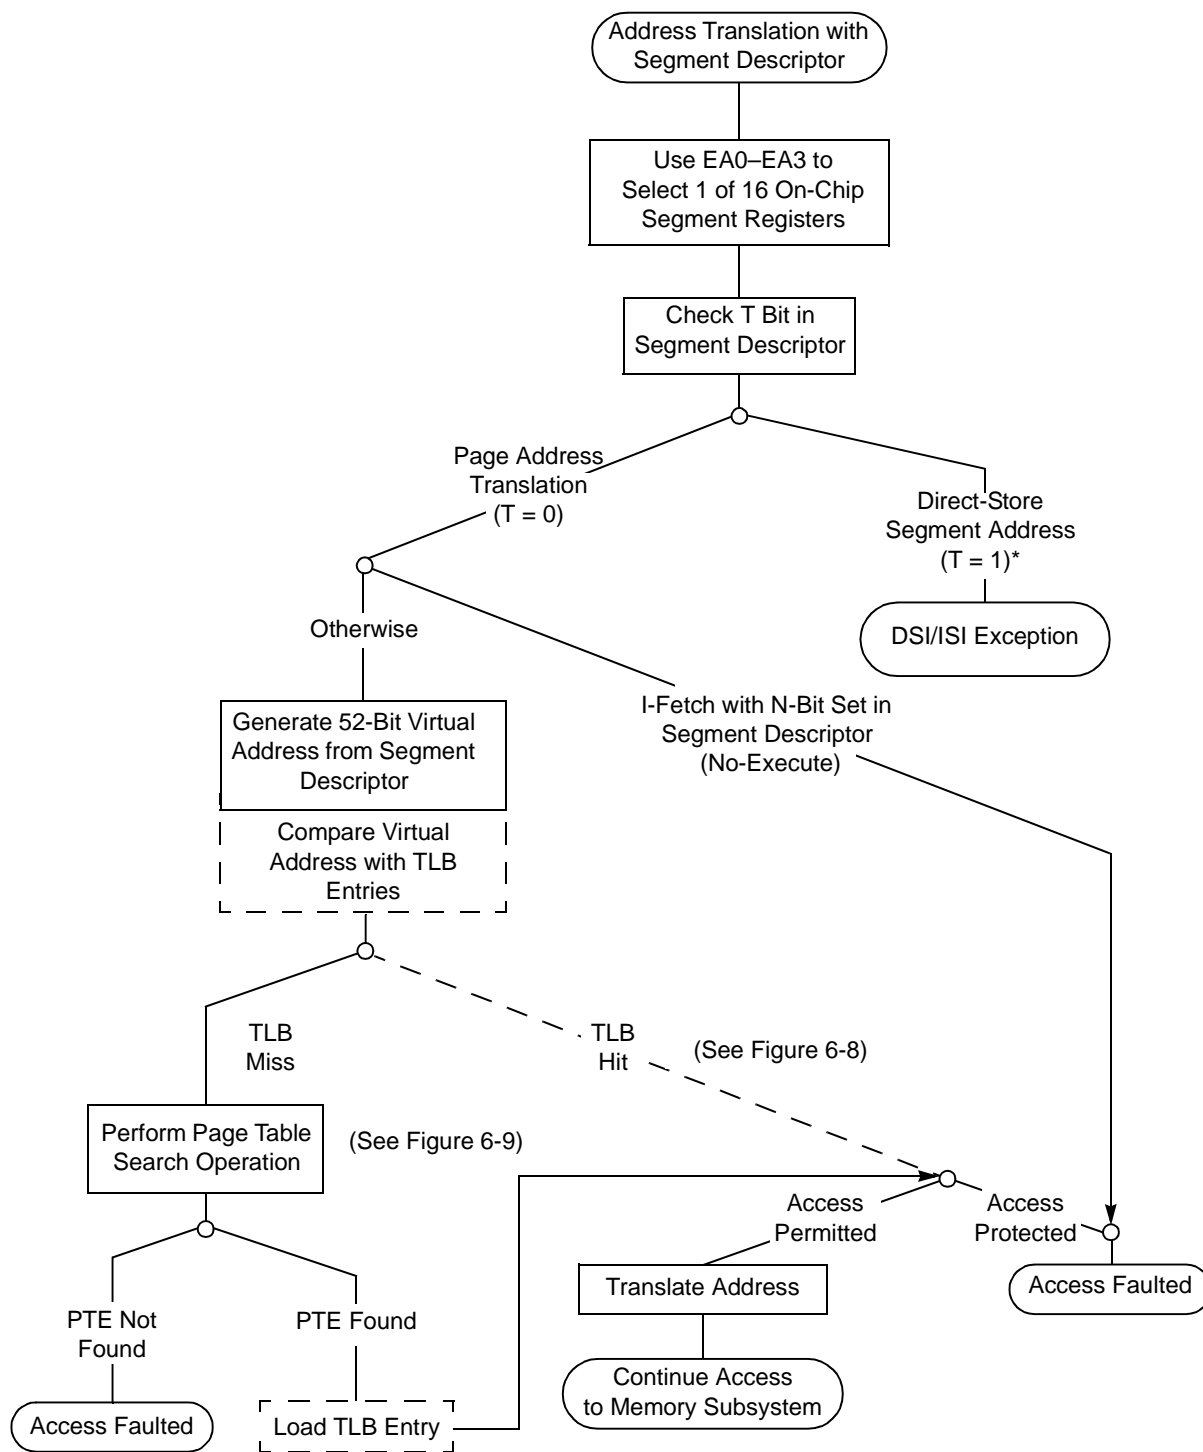


Figure 6-5. General Flow of Address Translation (Real Addressing Mode and Block)

6.1.6.2 Page Address Translation Selection

If address translation is enabled (real addressing mode not selected) and the effective address information does not match with a BAT array entry, then the segment descriptor must be located. Once the segment descriptor is located, the T bit in the segment descriptor selects whether the translation is to a page or to a direct-store interface segment, as shown in Figure 6-6. Note that the G2 core does not implement the direct-store interface, and accesses to these segments cause a DSI exception. In addition, Figure 6-6 also shows the way the no-execute protection is enforced; if the N bit in the segment descriptor is set and the access is an instruction fetch, the access is faulted as described in Chapter 7, “Memory Management,” in the *Programming Environments Manual*. Note that the figure shows the flow for these cases as described by the OEA and, therefore, the TLB references are shown as optional. Since the core implements TLBs, these branches are valid, and described in more detail throughout this chapter.



- - - Optional to the PowerPC architecture. Implemented in the MPC603e.

*In the case of instruction accesses, causes ISI exception.

Figure 6-6. General Flow of Page and Direct-Store Interface Address Translation

If the T bit in the corresponding segment descriptor is zero, page address translation is selected. The information in the segment descriptor is then used to generate the 52-bit virtual address. The virtual address is then used to identify the page address translation information (stored as page table entries (PTEs) in a page table in memory). For increased performance, the core has two TLBs to store recently-used PTEs on-chip.

If an access hits in the appropriate TLB, the page translation occurs and the physical address bits are forwarded to the memory subsystem. If the required PTE is not resident, the MMU requires a search of the page table. In this case, the core traps to one of three exception handlers for the system software to perform the page table search. If the PTE is successfully matched, a new TLB entry is created and the page translation is once again attempted. This time, the TLB is guaranteed to hit. Once the PTE is located, the access is qualified with the appropriate protection bits. If the access is a protection violation (not allowed), an exception (instruction access or data access) is generated.

If the PTE is not found by the table search operation, a page fault condition exists, and the TLB miss exception handlers synthesize either an ISI or DSI exception to handle the page fault.

6.1.7 MMU Exceptions Summary

In order to complete any memory access, the effective address must be translated to a physical address. In the G2 core, an MMU exception condition occurs if this translation fails for one of the following reasons:

- Page fault—There is no valid page table entry to identify the page specified by the effective address (and segment descriptor) and there is no valid BAT translation.
- An address translation is found but the access is not allowed by the memory protection mechanism.

Additionally, because the core relies on software to perform table search operations, the processor also takes an exception when:

- There is a miss in the corresponding (instruction or data) TLB.
- The page table requires an update to the changed (C) bit.

The state saved by the processor for each of these exceptions contains information that identifies the address of the failing instruction. Refer to Chapter 5, “Exceptions,” for a more detailed description of exception processing.

Because a page fault condition (PTE not found in the page tables in memory) is detected by the software that performs the table search operation (and not the core hardware), it does not cause a G2 core exception, in the strictest sense, in that exception processing as described in Chapter 5, “Exceptions,” does not occur. However, in order to maintain architectural compatibility with software written for other devices that implement the PowerPC architecture, the software that detects this condition should synthesize an

exception by setting the appropriate bits in the DSISR or SRR1 and branching to the ISI or DSI exception handler. Refer to Section 6.5.2, “Implementation-Specific Table Search Operation,” for more information and examples of this exception software. The remainder of this chapter assumes that the table search software emulates this exception and refers to this condition as an exception.

The translation exception conditions defined by the OEA for 32-bit implementations cause either the ISI or the DSI exception to be taken as shown in Table 6-3.

Table 6-3. Translation Exception Conditions

Condition	Description	Exception
Page fault (no PTE found)	No matching PTE found in page tables (and no matching BAT array entry)	I access: ISI exception ¹ SRR1[1] = 1
		D access: DSI exception ¹ DSISR[1] = 1
Block protection violation	Conditions described for block in “Block Memory Protection” in Chapter 7, “Memory Management,” in the <i>Programming Environments Manual</i> .”	I access: ISI exception SRR1[4] = 1
		D access: DSI exception DSISR[4] = 1
Page protection violation	Conditions described for page in “Page Memory Protection” in Chapter 7, “Memory Management,” in the <i>Programming Environments Manual</i> .”	I access: ISI exception ² SRR1[4] = 1
		D access: DSI exception ² DSISR[4] = 1
No-execute protection violation	Attempt to fetch instruction when SR[N] = 1	ISI exception SRR1[3] = 1
Instruction fetch from direct-store segment	Attempt to fetch instruction when SR[T] = 1	ISI exception SRR1[3] = 1
Data access to direct-store segment (including floating-point accesses) Note: This is a G2 core-specific condition	Attempt to perform load or store (including floating-point load or store) when SR[T] = 1	DSI exception DSISR[5] = 1
Instruction fetch from guarded memory with MSR[IR] = 1	Attempt to fetch instruction when MSR[IR] = 1 and either matching xBAT[G] = 1, or no matching BAT entry and PTE[G] = 1.	ISI exception SRR1[3] = 1

¹ The G2 core hardware does not vector to these exceptions automatically. It is assumed that the software that performs the table search operation vectors to these exceptions and sets the appropriate bits when a page fault condition occurs.

² The table search software can also vector to these exception conditions.

In addition to the translation exceptions, there are other MMU-related conditions (some of them defined as implementation-specific and, therefore, not required by the architecture) that can cause an exception to occur in the G2 core. These exception conditions map to the processor exception as shown in Table 6-4. For example, the G2 core also defines three exception conditions to support software table searching. The only exception conditions that occur when MSR[DR] = 0, are the conditions that cause the alignment exception for

data accesses. For more detailed information about the conditions that cause the alignment exception (in particular for string/multiple instructions), see Section 5.5.6, “Alignment Exception (0x00600).”

Table 6-4. Other MMU Exception Conditions

Condition	Description	Exception
TLB miss for an instruction fetch	No matching entry found in ITLB	Instruction TLB miss exception SRR1[13] = 1 MSR[14] = 1
TLB miss for a data load access	No matching entry found in DTLB for data load access	Data TLB miss on load exception SRR1[13] = 0 SRR1[15] = 1 MSR[14] = 1
TLB miss for a data store, or store and C = 0	No matching entry found in DTLB for data store access or matching DLTB entry has C = 0 and the access is a store	Data TLB miss on store exception, or store and C = 0 SRR1[13] = 0 SRR1[15] = 0 MSR[14] = 1
dcbz with W = 1 or I = 1	dcbz instruction to write-through or cache-inhibited segment or block	Alignment exception (not required by architecture for this condition)
dcbz when the data cache is locked	The dcbz instruction takes an alignment exception if the data cache is locked (HID0 bits 18 and 19) when it is executed	Alignment exception
lwarx , stwcx. , eciwx , or ecowx instruction to direct-store segment	Reservation instruction or external control instruction when SR[T] = 1	DSI exception DSISR[5] = 1
Floating-point load or store to direct-store segment	FP memory access when SR[T] = 1	See data access to direct-store segment in Table 6-3
Load or store that results in a direct-store error	Does not occur in G2 core	Does not apply
eciwx or ecowx attempted when external control facility disabled	eciwx or ecowx attempted with EAR[E] = 0	DSI exception DSISR[11] = 1
lmw , stmw , lswi , lswx , stswi , or stswx instruction attempted in little-endian mode	lmw , stmw , lswi , lswx , stswi , or stswx instruction attempted while MSR[LE] = 1.	Alignment exception
Operand misalignment	Translation enabled and operand is misaligned as described in Chapter 5, “Exceptions.”	Alignment exception (some of these cases are implementation-specific)

Note that some exception conditions depend on whether the memory area is set up as write-through (W = 1) or cache-inhibited (I = 1). These bits are described fully in “Memory/Cache Access Attributes” in Chapter 5, “Cache Model and Memory Coherency,” in the *Programming Environments Manual*. Refer to Chapter 5, “Exceptions,” and to Chapter 6, “Exceptions,” in the *Programming Environments Manual* for a complete description of the SRR1 and DSISR bit settings for these exceptions.

6.1.8 MMU Instructions and Register Summary

The MMU instructions and registers provide the operating system with the ability to set up the block address translation areas and the page tables in memory.

Note that because the implementation of TLBs is optional, the instructions that refer to these structures are also optional. However, because these structures serve as caches of the page table, the architecture specifies a software protocol for maintaining coherency between these caches and the tables in memory whenever changes are made to the tables in memory. When the tables in memory are changed, the operating system purges these caches of the corresponding entries, allowing the translation caching mechanism to refetch from the tables when the corresponding entries are required.

Note that the G2 core implements all TLB-related instructions except **tlbia**, which is treated as an illegal instruction. The G2 core also uses some implementation-specific instructions to load two on-chip TLBs.

Because the MMU specification for these processors is so flexible, it is recommended that the software that uses these instructions and registers be encapsulated into subroutines to minimize the impact of migrating across the family of implementations.

Table 6-5 summarizes G2 core instructions that specifically control the MMU. For more detailed information about the instructions, refer to Chapter 3, “Instruction Set Model,” in this book and Chapter 8, “Instruction Set,” in the *Programming Environments Manual*.

Table 6-5. Instruction Summary—MMU Control

Instruction	Description
mtsr SR,rS	Move to Segment Register SR[SR#]← rS
mtsrin rS,rB	Move to Segment Register Indirect SR[rB[0–3]]← rS
mfsr rD,SR	Move from Segment Register rD←SR[SR#]
mfsrin rD,rB	Move from Segment Register Indirect rD←SR[rB[0–3]]
tlbie rB ¹	TLB Invalidate Entry For effective address specified by rB, TLB[V]← 0 The tlbie instruction invalidates both TLB entries indexed by the EA, and operates on both the instruction and data TLBs simultaneously invalidating four TLB entries. The index corresponds to bits 15–19 of the EA. Software must ensure that instruction fetches or memory references to the virtual pages specified by the tlbie instruction have been completed prior to executing the tlbie instruction.
tlbsync ¹	TLB Synchronize Synchronizes the execution of all other tlbie instructions in the system. In the G2 core, when the <code>core_tlbsync</code> signal is negated, instruction execution may continue or resume after the completion of a tlbsync instruction. When the <code>core_tlbsync</code> signal is asserted, instruction execution stops after the completion of a tlbsync instruction. For a complete description of the <code>core_tlbsync</code> signal, refer to Section 8.3.11.5, “TLBI Sync (<code>core_tlbsync</code>)—Input.”

Table 6-5. Instruction Summary—MMU Control (continued)

Instruction	Description
tlbli (implementation-specific)	Load Instruction TLB Entry Loads the contents of the ICMP and RPA registers into the ITLB.
tlbld (implementation-specific)	Load Data TLB Entry Loads the contents of the DCMP and RPA registers into the DTLB.

¹ These instructions are defined by the PowerPC architecture, but are optional.

Table 6-6 summarizes the registers that the operating system uses to program the G2 core MMUs. These registers are accessible to supervisor-level software only. These registers are described in Chapter 2, “Register Set,” in the *Programming Environments Manual*. For G2 core-specific registers, see Chapter 2, “Register Model,” of this book.

Table 6-6. MMU Registers

Register	Description
Segment registers (SR0–SR15)	The sixteen 32-bit segment registers are present only in 32-bit implementations of the PowerPC architecture. The fields in the segment register are interpreted differently depending on the value of bit 0. The segment registers are accessed by the mtsr , mtsrin , mfsr , and mfsrin instructions.
BAT registers G2 core: (IBAT0U–IBAT3U, IBAT0L–IBAT3L, DBAT0U–DBAT3U, and DBAT0L–DBAT3L) G2_LE core: (IBAT0U–IBAT7U, IBAT0L–IBAT7L, DBAT0U–DBAT7U, and DBAT0L–DBAT7L)	The G2 core has 16 BAT registers, organized as 4 pairs of instruction BAT registers (IBAT0U–IBAT3U paired with IBAT0L–IBAT3L) and 4 pairs of data BAT registers (DBAT0U–DBAT3U paired with DBAT0L–DBAT3L). The G2_LE core has 32 BAT registers, organized as 8 pairs of instruction BAT registers (IBAT0U–IBAT7U paired with IBAT0L–IBAT7L) and 8 pairs of data BAT registers (DBAT0U–DBAT7U paired with DBAT0L–DBAT7L). The BAT registers are defined as 32-bit registers in 32-bit implementations. These are special-purpose registers that are accessed by the mtspr and mfspr instructions, regardless of the setting of HID2[13].
SDR1	The SDR1 register specifies the variable used in accessing the page tables in memory. SDR1 is defined as a 32-bit register for 32-bit implementations. This is a special-purpose register that is accessed by the mtspr and mfspr instructions.
Instruction TLB miss address and data TLB miss address registers (IMISS and DMISS)	When a TLB miss exception occurs, the IMISS or DMISS register contains the 32-bit effective address of the instruction or data access, respectively, that caused the miss. Note that the G2 core always loads a big-endian address into the DMISS register. These registers are implementation-specific.
Primary and secondary hash address registers (HASH1 and HASH2)	The HASH1 and HASH2 registers contain the primary and secondary PTEG addresses that correspond to the address causing a TLB miss. These PTEG addresses are automatically derived by the core by performing the primary and secondary hashing function on the contents of IMISS or DMISS, for an ITLB or DTLB miss exception, respectively. These registers are implementation-specific.

Table 6-6. MMU Registers (continued)

Register	Description
Instruction and data PTE compare registers (ICMP and DCMP)	The ICMP and DCMP registers contain the word to be compared with the first word of a PTE in the table search software routine to determine if a PTE contains the address translation for the instruction or data access. The contents of ICMP and DCMP are automatically derived by the core when a TLB miss exception occurs. These registers are implementation-specific.
Required physical address register (RPA)	The system software loads a TLB entry by loading the second word of the matching PTE entry into the RPA register and then executing the tlbli or tlbld instruction (for loading the ITLB or DTLB, respectively). This register is implementation-specific.

Note that the G2 core contains other features that do not specifically control the MMU, but are implemented to increase performance and flexibility. These are:

- Complete set of shadow segment registers for the instruction MMU. These registers are invisible to the programming model, as described in Section 6.4.3, “TLB Description.”
- Temporary GPR0–GPR3. These registers are available as **r0–r3** when MSR[TGPR] is set. The core automatically sets MSR[TGPR] whenever one of the three TLB miss exceptions occurs, allowing these exception handlers to have four registers that are used as scratchpad space, without having to save or restore this part of the machine state that existed when the exception occurred. Note that MSR[TGPR] is restored to the value in SRR1 when the **rfi** instruction is executed. Refer to Section 6.5.2, “Implementation-Specific Table Search Operation,” for code examples that take advantage of these registers.

In addition, the G2 core also automatically saves the values of CR[CR0] of the executing context to SRR1[0–3] whenever one of the three TLB miss exceptions occurs. Thus, the exception handler can set CR[CR0] bits and branch accordingly in the exception handler routine, without having to save the existing CR[CR0] bits. However, the exception handler must restore these bits to CR[CR0] before executing the **rfi** instruction. There are also four other bits saved in SRR1 whenever a TLB miss exception occurs that give information about whether the access was an instruction or data access; and if it was a data access, whether it was for a load or a store instruction. Also, these bits give some information related to the protection attributes for the access, and which set in the TLB will be replaced when the next TLB entry is loaded. Refer to Section 6.5.2.1, “Resources for Table Search Operations,” for more information on these bits and their use.

6.2 Real Addressing Mode

If address translation is disabled (MSR[IR] = 0 or MSR[DR] = 0) for a particular access, the effective address is treated as the physical address and is passed directly to the memory

subsystem as described in Chapter 7, “Memory Management,” in the *Programming Environments Manual*.

Note that the default WIMG bits (0b0011) cause data accesses to be considered cacheable ($I = 0$) and, thus, load and store accesses are weakly ordered. This is the case, even if the data cache is disabled in the HID0 register (as it is out of hard reset). If I/O devices require load and store accesses to occur in strict program order (strongly ordered), translation must be enabled so that the corresponding I bit can be set. Also, for instruction accesses, the default memory access mode bits (WIMG) are 0b0001. That is, instruction accesses are considered cacheable ($I = 0$), and the memory is guarded. Again, instruction cache accesses are considered cacheable even if the instruction cache is disabled in the HID0 register (as it is out of hard reset). The W and M bits have no effect on the instruction cache.

For information on the synchronization requirements for changes to MSR[IR] and MSR[DR], refer to “Synchronization Requirements for Special Registers and for Lookaside Buffers” in Chapter 2, “Register Set,” in the *Programming Environments Manual*.

6.3 Block Address Translation

The block address translation (BAT) mechanism in the OEA provides a way to map ranges of effective addresses larger than a single page into contiguous areas of physical memory. Such areas can be used for data that is not subject to normal virtual memory handling (paging), such as a memory-mapped display buffer or an extremely large array of numerical data.

The software model for block address translation in the G2 core is described in Chapter 7, “Memory Management,” in the *Programming Environments Manual* for 32-bit implementations. However, note that for improved performance, the G2_LE core contains twice as many BAT registers as the G2 core, as shown in Figure 6-2 and Figure 6-3.

Implementation Note—The BAT registers are not initialized by the hardware after the power-up or reset sequence. Consequently, all valid bits in both instruction and data BAT areas must be explicitly cleared before setting any BAT area for the first time and before enabling translation. Also, note that software must avoid overlapping blocks while updating a BAT area or areas. Even if translation is disabled, multiple BAT area hits (with the valid bits set) can corrupt the remaining portion (any bits except the valid bits) of the BAT registers.

Thus, multiple BAT hits (with valid bits set) are considered a programming error whether translation is enabled or disabled, and can lead to unpredictable results if translation is enabled, (or if translation is disabled, when translation is eventually enabled). For the case of unused BATs (if translation is to be enabled), it is sufficient precaution to simply clear the valid bits of the unused BAT entries.

6.4 Memory Segment Model

The G2 core adheres to the memory segment model as defined in Chapter 7, “Memory Management,” in the *Programming Environments Manual* for 32-bit implementations. Memory in the OEA is divided into 256-Mbyte segments. This segmented memory model provides a way to map 4-Kbyte pages of effective addresses to 4-Kbyte pages in physical memory (page address translation), while providing the programming flexibility afforded by a large virtual address space (52 bits).

The segment/page address translation mechanism may be superseded by the BAT mechanism described in Section 6.3, “Block Address Translation.” If not, the translation proceeds in the following two steps:

1. From effective address to the virtual address (which never exists as a specific entity, but can be considered to be the concatenation of the virtual page number and the byte offset within a page).
2. From virtual address to physical address.

The following section highlights those areas of the memory segment model defined by the OEA that are specific to the G2 core.

6.4.1 Page History Recording

Referenced (R) and changed (C) bits reside in each PTE to keep history information about the page. They are maintained by a combination of the core hardware and the table search software. The operating system uses this information to determine which areas of memory to write back to disk when new pages must be allocated in main memory. Referenced and changed recording is performed only for accesses made with page address translation and not for translations made with the BAT mechanism or for accesses that correspond to direct-store interface (T = 1) segments. Furthermore, R and C bits are maintained only for accesses made while address translation is enabled (MSR[IR] = 1 or MSR[DR] = 1).

In the G2 core, the referenced and changed bits are updated as follows:

- For TLB hits, the C bit is updated according to Table 6-7.
- For TLB misses, when a table search operation is in progress to locate a PTE, the R and C bits are updated (set, if required) to reflect the status of the page based on this access.

Table 6-7 shows that the status of the C bit in the TLB entry (in the case of a TLB hit) is what causes the processor to update the C bit in the PTE (the R bit is assumed to be set in the page tables if there is a TLB hit). Therefore, when software clears the R and C bits in the page tables in memory, it must invalidate the TLB entries associated with the pages whose referenced and changed bits were cleared.

Table 6-7. Table Search Operations to Update History Bits—TLB Hit Case

R and C Bits in TLB Entry	Processor Action	
00	Combination does not occur	
01	Combination does not occur	
10	Read: Write:	No special action Table search operation required to update C. Causes a data TLB miss on store exception.
11	No special action for read or write	

The G2 core causes the R bit to be set for the execution of the **dcbt** or **dcbtst** instruction to that page (by causing a TLB miss exception to load the TLB entry in the case of a TLB miss). However, neither of these instructions causes the C bit to be set.

As defined by the PowerPC architecture, the referenced and changed bits are updated as if address translation were disabled (real addressing mode translation). Additionally, these updates should be performed with single-beat read and byte write transactions on the bus.

6.4.1.1 Referenced Bit

The referenced (R) bit of a page is located in the PTE in the page table. Every time a page is referenced (with a read or write access) and the R bit is zero, the R bit is then set in the page table. The OEA specifies that the referenced bit may be set immediately, or the setting may be delayed until the memory access is determined to be successful. Because the reference to a page is what causes a PTE to be loaded into the TLB, the referenced bit in all G2 core TLB entries is effectively always set. The processor never automatically clears the referenced bit.

The referenced bit is only a hint to the operating system about the activity of a page. At times, the referenced bit may be set by software although the access was not logically required by the program, or even if the access was prevented by memory protection. Examples of this in these systems include the following:

- Fetching of instructions not subsequently executed
- Accesses generated by an **lswx** or **stswx** instruction with a zero length
- Accesses generated by a **stwcx.** instruction when no store is performed because a reservation does not exist
- Accesses that cause exceptions and are not completed

6.4.1.2 Changed Bit

The changed bit of a page is located both in the PTE in the page table and in the copy of the PTE loaded into the TLB (if a TLB is implemented, as in the G2 core). Whenever a data store instruction is executed successfully, if the TLB search (for page address translation)

results in a hit, the changed bit in the matching TLB entry is checked. If it is already set, the processor does not change the C bit. If the TLB changed bit is 0, it is set and a table search operation is performed to also set the C bit in the corresponding PTE in the page table. The G2 core causes a data TLB miss on store exception for this case so that the software can perform the table search operation for setting the C bit. Refer to Section 6.5.2, “Implementation-Specific Table Search Operation,” for an example code sequence that handles these conditions.

The changed bit (in both the TLB and PTE in the page tables) is set only when a store operation is allowed by the page memory protection mechanism and all conditional branches occurring earlier in the program have been resolved (such that the store is guaranteed to be in the execution path). Furthermore, the following conditions may cause the C bit to be set:

- The execution of an **stwcx.** instruction is allowed by the memory protection mechanism, but a store operation is not performed because no reservation exists.
- The execution of an **stswx** instruction is allowed by the memory protection mechanism, but a store operation is not performed because the specified length is zero.
- The store operation is not performed because an exception occurs before the store is performed.

Again, note that although the execution of the **dcbt** and **dcbtst** instructions may cause the R bit to be set, they never cause the C bit to be set.

6.4.1.3 Scenarios for Referenced and Changed Bit Recording

This section provides a summary of the model (defined by the OEA) that is used by the processors for maintaining the referenced and changed bits. In some scenarios, the bits are guaranteed to be set by the processor, in some scenarios, the architecture allows that the bits may be set (not absolutely required), and in some scenarios, the bits are guaranteed to not be set.

In implementations that do not maintain the R and C bits in hardware (such as the G2 core), software assistance is required. For these processors, the information in this section still applies, except that the software performing the updates is constrained to the rules described (that is, must set bits shown as guaranteed to be set and must not set bits shown as guaranteed to not be set).

Table 6-8 defines a prioritized list of the R and C bit settings for all scenarios. The entries in the table are prioritized from top to bottom, such that a matching scenario occurring closer to the top of the table takes precedence over a matching scenario closer to the bottom of the table. For example, if an **stwcx.** instruction causes a protection violation and there is no reservation, the C bit is not altered, as shown for the protection violation case. Note that in the table, load operations include those generated by load instructions, by the **eciwx**

instruction, and by the cache management instructions that are treated as a load with respect to address translation. Similarly, store operations include those operations generated by store instructions, by the **ecowx** instruction, and by the cache management instructions that are treated as a store with respect to address translation. In the columns for the G2 core, the combination of the core itself and the software used to search the page tables (described in Section 6.5.2, “Implementation-Specific Table Search Operation”) is assumed.

Table 6-8. Model for Guaranteed R and C Bit Settings

Priority	Scenario	R Bit Set		C Bit Set	
		OEA	G2 Core	OEA	G2 Core
1	No-execute protection violation	No	No	No	No
2	Page protection violation	Maybe	Yes	No	No
3	Out-of-order instruction fetch or load operation	Maybe	No	No	No
4	Out-of-order store operation for instructions that will cause no other kind of precise exception (in the absence of system-caused, imprecise, or floating-point assist exceptions)	Maybe ¹	No	No	No
5	All other out-of-order store operations	Maybe ¹	No	Maybe ¹	No
6	Zero-length load (lswx)	Maybe	Yes	No	No
7	Zero-length store (stswx)	Maybe ¹	Yes	Maybe ¹	Yes
8	Store conditional (stwcx.) that does not store	Maybe ¹	Yes	Maybe ¹	Yes
9	In-order instruction fetch	Yes ²	Yes	No	No
10	Load instruction or eciwx	Yes	Yes	No	No
11	Store instruction, ecowx or dcbz instruction	Yes	Yes	Yes	Yes
12	dcbt , dcbtst , dcbst , or dcbf instruction	Maybe	Yes	No	No
13	icbi instruction	Maybe ¹	No	No ¹	No
14	dcbi ³ instruction	Maybe ¹	Yes	Maybe ¹	Yes

¹ If C is set, R is guaranteed to also be set.

² This includes the case when the instruction was fetched out-of-order and R was not set (does not apply for the G2 core).

³ The **dcbi** instruction should never be used on the G2 core.

For more information, see “Page History Recording” in Chapter 7, “Memory Management,” of the *Programming Environments Manual*.

6.4.2 Page Memory Protection

The G2 core implements page memory protection as it is defined in Chapter 7, “Memory Management,” in the *Programming Environments Manual*.

6.4.3 TLB Description

This section describes the hardware resources provided in the G2 core to facilitate the page address translation process. Note that the hardware implementation of the MMU is not specified by the architecture, and while this description applies to the G2 core, it does not necessarily apply to other processors of this family.

6.4.3.1 TLB Organization

Because the G2 core has two MMUs (IMMU and DMMU) that operate in parallel, some of the MMU resources are shared, and some are actually duplicated (shadowed) in each MMU to maximize performance. Figure 6-7 shows the relationships between these resources within both the IMMU and DMMU and how the various portions of the effective address are used in the address translation process.

While both MMUs can be accessed simultaneously (both sets of segment registers and TLBs can be accessed in the same clock), when there is an exception condition, only one exception is reported at a time. ITLB miss exceptions are reported when there are no more instructions to be dispatched or retired (the pipeline is empty). Refer to Chapter 7, “Instruction Timing,” for more detailed information about the internal pipelines and the reporting of exceptions.

As TLB entries are on-chip copies of PTEs in the page tables in memory, they are similar in structure. TLB entries consist of two words; the high-order word contains the VSID and API fields of the high-order word of the PTE and the low-order word contains the RPN, C bit, WIMG bits, and PP bits (as in the low-order word of the PTE). In order to uniquely identify a TLB entry as the required PTE, the TLB entry also contains five more bits of the page index, EA[10–14] (in addition to the API bits of the PTE).

When an instruction or data access occurs, the effective address is routed to the appropriate MMU. EA[0–3] select 1 of the 16 segment registers and the remaining effective address bits and the virtual address from the segment register is passed to the TLB. EA[15–19] then select two entries in the TLB; the valid bit is checked and EA[10–14], VSID, and API fields (EA[4–9]) for the access are then compared with the corresponding values in the TLB entries. If one of the entries hits, the PP bits are checked for a protection violation, and the C bit is checked. If these bits do not cause an exception, the RPN value is passed to the memory subsystem and the WIMG bits are then used as attributes for the access.

Also, note that the segment registers do not have a valid bit, and so they should also be initialized before translation is enabled.

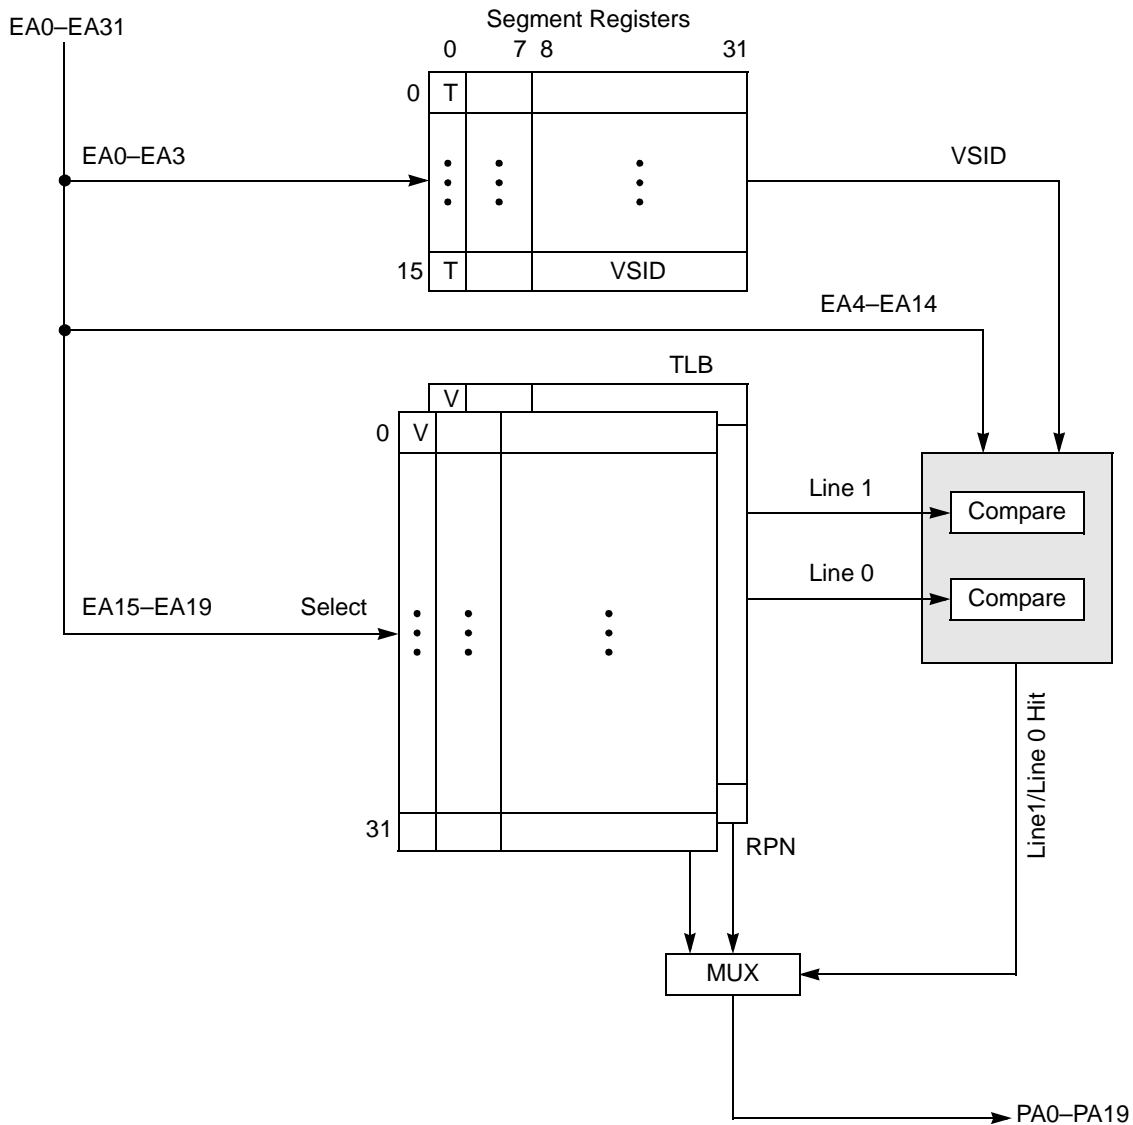


Figure 6-7. Segment Register and TLB Organization

6.4.3.2 TLB Entry Invalidation

For processors, such as the G2 core, that implement TLB structures to maintain on-chip copies of the PTEs that are resident in physical memory, the optional **tlbie** instruction provides a way to invalidate the TLB entries. Note that the execution of the **tlbie** instruction in the G2 core invalidates four entries—both the ITLB entries indexed by EA[15–19] and both the indexed entries of the DTLB.

The architecture allows **tlbie** to optionally enable a TLB invalidate signaling mechanism in hardware so that other processors also invalidate their resident copies of the matching PTE. The G2 core does not signal the TLB invalidation to other processors and does not perform any action when a TLB invalidation is performed by another processor.

The **tlbsync** instruction causes instruction execution to stop if the `core_tlbsync` input signal is also asserted. If `core_tlbsync` is negated, instruction execution may continue or resume after the completion of a **tlbsync** instruction. Section 8.3.11.5, “TLBI Sync (`core_tlbsync`)—Input,” describes the TLB synchronization mechanism in further detail.

The **tlbia** instruction is not implemented on the G2 core and when its opcode is encountered, an illegal instruction program exception is generated. To invalidate all entries of both TLBs, 32 **tlbie** instructions must be executed, incrementing the value in EA[15–19] by 1 each time. See Chapter 8, “Instruction Set,” in the *Programming Environments Manual* for detailed information about the **tlbie** instruction.

6.4.4 Page Address Translation Summary

Figure 6-8 provides the detailed flow for the page address translation mechanism. The figure includes the checking of the N bit in the segment descriptor and then expands on the TLB Hit branch of Figure 6-6. The detailed flow for the TLB Miss branch is described in Section 6.5.1, “Page Table Search Operation—Conceptual Flow.” Note that as in the case of block address translation, if the **dcbz** instruction is attempted to be executed either in write-through mode or as cache-inhibited ($W = 1$ or $I = 1$), the alignment exception is generated. The checking of memory protection violation conditions for page address translation is described in Chapter 7, “Memory Management,” in the *Programming Environments Manual* for 32-bit implementations.

6.5 Page Table Search Operation

As stated earlier, the operating system must synthesize the table search algorithm for setting up the tables. The G2 core TLB miss exception handlers also use this algorithm (with the assistance of some hardware-generated values) to load TLB entries when TLB misses occur, as described in Section 6.5.2, “Implementation-Specific Table Search Operation.”

6.5.1 Page Table Search Operation—Conceptual Flow

The table search process for a processor of this family varies slightly for 64- and 32-bit implementations. The main differences are the address ranges and PTE formats specified. See the *Programming Environments Manual* for the PTE format. An outline of the page table search process performed by a 32-bit implementation is as follows:

1. The 32-bit physical address of the primary PTEG is generated as described in Chapter 7, “Memory Management,” in the *Programming Environments Manual* for 32-bit implementations.
2. The first PTE (PTE0) in the primary PTEG is read from memory. PTE reads should occur with an implied WIM memory/cache mode control bit setting of 0b001. Therefore, they are considered cacheable and burst in from memory and placed in the cache.

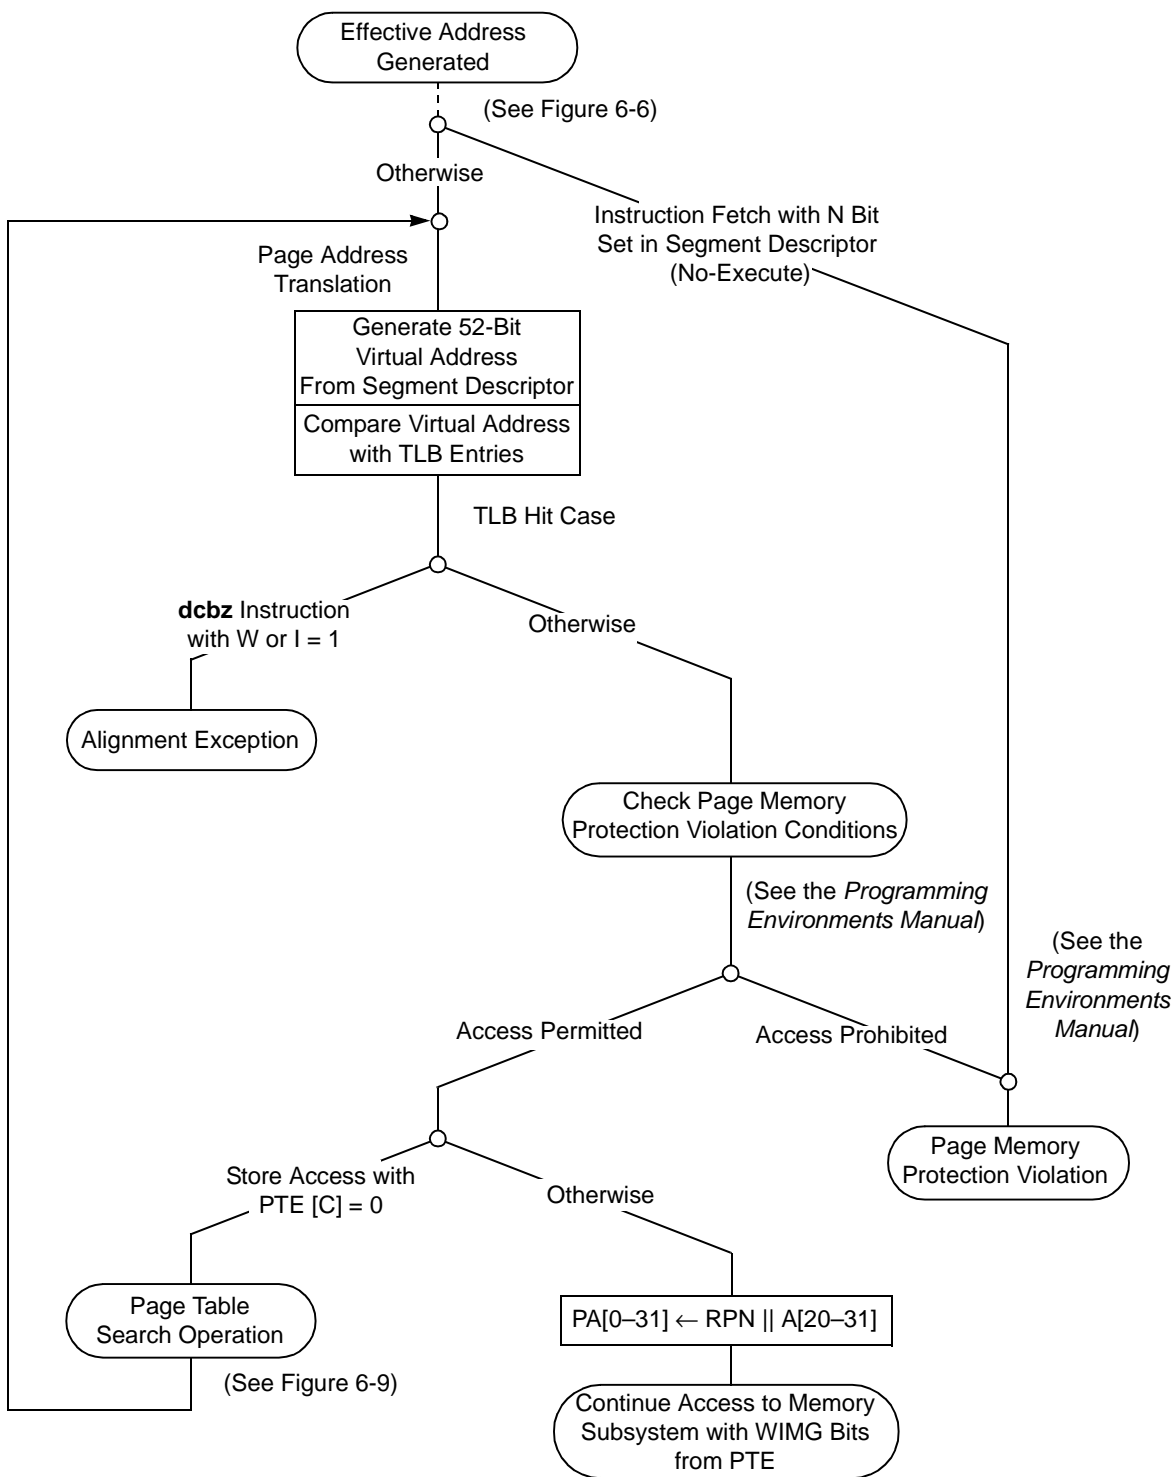


Figure 6-8. Page Address Translation Flow for 32-Bit Implementations—TLB Hit

3. The PTE in the selected PTEG is tested for a match with the virtual page number (VPN) of the access. The VPN is the VSID concatenated with the page index field of the virtual address. For a match to occur, the following must be true:
 - PTE[H] = 0
 - PTE[V] = 1
 - PTE[VSID] = VA[0–23]
 - PTE[API] = VA[24–29]
4. If a match is not found, step 3 is repeated for each of the other seven PTEs in the primary PTEG. If a match is found, the table search process continues as described in step 8. If a match is not found within the eight PTEs of the primary PTEG, the address of the secondary PTEG is generated.
5. The first PTE (PTE0) in the secondary PTEG is read from memory. Again, because PTE reads typically have a WIM bit combination of 0b001, an entire cache line is burst into the on-chip cache.
6. The PTE in the selected secondary PTEG is tested for a match with the virtual page number (VPN) of the access. For a match to occur, the following must be true:
 - PTE[H] = 1
 - PTE[V] = 1
 - PTE[VSID] = VA[0–23]
 - PTE[API] = VA[24–29]
7. If a match is not found, step 6 is repeated for each of the other seven PTEs in the secondary PTEG.
8. If a match is found, the PTE is written into the on-chip TLB and the R bit is updated in the PTE in memory (if necessary). If there is no memory protection violation, the C bit is also updated in memory and the table search is complete.
9. If no match is found in the eight PTEs of the secondary PTEG, the search fails and a page fault exception condition occurs (either an ISI exception or a DSI exception). Note that the software routines that implement this algorithm must synthesize this condition by appropriately setting the SRR1 or DSISR and branching to the ISI or DSI handler routine.

Reads from memory for table search operations should be performed as global (but not exclusive), cacheable operations, and can be loaded into the on-chip cache.

Figure 6-9 and Figure 6-10 provide conceptual flow diagrams of primary and secondary page table search operations as described in the OEA for 32-bit processors. Recall that the architecture allows implementations to perform the page table search operations automatically (in hardware) or with software assistance (may be required), as is the case with the G2 core. Also, the elements in the figure that apply to TLBs are shown as optional because TLBs are not required by the architecture.

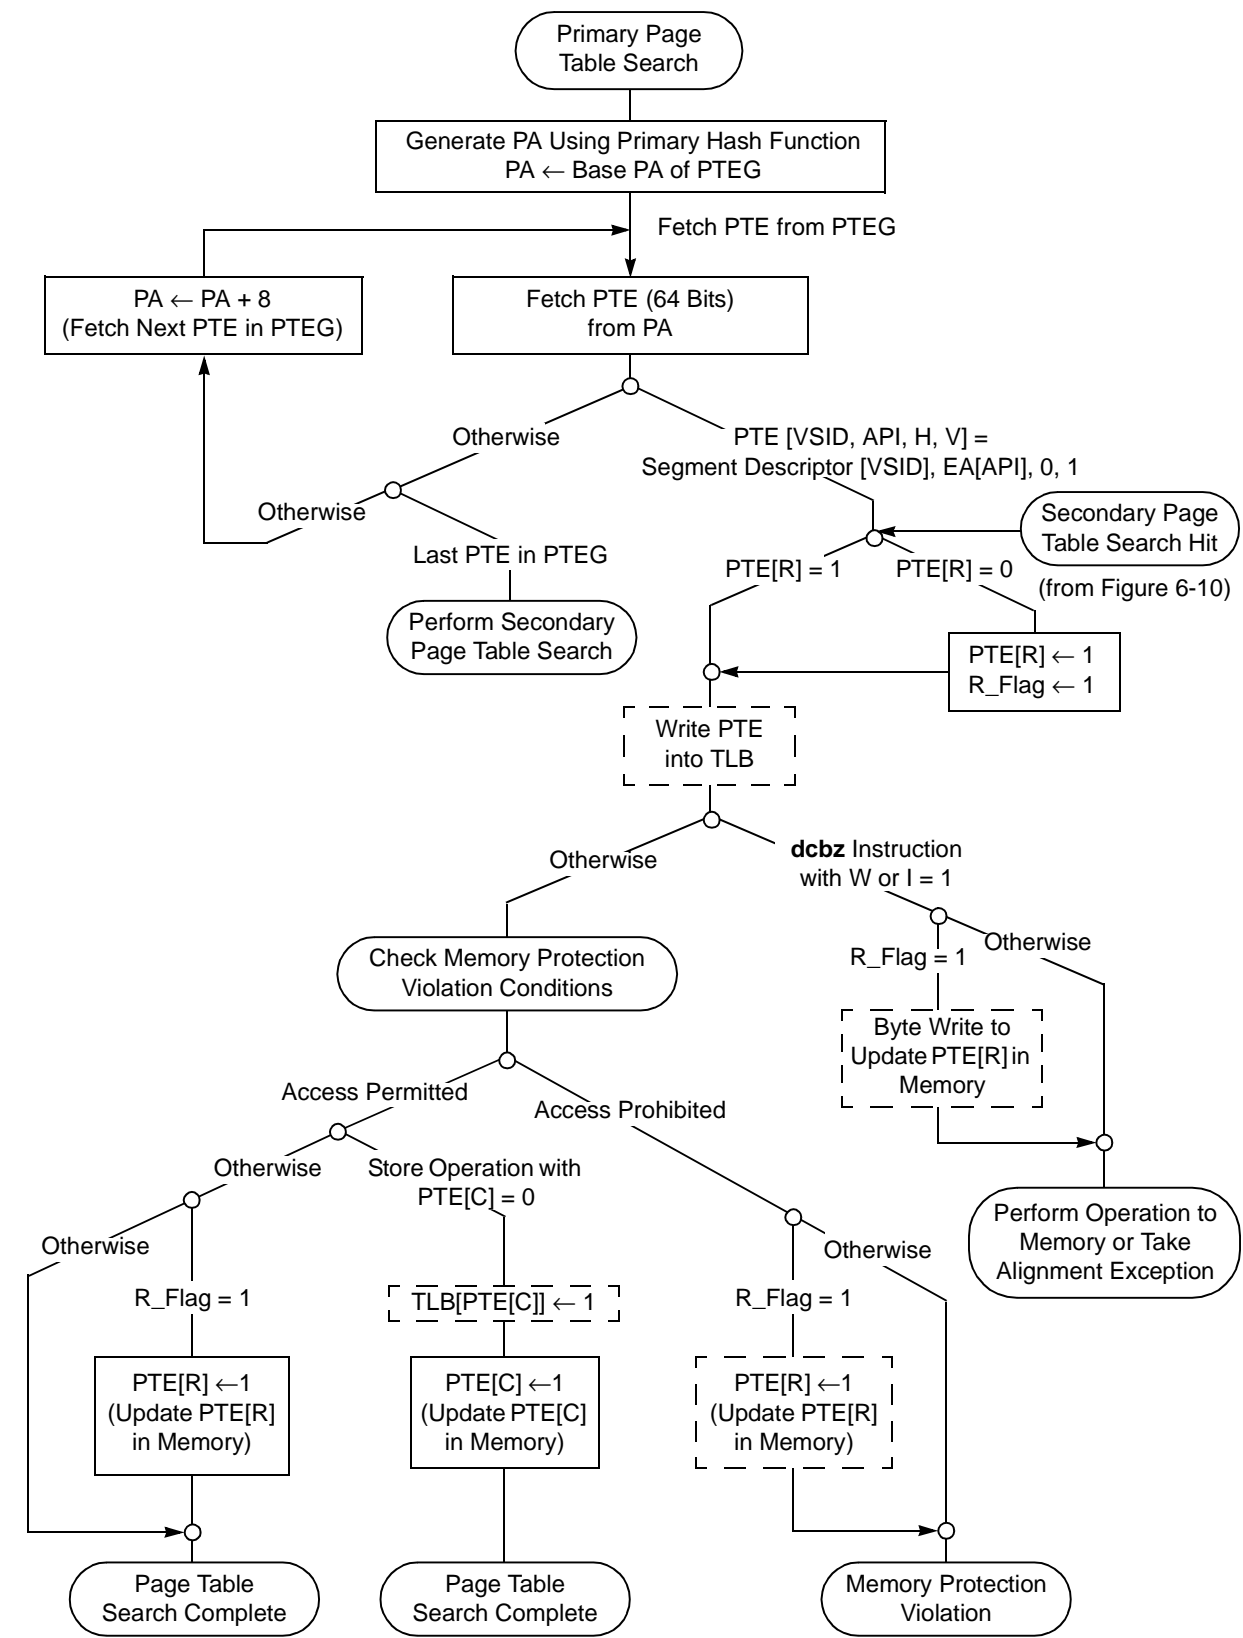


Figure 6-9. Primary Page Table Search—Conceptual Flow

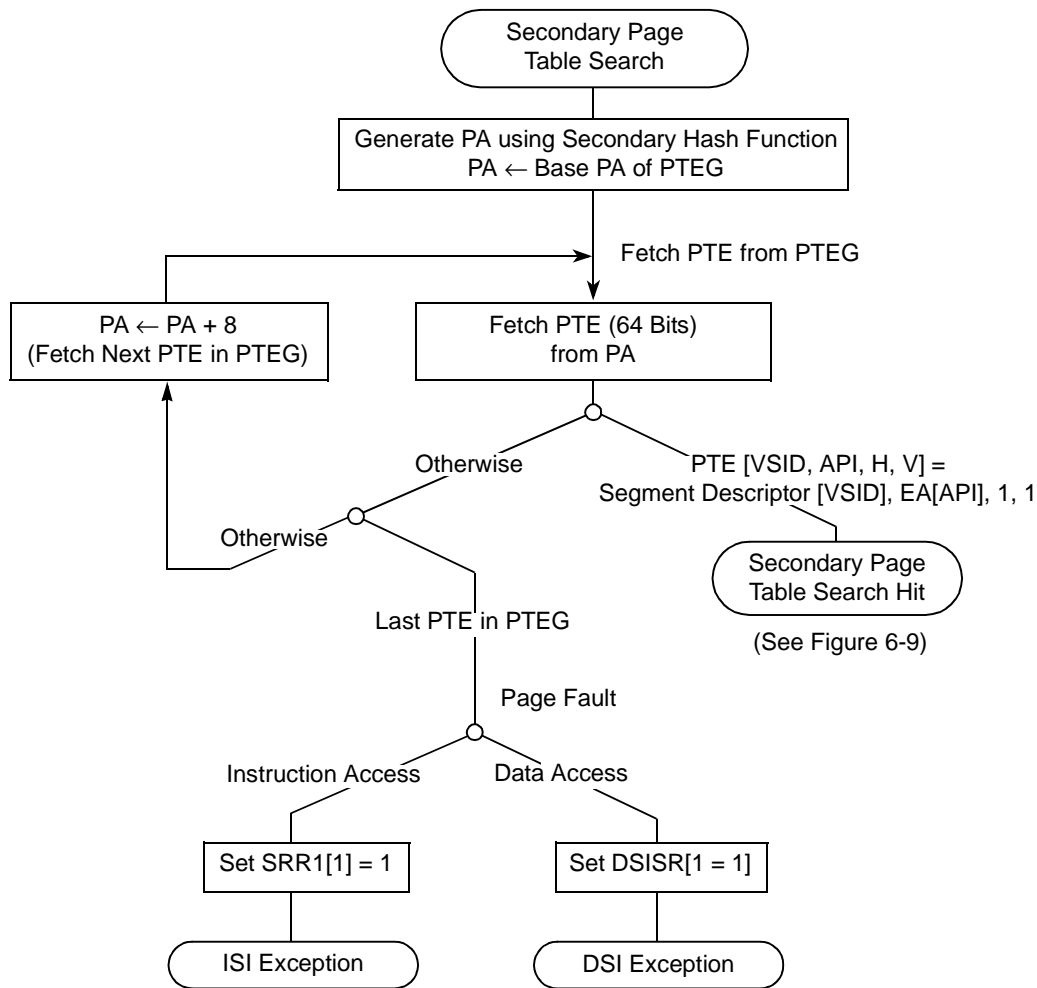


Figure 6-10. Secondary Page Table Search Flow—Conceptual Flow

Figure 6-9 shows the case of a **dcbz** instruction that is executed with $W = 1$ or $I = 1$, and that the **R** bit may be updated in memory (if required) before the operation is performed or the alignment exception occurs. The **R** bit may also be updated by a memory protection violation.

6.5.2 Implementation-Specific Table Search Operation

The G2 core has a set of implementation-specific registers, exceptions, and instructions that facilitate very efficient software searching of the page tables in memory. This section describes those resources and provides three example code sequences that can be used in a G2 core system for an efficient search of the translation tables in software. These three code sequences can be used as handlers for the three exceptions requiring access to the PTEs in the page tables in memory—instruction TLB miss, data TLB miss on load, and data TLB miss on store exceptions.

6.5.2.1 Resources for Table Search Operations

In addition to setting up the translation page tables in memory, the system software must assist the processor in loading PTEs into the on-chip TLBs. When a required TLB entry is not found in the appropriate TLB, the processor vectors to one of the three TLB miss exception handlers so that the software can perform a table search operation and load the TLB. When this occurs, the processor automatically saves information about the access and the executing context. Table 6-9 provides a summary of the implementation-specific exceptions, registers, and instructions that can be used by the TLB miss exception handler software in G2 core systems. Refer to Chapter 5, “Exceptions,” for more information about exception processing.

Table 6-9. Implementation-Specific Resources for Table Search Operations

Resource	Name	Description
Exceptions	Instruction TLB miss exception (vector offset 0x1000)	No matching entry found in ITLB
	Data TLB miss on load exception (vector offset 0x1100)	No matching entry found in DTLB for a load data access
	Data TLB miss on store exception—also caused when changed bit must be updated (vector offset 0x1200)	No matching entry found in DTLB for a store data access or matching DTLB entry has C = 0 and access is a store
Registers	IMISS and DMISS	When a TLB miss exception occurs, the IMISS or DMISS register contains the 32-bit effective address of the instruction or data access that caused the miss exception.
	ICMP and DCMP	The ICMP and DCMP registers contain the word to be compared with the first word of a PTE in the table search software routine to determine if a PTE contains the address translation for the instruction or data access. The contents of ICMP and DCMP are automatically derived by the core when a TLB miss exception occurs.
	HASH1 and HASH2	The HASH1 and HASH2 registers contain the primary and secondary PTEG addresses that correspond to the address causing a TLB miss. These PTEG addresses are automatically derived by the core by performing the primary and secondary hashing function on the contents of IMISS or DMISS, for an ITLB or DTLB miss exception, respectively.
	RPA	The system software loads a TLB entry by loading the second word of the matching PTE entry into the RPA register and then executing the tlbli or tlbld instruction (for loading the ITLB or DTLB, respectively).
Instructions	tlbli rB	Loads the contents of the ICMP and RPA registers into the ITLB entry selected by <ea> and SRR1[WAY]
	tlbld rB	Loads the contents of the DCMP and RPA registers into the DTLB entry selected by <ea> and SRR1[WAY]

In addition, the G2 core contains the following features that do not specifically control the MMU, but that are implemented to increase performance and flexibility in the software table search routines whenever one of the three TLB miss exceptions occurs:

- Temporary GPR0–GPR3. These registers are available as **r0–r3** when MSR[TGPR] is set. The G2 core automatically sets MSR[TGPR] for these cases, allowing these exception handlers to have four registers that are used as scratchpad space, without having to save or restore this part of the machine state that existed when the exception occurred. Note that MSR[TGPR] is cleared when the **rfi** instruction is executed because the old MSR value (with MSR[TGPR] = 0) saved in SRR1 is restored. Refer to Section 6.5.2.2, “Software Table Search Operation,” for code examples that take advantage of these registers.
- Also, the core automatically saves the values of CR[CR0] of the executing context to SRR1[0–3]. Thus, the exception handler can set CR[CR0] bits and branch accordingly in the exception handler routine, without having to save the existing CR[CR0] bits. However, the exception handler must restore these bits to CR[CR0] before executing the **rfi** instruction or branching to the DSI or ISI exception handler. In addition, SRR1[CRF0] must be cleared before branching to the DSI exception handler on a data access page fault. For an instruction access page fault, SRR1[0, 2–3] must be cleared before branching to the ISI handler. See Figure 6-17 for synthesizing a page fault exception when no PTE is found.
- SRR1[D/I] identifies an instruction or data miss, and SRR1[L/S] identifies a load or store miss. SRR1[WAY] identifies the associativity class of the TLB entry selected for replacement by the LRU algorithm. The software can change this value, effectively overriding the replacement algorithm. The SRR1[KEY] bit is used by the table search software to determine if there is a protection violation associated with the access (useful on data write misses for determining if the C bit should be updated in the table). Table 6-10 summarizes the SRR1 bits updated whenever one of the three TLB miss exceptions occurs.

Table 6-10. Implementation-Specific SRR1 Bits

Bits	Name	Function
0–3	CRF0	Condition register field 0 bits
12	KEY	Key for TLB miss (either Ks or Kp from segment register, depending on whether the access is a user or supervisor access).
13	D/I	Set if instruction TLB miss
14	WAY	Next TLB set to be replaced (set per LRU)
15	S/L	Set if data TLB miss was for a load instruction

The key bit saved in SRR1 is derived as follows:

Select KEY from segment register:

If MSR[PR] = 0, KEY = K_s

If MSR[PR] = 1, KEY = K_p

The rest of this section describes the format of the implementation-specific SPRs used by the TLB miss exception handlers. These registers can be accessed by supervisor-level instructions only. Because DMISS, IMISS, DCMP, ICMP, HASH1, HASH2, and RPA are used to access the translation tables for software table search operations, they should only be accessed when address translation is disabled (MSR[IR] = 0 and MSR[DR] = 0). Note that MSR[IR] and MSR[DR] are cleared whenever an exception occurs.

6.5.2.1.1 Data and Instruction TLB Miss Address Registers (DMISS and IMISS)

The DMISS and IMISS registers have the same format as shown in Figure 6-11. They are loaded automatically on a data or instruction TLB miss. The DMISS and IMISS contain the effective page address of the access which caused the TLB miss exception. The contents are used by the processor when calculating the values of HASH1 and HASH2, and by the **tlbld** and **tlbli** instructions when loading a new TLB entry. Note that the core always loads a big-endian address into the DMISS register. These registers are both read- and write-accessible. However, great caution should be used when writing to these registers.

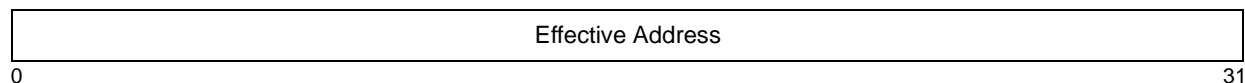


Figure 6-11. DMISS and IMISS Registers

6.5.2.1.2 Data and Instruction TLB Compare Registers (DCMP and ICMP)

The DCMP and ICMP registers are shown in Figure 6-12. These registers contain the first word in the required PTE. The contents are constructed automatically from the contents of the segment registers and the effective address (DMISS or IMISS) when a TLB miss exception occurs. Each PTE read from the tables in memory during the table search process should be compared with this value to determine whether or not the PTE is a match. Upon execution of a **tlbld** or **tlbli** instruction, the contents of the DCMP or ICMP register is loaded into the first word of the selected TLB entry.

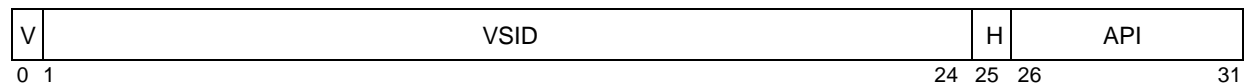


Figure 6-12. DCMP and ICMP Registers

Table 6-11 describes the bit settings for the DCMP and ICMP registers.

Table 6-11. DCMP and ICMP Bit Settings

Bits	Name	Description
0	V	Valid bit. Set by the processor on a TLB miss exception.
1–24	VSID	Virtual segment ID. Copied from VSID field of corresponding segment register.
25	H	Hash function identifier. Cleared by the processor on a TLB miss exception.
26–31	API	Abbreviated page index. Copied from API of effective address.

6.5.2.1.3 Primary and Secondary Hash Address Registers (HASH1 and HASH2)

HASH1 and HASH2 contain the physical addresses of the primary and secondary PTEGs for the access that caused the TLB miss exception. Only bits 7–25 differ between them. For convenience, the processor automatically constructs the full physical address by routing bits 0–6 of SDR1 into HASH1 and HASH2 and clearing the lower six bits. These registers are read-only and are constructed from the contents of the DMISS or IMISS register. The format for HASH1 and HASH2 is shown in Figure 6-13.

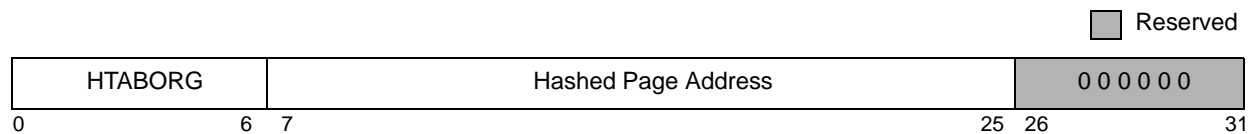


Figure 6-13. HASH1 and HASH2 Registers

Table 6-12 describes the bit settings of the HASH1 and HASH2 registers.

Table 6-12. HASH1 and HASH2 Bit Settings

Bits	Name	Description
0–6	HTABORG[0–6]	Copy of the upper 7 bits of the HTABORG field from SDR1
7–25	Hashed page address	Address bits 7–25 of the PTEG to be searched
26–31	—	Reserved

6.5.2.1.4 Required Physical Address Register (RPA)

The RPA is shown in Figure 6-14. During a page table search operation, the software must load the RPA with the second word of the correct PTE. When the **tlbld** or **tlbli** instruction is executed, data from IMISS and ICMP (or DMISS and DCMP) and the RPA registers is merged and loaded into the selected TLB entry. The TLB entry is selected by the effective address of the access (loaded by the table search software from the DMISS or IMISS register) and SRR1[WAY].

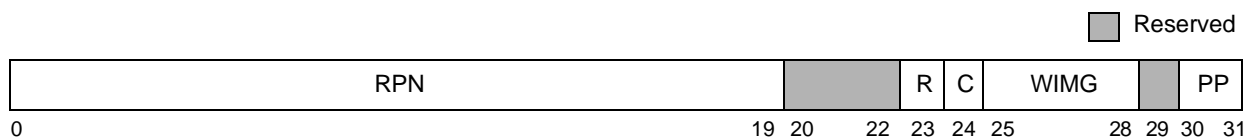


Figure 6-14. Required Physical Address (RPA) Register

Table 6-13 describes the bit settings of the RPA register.

Table 6-13. RPA Bit Settings

Bits	Name	Description
0–19	RPN	Physical page number from PTE
20–22	—	Reserved
23	R	Referenced bit from PTE
24	C	Changed bit from PTE
25–28	WIMG	Memory/cache access attribute bits
29	—	Reserved
30–31	PP	Page protection bits from PTE

6.5.2.2 Software Table Search Operation

When a TLB miss occurs, the instruction or data MMU loads IMISS or DMISS, with the effective address of the access. The processor completes all instructions ahead of the instruction that caused the exception, status information is saved in SRR1, and one of the three TLB miss exceptions is taken. In addition, the processor loads ICMP or DCMP with the value to be compared with the first word of PTEs in the tables in memory.

The software should then access the first PTE at the address pointed to by HASH1. The first word of the PTE should be loaded and compared to the contents of DCMP or ICMP. If there is a match, the required PTE has been found and the second word of the PTE is loaded from memory into RPA. Then the **tlbli** or **tblld** instruction is executed, which loads the contents of ICMP or DCMP and RPA into the selected TLB entry. The TLB entry is selected by the effective address of the access and SRR1[WAY].

If the comparison does not match, the PTEG address is incremented to point to the next PTE in the table, and the above sequence is repeated. If none of the eight PTEs in the primary PTEG matches, the sequence is then repeated using the secondary PTEG (at the address contained in HASH2).

If the PTE is also not found in the eight entries of the secondary page table, a page fault condition exists and a page fault exception must be synthesized. Thus, the appropriate bits must be set in SRR1 (or DSISR) and the TLB miss handler must branch to either the ISI or DSI exception handler, which handles the page fault condition.

The following section provides a flow diagram outlining some example software that can be used to handle the three TLB miss exceptions and sample assembly language that implements that flow.

6.5.2.2.1 Flow for Example Exception Handlers

Figure 6-15 shows the flow for the example TLB miss exception handlers. The flow shown is common for the three exception handlers, except that the IMISS and ICMP registers are used for the instruction TLB miss exception while the DMISS and DCMP registers are used for the two data TLB miss exceptions. Also, for the cases of store instructions that cause either a TLB miss or require a table search operation to update the C bit, the flow shows that the C bit is set in both the TLB entry and PTE in memory. Finally, in the case of a page fault (no PTE found in the table search operation), the setup for the ISI or DSI exception is slightly different for these two cases.

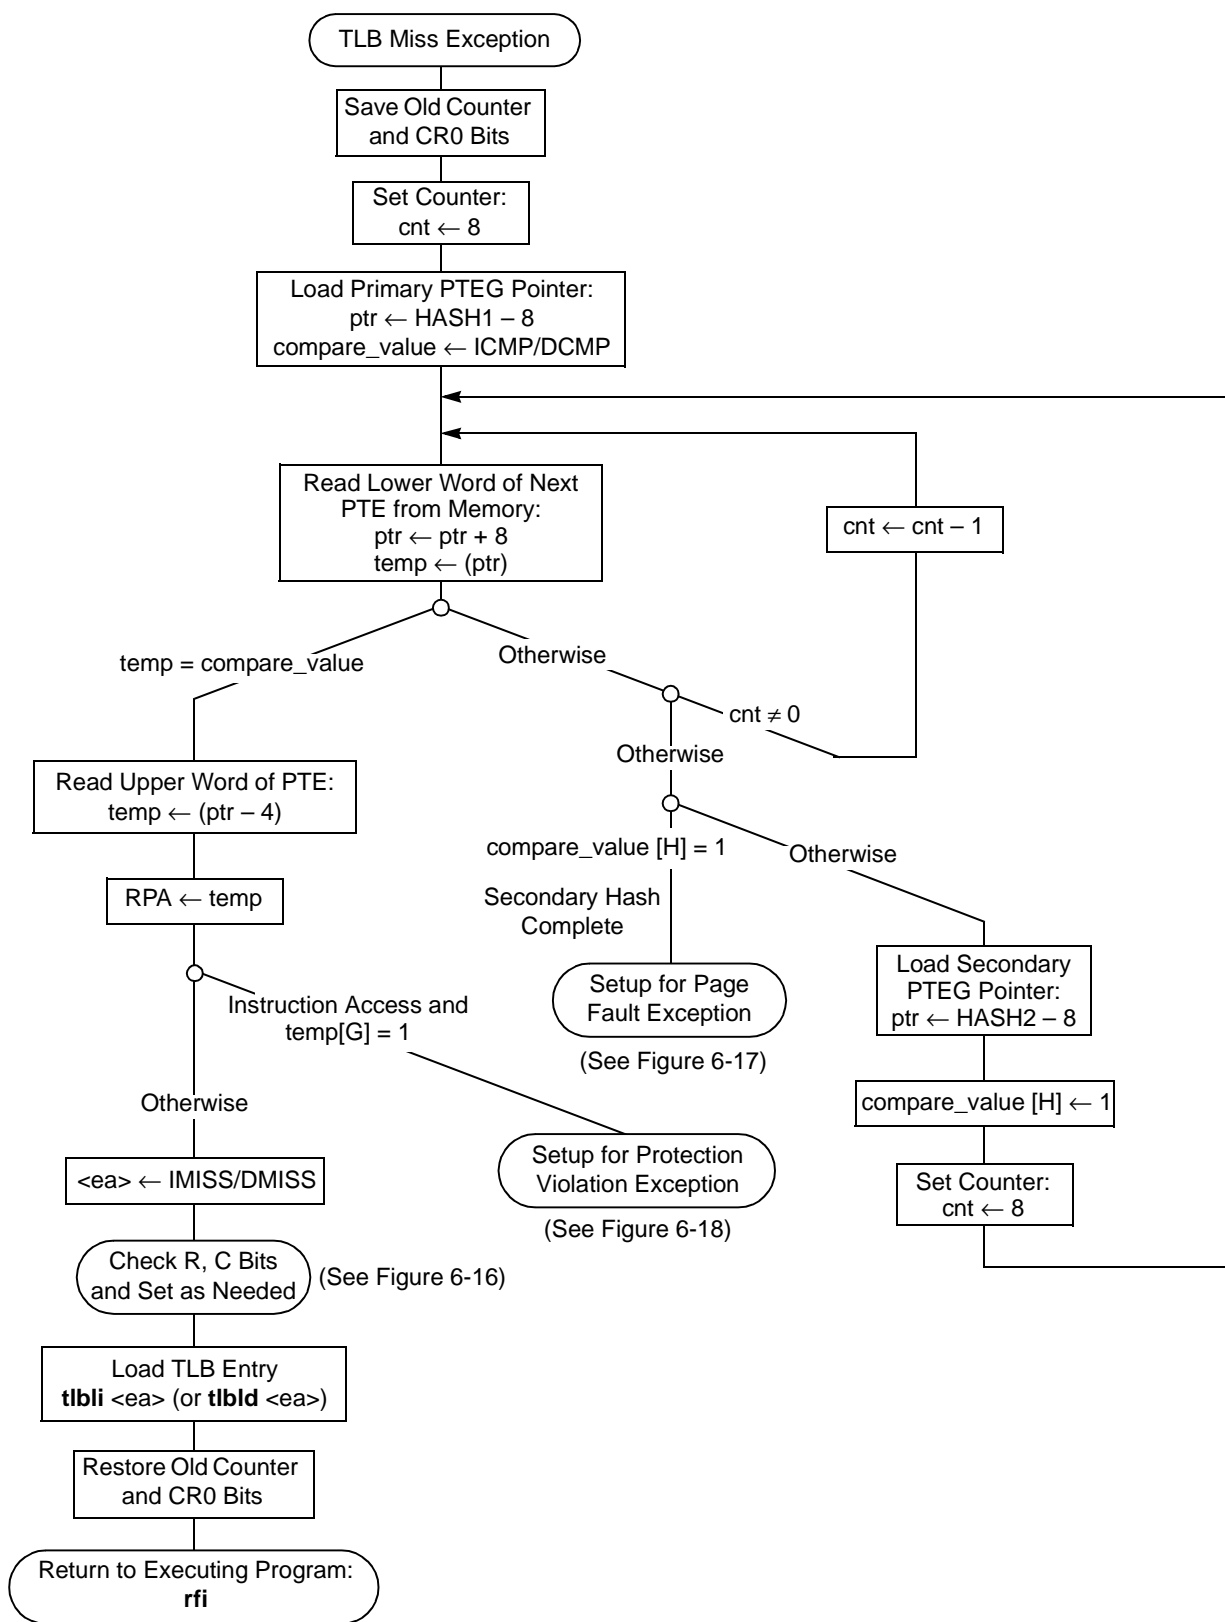


Figure 6-15. Flow for Example Software Table Search Operation

The flow for checking the R and C bits and setting them appropriately is shown in Figure 6-16.

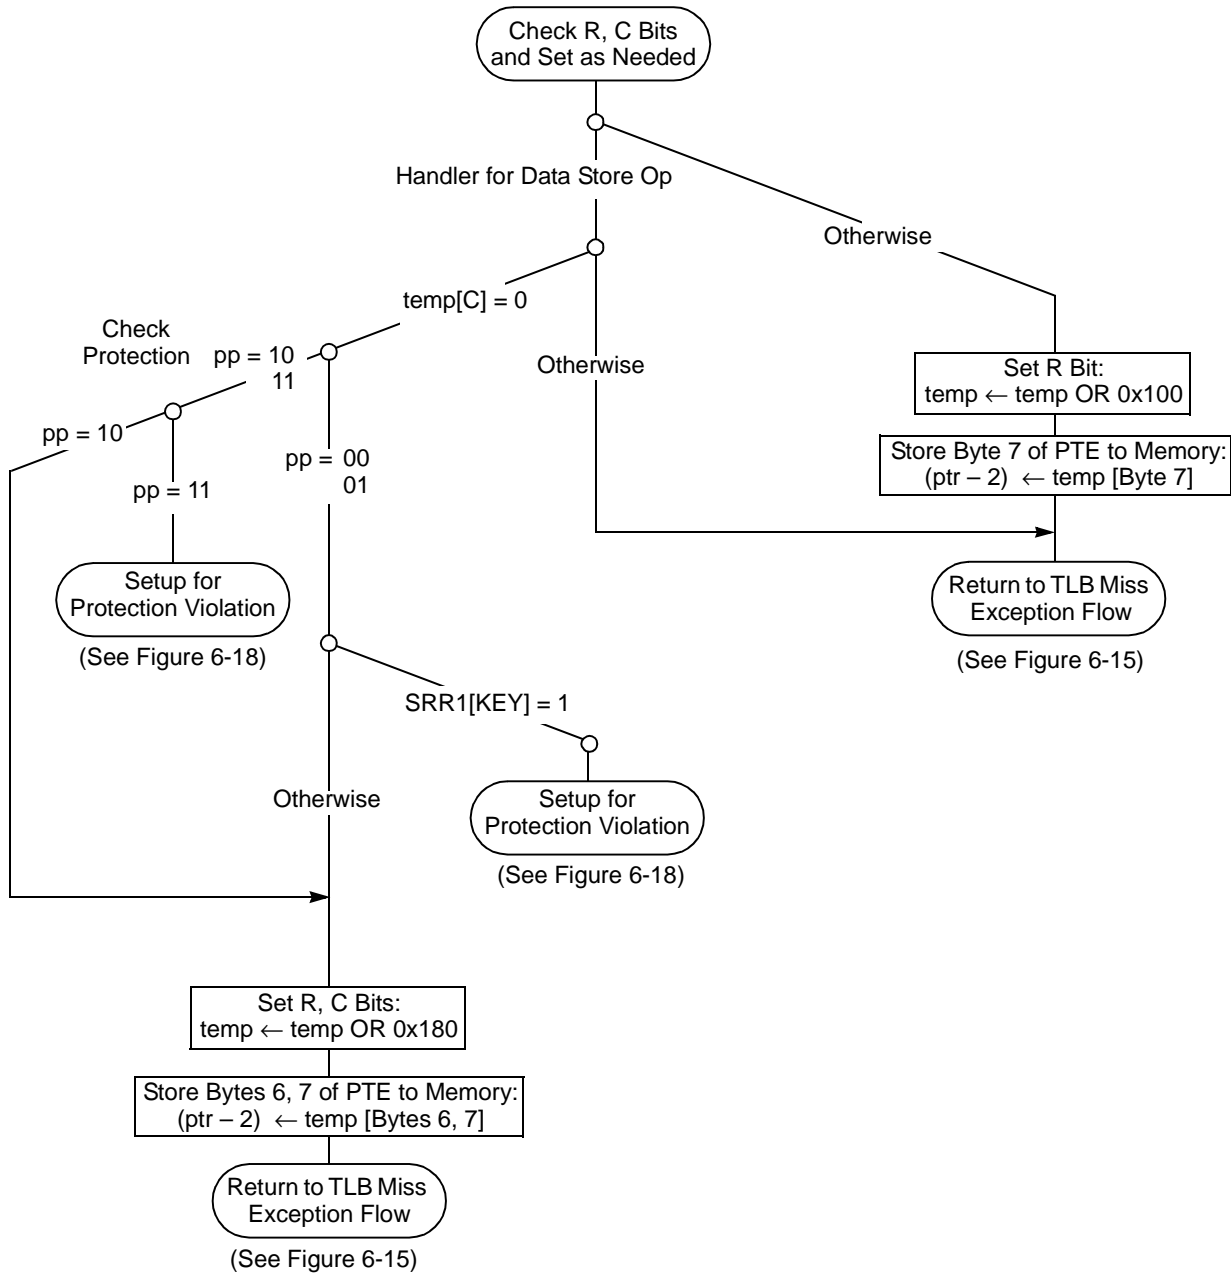


Figure 6-16. Check and Set R and C Bit Flow

Figure 6-17 shows the flow for synthesizing a page fault exception when no PTE is found.

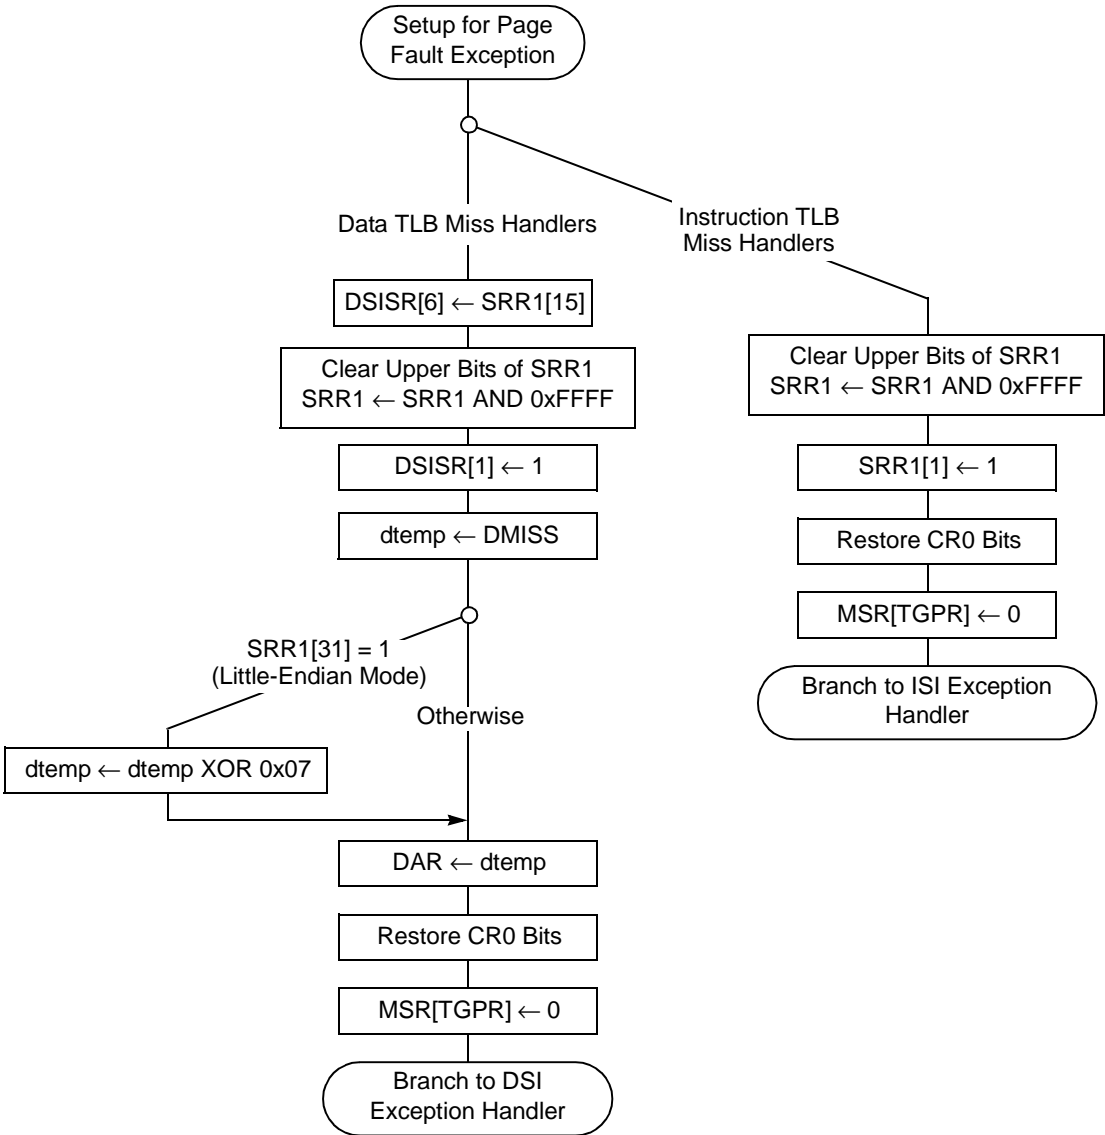


Figure 6-17. Page Fault Setup Flow

Figure 6-18 shows the flow for managing the cases of a TLB miss on an instruction access to guarded memory, and a TLB miss when $C = 0$ and a protection violation exists. The setup for these protection violation exceptions is very similar to that of page fault conditions (as shown in Figure 6-17) except that different bits in SRR1 (and DSISR) are set.

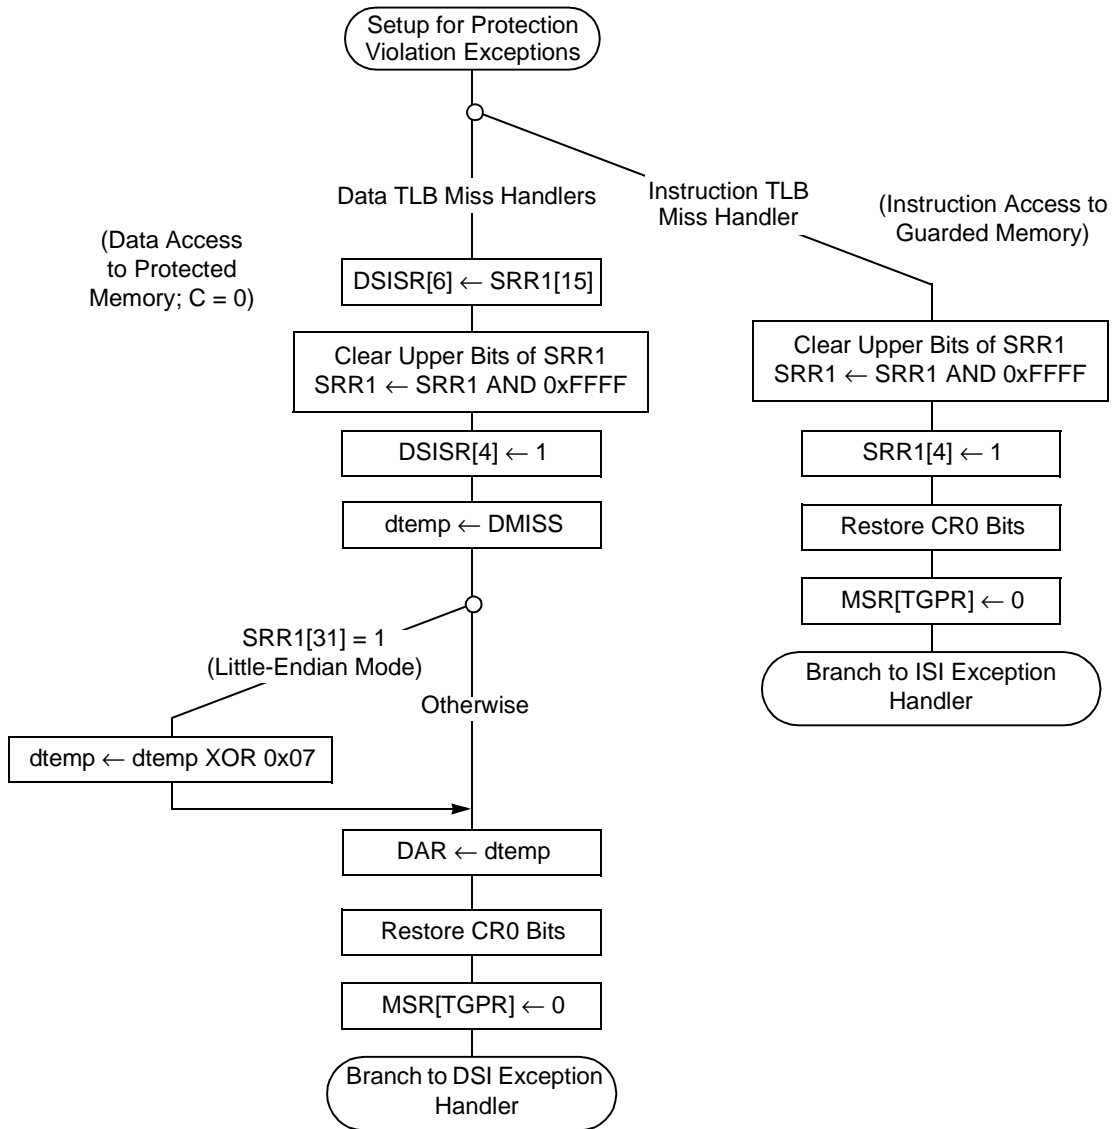


Figure 6-18. Setup for Protection Violation Exceptions

6.5.2.2.2 Code for Example Exception Handlers

This section provides assembly language examples that implement the flow diagrams described above. Note that although these routines fit into a few cache lines, they are supplied only as functional examples; they could be further optimized for faster performance.

```
# TLB software load for G2 core
#
# New Instructions:
#     tlbld          - write the dtlb with the pte in rpa reg
#     tlbli          - write the itlb with the pte in rpa reg
# New SPRs
#     dmiss          - address of dstream miss
#     imiss          - address of istream miss
#     hash1          - address primary hash PTEG address
#     hash2          - returns secondary hash PTEG address
#     iCmp           - returns the primary istream compare value
#     dCmp           - returns the primary dstream compare value
#     rpa            - the second word of pte used by tlbld
#
# gpr r0..r3 are shadowed
#
# there are three flows.
#     tlbDataMiss    - tlb miss on data load
#     tlbCeq0        - tlb miss on data store or store with tlb change bit
#                     == 0
#     tlbInstrMiss   - tlb miss on instruction fetch
#+
# place labels for rel branches
#-
#.machine PPC_603e
.set    r0, 0
.set    r1, 1
.set    r2, 2
.set    r3, 3
.set    dMiss, 976
.set    dCmp, 977
.set    hash1, 978
.set    hash2, 979
.set    iMiss, 980
.set    iCmp, 981
.set    rpa, 982
.set    c0, 0
.set    dar, 19
.set    dsisr, 18
.set    srr0, 26
.set    srr1, 27
.
.csect tlbmiss[PR]
vec0:
.globl vec0
```

```

.org    vec0+0x300

vec300:
.org    vec0+0x400
vec400:
#+
# Instruction TB miss flow
# Entry:
#     Vec = 1000
#     srr0      -> address of instruction that missed
#     srr1      -> 0:3=cr0 4=lru way bit 16:31 = saved MSR
#     msr<tgpr> -> 1
#     iMiss     -> ea that missed
#     iCmp      -> the compare value for the va that missed

#     hash1     -> pointer to first hash pteg
#     hash2     -> pointer to second hash pteg
#
# Register usage:
#     r0 is saved counter
#     r1 is junk
#     r2 is pointer to pteg
#     r3 is current compare value

.org    vec0+0x1000

tlbInstrMiss:
    mfspr    r2, hash1      # get first pointer
    addi     r1, 0, 8       # load 8 for counter
    mfctr    r0             # save counter
    mfspr    r3, iCmp      # get first compare value
    addi     r2, r2, -8     # pre dec the pointer
im0:    mtctr    r1         # load counter
im1:    lwzu    r1, 8(r2)   # get next pte
        cmp     c0, r1, r3  # see if found pte
        bdnz   eq, im1     # dec count br if cmp ne and if count not zero

        bne    instrSecHash # if not found set up second hash or exit
        l      r1, +4(r2)   # load tlb entry lower-word
        andi.  r3, r1, 8    # check G bit
        bne    doISIP      # if guarded, take an ISI
        mtctr  r0          # restore counter
        mfspr  r0, iMiss    # get the miss address for the tlbli
        mfspr  r3, srr1     # get the saved cr0 bits
        mtcrf  0x80, r3    # restore CR0
        mtspr  rpa, r1     # set the pte
        ori   r1, r1, 0x100 # set reference bit
        srwi  r1, r1, 8    # get byte 7 of pte
        tlbli r0          # load the itlb
        stb   r1, +6(r2)   # update page table
        rfi                    # return to executing program

#+
# Register usage:
#     r0 is saved counter
#     r1 is junk
    
```



Page Table Search Operation

```
#      r2 is pointer to pteg
#      r3 is current compare value
#-

instrSecHash:
    andi.   r1, r3, 0x0040 # see if we have done second hash
    bne     doISI          # if so, go to ISI exception
    mfspr   r2, hash2     # get the second pointer
    ori     r3, r3, 0x0040 # change the compare value
    addi    r1, 0, 8      # load 8 for counter
    addi    r2, r2, -8    # pre dec for update on load
    b       im0           # try second hash

#+
# entry Not Found: synthesize an ISI exception
# guarded memory protection violation: synthesize an ISI exception
# Entry:
#      r0 is saved counter
#      r1 is junk
#      r2 is pointer to pteg
#      r3 is current compare value
#
doISIp:
    mfspr   r3, srr1      # get srr1
    andi.   r2, r3, 0xffff # clean upper srr1
    addis   r2, r2, 0x0800 # or in srr<4> = 1 to flag prot violation
    b       isil:

doISI:
    mfspr   r3, srr1      # get srr1
    andi.   r2, r3, 0xffff # clean srr1
    addis   r2, r2, 0x4000 # or in srr1<1> = 1 to flag pte not found
isil      mtctr   r0       # restore counter
           mtspr   srr1, r2 # set srr1
           mfmsr   r0       # get msr
           xoris   r0, r0, 0x8000 # flip the msr<tgpr> bit
           mtcrf   0x80, r3 # restore CR0
           mtmsr   r0       # flip back to the native gprs
           b       vec400   # go to instr. access exception

#
#+
# Data TLB miss flow
# Entry:
#      Vec = 1100
#      srr0      -> address of instruction that caused data tlb miss
#      srr1      -> 0:3=cr0 4=lru way bit 5=1 if store 16:31 = saved MSR
#      msr<tgpr> -> 1
#      dMiss     -> ea that missed
#      dCmp      -> the compare value for the va that missed
#      hash1     -> pointer to first hash pteg
#      hash2     -> pointer to second hash pteg
#
# Register usage:
#      r0 is saved counter
#      r1 is junk
```

```

#       r2 is pointer to pteg
#       r3 is current compare value
#-

.csect  tlbmiss[PR]
.org    vec0+0x1100

tlbDataMiss:
    mfspr  r2, hash1      # get first pointer
    addi   r1, 0, 8      # load 8 for counter
    mfctr  r0             # save counter
    mfspr  r3, dCmp      # get first compare value
    addi   r2, r2, -8    # pre dec the pointer
dm0:    mtctr  r1         # load counter
dm1:    lwzu   r1, 8(r2)  # get next pte
    cmp    c0, r1, r3    # see if found pte
    bdnzfb 0, dm1       # dec count br if cmp ne and if count not zero
    bne    dataSecHash  # if not found set up second hash or exit
    l      r1, +4(r2)    # load tlb entry lower-word
    mtctr  r0           # restore counter
    mfspr  r0, dMiss     # get the miss address for the tlbld
    mfspr  r3, srr1     # get the saved cr0 bits
    mtcrfb 0x80, r3     # restore CR0
    mtspr  rpa, r1     # set the pte
    ori    r1, r1, 0x100 # set reference bit
    srw    r1, r1, 8    # get byte 7 of pte
    tlbld  r0           # load the dtlb
    stb    r1, +6(r2)   # update page table
    rfi                          # return to executing program

#+
# Register usage:
#   r0 is saved counter
#   r1 is junk
#   r2 is pointer to pteg
#   r3 is current compare value
#-

dataSecHash:
    andi.  r1, r3, 0x0040 # see if we have done second hash
    bne    doDSI          # if so, go to DSI exception
    mfspr  r2, hash2      # get the second pointer
    ori    r3, r3, 0x0040 # change the compare value
    addi   r1, 0, 8      # load 8 for counter
    addi   r2, r2, -8    # pre dec for update on load
    b      dm0           # try second hash

#

#+
# C=0 in dtlb and dtlb miss on store flow
# Entry:
#   Vec = 1200
#   srr0    -> address of store that caused the exception
#   srr1    -> 0:3=cr0 4=lru way bit 5=1 16:31 = saved MSR
#   msr<tgpr> -> 1
    
```



Page Table Search Operation

```
#      dMiss      -> ea that missed
#      dCmp       -> the compare value for the va that missed
#      hash1      -> pointer to first hash pteg
#      hash2      -> pointer to second hash pteg
#
# Register usage:
#      r0 is saved counter
#      r1 is junk
#      r2 is pointer to pteg
#      r3 is current compare value
#-

.csect  tlbmiss[PR]
.org    vec0+0x1200

tlbCeq0:
    mfspr    r2, hash1      # get first pointer
    addi     r1, 0, 8       # load 8 for counter
    mfctr    r0             # save counter
    mfspr    r3, dCmp      # get first compare value
    addi     r2, r2, -8     # pre dec the pointer
ceq0:      mtctr    r1       # load counter
ceq1:      lwzu     r1, 8(r2) # get next pte
           cmp      c0, r1, r3 # see if found pte
           bdnzfb  0, ceq1    # dec count br if cmp ne and if count not zero
           bne     cEq0SecHash # if not found set up second hash or exit
           l       r1, +4(r2)  # load tlb entry lower-word
           andi.   r3, r1, 0x80 # check the C-bit
           beq     cEq0ChkProt # if (C==0) go check protection modes
ceq2:      mtctr    r0       # restore counter
           mfspr   r0, dMiss   # get the miss address for the tlbld
           mfspr   r3, srr1    # get the saved cr0 bits
           mtcrf   0x80, r3    # restore CR0
           mtspr   rpa, r1     # set the pte
           tlbld   r0          # load the dtlb
           rfi     # return to executing program

#+
# Register usage:
#      r0 is saved counter
#      r1 is junk
#      r2 is pointer to pteg
#      r3 is current compare value
#-

cEq0SecHash:
    andi.    r1, r3, 0x0040 # see if we have done second hash
    bne     doDSI          # if so, go to DSI exception
    mfspr   r2, hash2      # get the second pointer
    ori     r3, r3, 0x0040 # change the compare value
    addi    r1, 0, 8       # load 8 for counter
    addi    r2, r2, -8     # pre dec for update on load
    b       ceq0          # try second hash
```

```

#+
# entry found and PTE(c-bit==0):
# (check protection before setting PTE(c-bit))
# Register usage:
#     r0 is saved counter
#     r1 is PTE entry
#     r2 is pointer to pteg
#     r3 is trashed
#-

ceq0ChkProt:
    rlwinm. r3,r1,30,0,1    # test PP
    bge-   chk0            # if (PP==00 or PP==01) goto chk0:
    andi.  r3,r1,1         # test PP[0]
    beq+   chk2            # return if PP[0]==0
    b      doDSIp         # else DSIP

chk0:    mfspr   r3,srr1    # get old msr
    andis. r3,r3,0x0008    # test the KEY bit (SRR1-bit 12)
    beq    chk2            # if (KEY==0) goto chk2:
    b      doDSIp         # else DSIP

chk2:    ori     r1, r1, 0x180 # set reference and change bit
    sth     r1, 6(r2)        # update page table
    b      ceq2            # and back we go

#

#+
# entry Not Found: synthesize a DSI exception
# Entry:
#     r0 is saved counter
#     r1 is junk
#     r2 is pointer to pteg
#     r3 is current compare value
#

doDSI:
    mfspr   r3, srr1        # get srr1
    rlwinm  r1, r3, 9,6,6   # get srr1<flag> to bit 6 for load/store, zero
                          # rest
    addis   r1, r1, 0x4000  # or in dsisr<1> = 1 to flag pte not found
    b      dsil:

doDSIp:
    mfspr   r3, srr1        # get srr1
    rlwinm  r1, r3, 9,6,6   # get srr1<flag> to bit 6 for load/store, zero
                          # rest
    addis   r1, r1, 0x0800  # or in dsisr<4> = 1 to flag prot violation

dsil:    mtctr   r0          # restore counter
    andi.   r2, r3, 0xffff  # clear upper bits of srr1
    mtspr   srr1, r2        # set srr1
    mtspr   dsisr, r1       # load the dsisr
    mfspr   r1, dMiss       # get miss address
    rlwinm. r2,r2,0,31,31   # test LE bit
    
```

```

        beq      dsi2:          # if little endian then:
        xor      r1,r1,0x07    # de-mung the data address

dsi2:   mtspr    dar, r1       # put in dar
        mfmsr   r0            # get msr
        xoris   r0, r0, 0x2    # flip the msr<tgpr> bit
        mtcrf   0x80, r3      # restore CR0
        mtmsr   r0            # flip back to the native gprs
        b       vec300        # branch to DSI exception

```

6.5.3 Page Table Updates

TLBs are defined as noncoherent caches of the PTEs. TLB entries must be flushed explicitly with the TLB invalidate entry instruction (**tlbie**) whenever the corresponding PTE is modified. Because the G2 core is intended primarily for uniprocessor environments, it does not provide coherency checking for TLBs between multiple processors. If the G2 core is used in a multiprocessor environment where TLB coherency is required, synchronization must be implemented in software.

Processors may write referenced and changed bits with unsynchronized, atomic byte store operations. Note that each V, R, and C bits reside in a distinct byte of a PTE. Therefore, extreme care must be taken to use byte writes when updating only one of these bits.

Explicitly altering certain MSR bits (using the **mtmsr** instruction), PTEs, or certain system registers, may have the side effect of changing the effective or physical addresses from which the current instruction stream is being fetched. This kind of side effect is defined as an implicit branch. Implicit branches are not supported and an attempt to perform one causes boundedly-undefined results. Therefore, PTEs must not be changed in a manner that causes an implicit branch. Chapter 2, “Register Set,” in the *Programming Environments Manual*, lists the possible implicit branch conditions that can occur when system registers and MSR bits are changed.

6.5.4 Segment Register Updates

Synchronization requirements for using the move to segment register instructions (**mtsr** and **mtsrin**) are described in “Synchronization Requirements for Special Registers and for Lookaside Buffers” in Chapter 2, “Register Set,” in the *Programming Environments Manual*.

Chapter 7

Instruction Timing

This chapter describes how the G2 core processor fetches, dispatches, and executes instructions and how it reports the results of instruction execution. It gives detailed descriptions of how the G2 core execution units work, and how those units interact with other parts of the processor, such as the instruction fetching mechanism, register files, and caches. It gives examples of instruction sequences, showing potential bottlenecks and how to minimize their effects. Finally, it includes tables that identify the unit that executes each instruction implemented on the core, the latency for each instruction, and other information that is useful for the assembly language programmer.

7.1 Terminology and Conventions

This section provides an alphabetical glossary of terms used in this chapter. These definitions are provided as a review of commonly used terms and as a way to point out specific ways these terms are used in this chapter.

- **Branch prediction**—The process of guessing whether a branch will be taken. Such predictions can be correct or incorrect; the term predicted as it is used here does not imply that the prediction is correct (successful). The PowerPC architecture defines a means for static branch prediction as part of the instruction encoding.
- **Branch resolution**—The determination of whether a branch is taken or not taken. A branch is said to be resolved when the processor can determine which instruction path to take. If the branch is resolved as predicted, the instructions following the predicted branch that may have been speculatively executed can complete (see completion). If the branch is not resolved as predicted, instructions on the mispredicted path, and any results of speculative execution, are purged from the pipeline and fetching continues from the nonpredicted path.
- **Completion**—Completion occurs when an instruction has finished executing, written back any results, and is removed from the completion queue (CQ). When an instruction completes, it is guaranteed that this instruction and all previous instructions can cause no exceptions.
- **Fall-through (branch fall-through)**—A not-taken branch. On the G2 core, fall-through branch instructions are removed from the instruction stream at dispatch. That is, these instructions are allowed to fall through the instruction queue through

the dispatch mechanism, without either being passed to an execution unit and or given a position in the CQ.

- **Fetch**—The process of bringing instructions from memory (such as a cache or system memory) into the instruction queue.
- **Finish**—Finishing occurs in the last cycle of execution. In this cycle, the CQ entry is updated to indicate that the instruction has finished executing.
- **Folding (branch folding)**—The replacement of a branch instruction with target instructions and any instructions along the not-taken path, when a branch is either taken or predicted as taken.
- **Latency**—The number of clock cycles necessary to execute an instruction and make ready the results of that execution for a subsequent instruction.
- **Pipeline**—In the context of instruction timing, the term pipeline refers to the interconnection of the stages. The events necessary to process an instruction are broken into several cycle-length tasks to allow work to be performed on several instructions simultaneously—analogue to an assembly line. As an instruction is processed, it passes from one stage to the next. When it does, the stage becomes available for the next instruction.

Although an individual instruction may take many cycles to complete (the number of cycles is called instruction latency), pipelining makes it possible to overlap the processing so that the throughput (number of instructions completed per cycle) is greater than if pipelining were not implemented.

- **Program order**—The order of instructions in an executing program. More specifically, this term is used to refer to the original order in which program instructions are fetched into the instruction queue from the cache.
- **Rename register**—Temporary buffers used by instructions that have finished execution but have not completed.
- **Reservation station**—A buffer between the dispatch and execute stages that allows instructions to be dispatched even though the results of instructions on which the dispatched instruction may depend are not available.
- **Retirement**—Removal of the completed instruction from the CQ.
- **Stage**—The term stage is used in two different senses, depending on whether the pipeline is being discussed as a physical entity or a sequence of events. In the latter case, a stage is an element in the pipeline during which certain actions are performed, such as decoding the instruction, performing an arithmetic operation, or writing back the results. A stage is typically described as taking a processor clock cycle to perform its operation; however, some events (such as dispatch and write-back) happen instantaneously, and may be thought to occur at the end of the stage.

An instruction can spend multiple cycles in one stage. An integer multiply, for example, takes multiple cycles in the execute stage. When this occurs, subsequent instructions may stall.

In some cases, an instruction may also occupy more than one stage simultaneously, especially in the sense that a stage can be seen as a physical resource—for example, when instructions are dispatched they are assigned a place in the CQ at the same time they are passed to the execute stage. They can be said to occupy both the complete and execute stages in the same clock cycle.

- **Stall**—An occurrence when an instruction cannot proceed to the next stage.
- **Store Queue**—Holds store operations that have not been committed to memory, resulting from completed or retired instructions.
- **Superscalar**—A superscalar processor is one that can dispatch multiple instructions concurrently from a conventional linear instruction stream. In a superscalar implementation, multiple instructions can be in the same stage at the same time.
- **Throughput**—A measure of the number of instructions that are processed per cycle. For example, a series of double-precision floating-point multiply instructions has a throughput of one instruction per clock cycle.
- **Write-back**—Write-back (in the context of instruction handling) occurs when a result is written from the rename registers into the architectural registers (typically the GPRs and FPRs or the store queue).

7.2 Instruction Timing Overview

The G2 core design minimizes average instruction execution latency, the number of clock cycles it takes to fetch, decode, dispatch, and execute instructions and make the results available for a subsequent instruction. Some instructions, such as loads and stores, access memory and require additional clock cycles between the execute phase and the write-back phase. These latencies vary depending on whether the access is to cacheable or noncacheable memory, whether it hits in the L1 cache, whether the cache access generates a write-back to memory, whether the access causes a snoop hit from another device that generates additional activity, and other conditions that affect memory accesses.

The G2 core implements many features to improve throughput, such as pipelining, superscalar instruction dispatch, branch folding, removal of fall-through branches, two-level speculative branch handling, and multiple execution units that operate independently and in parallel.

As an instruction of load/store and floating-point units passes from stage to stage in a pipelined system, the following instruction can follow through the stages as the former instruction vacates them, allowing several instructions to be processed simultaneously. While it may take several cycles for an instruction to pass through all the stages, when the pipeline has been filled, one instruction can complete its work on every clock cycle.

Figure 7-1 represents a generic pipelined execution unit.

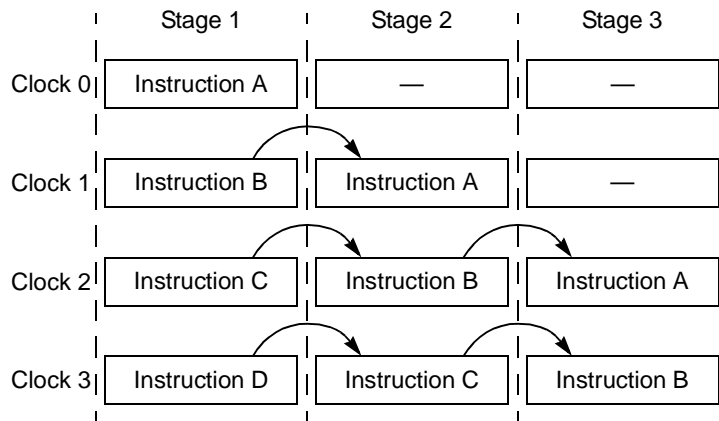


Figure 7-1. Pipelined Execution Unit

The entire path that instructions take through the fetch, decode/dispatch, execute, complete, and write-back stages is considered the G2 core master pipeline, and two of the core execution units (the FPU and LSU) are also multiple-stage pipelines.

The G2 core contains the following execution units that operate independently and in parallel:

- Branch processing unit (BPU)
- 32-bit integer unit (IU)—executes all integer instructions
- 64-bit floating-point unit (FPU)
- Load/store unit (LSU)
- System register unit (SRU)

The G2 core can retire two instructions on every clock cycle. In general, the core processes instructions in four stages—fetch, decode/dispatch, execute, and complete as shown in Figure 7-2. Note that the example of a pipelined execution unit in Figure 7-1 is similar to the three-stage FPU pipeline in Figure 7-2.

The instruction pipeline stages are described as follows:

- The instruction fetch stage includes the clock cycles necessary to request instructions from the memory system and the time the memory system takes to respond to the request. Instruction fetch timing depends on many variables, such as whether the instruction is in the branch target instruction cache, or in the on-chip instruction cache. Instruction fetch timing increases when it is necessary to fetch instructions from system memory. The variables that affect fetch timing include the processor-to-bus clock ratio, the amount of bus traffic, and whether any cache coherency operations are required.

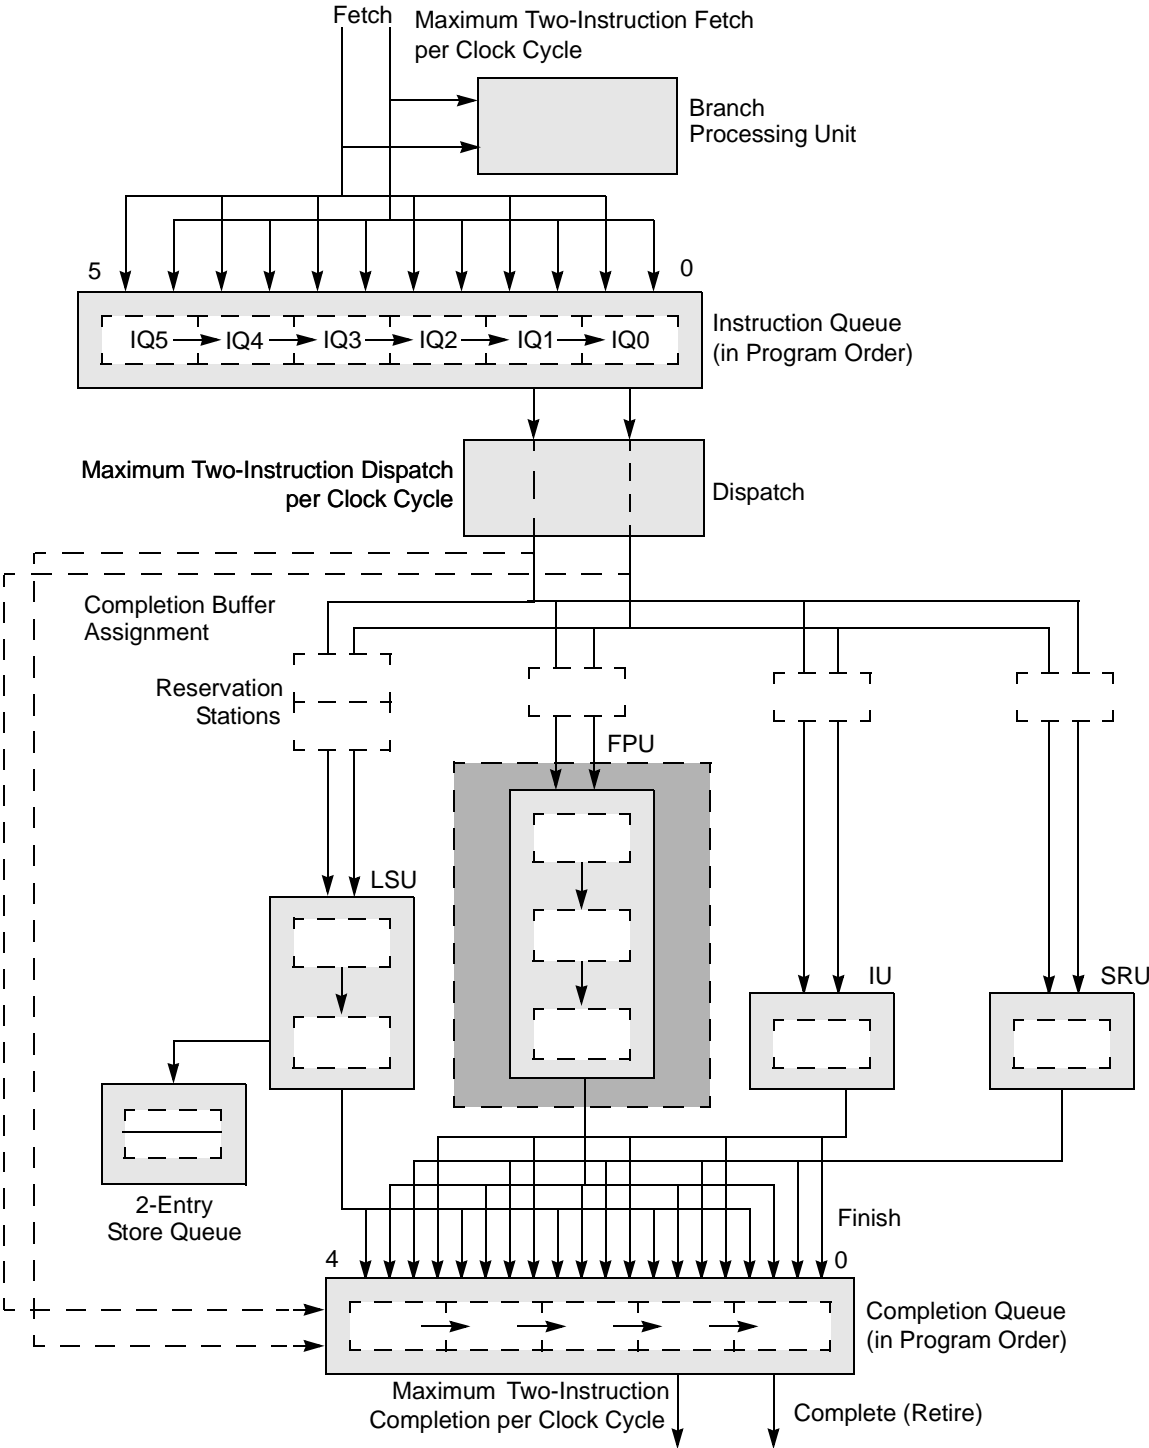


Figure 7-2. Instruction Flow Diagram

Because there are so many variables, unless otherwise specified, the instruction timing examples below assume optimal performance and that the instructions are available in the instruction queue in the same clock cycle that they are requested. The fetch stage ends when the instruction is dispatched.

- The decode/dispatch stage consists of the time it takes to fully decode the instruction and dispatch it from the instruction queue to the appropriate execution unit. Instruction dispatch requires the following:
 - Instructions can be dispatched only from the two lowest instruction queue entries, IQ0 and IQ1.
 - A maximum of two instructions can be dispatched per clock cycle.
 - Only one instruction can be dispatched to each execution unit per clock cycle.
 - There must be a vacancy in the specified execution unit.
 - A rename register must be available for each destination operand specified by the instruction.
 - For an instruction to dispatch, the appropriate execution unit must be available and there must be an open position in the CQ. If no entry is available, the instruction remains in the IQ.
- The execute stage consists of the time between dispatch to the execution unit (or reservation station) and the point at which the instruction vacates the execution unit. Most integer instructions have a one-cycle latency; results of these instructions can be used in the clock cycle after an instruction enters the execution unit. However, integer multiply and divide instructions take multiple clock cycles to complete. The IU can process all integer instructions.

The LSU and FPU are pipelined, as shown in Figure 7-2.

- The complete (complete/write-back) pipeline stage maintains the correct architectural machine state and commits it to the architectural registers at the proper time. If the completion logic detects an instruction containing an exception status, all following instructions are canceled, their execution results in rename registers are discarded, and the correct instruction stream is fetched.

The complete stage ends when the instruction is retired. Two instructions can be retired per cycle. Instructions are retired only from the two lowest CQ entries, CQ0 and CQ1.

The notation conventions used in the instruction timing examples are as follows:

- | | |
|--|---|
| | Fetch—The fetch stage includes the time between when an instruction is requested and when it is brought into the instruction queue. This latency can vary greatly, depending on whether the instruction is in the on-chip cache or system memory (in which case latency can be affected by bus speed and traffic on the system bus, and address translation dispatches). Therefore, in the examples in this chapter, the fetch stage is usually idealized; that is, an instruction is usually shown to be in the fetch stage when it is a valid instruction in the instruction queue. The instruction queue has six entries, IQ0–IQ5. |
|--|---|

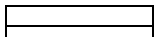



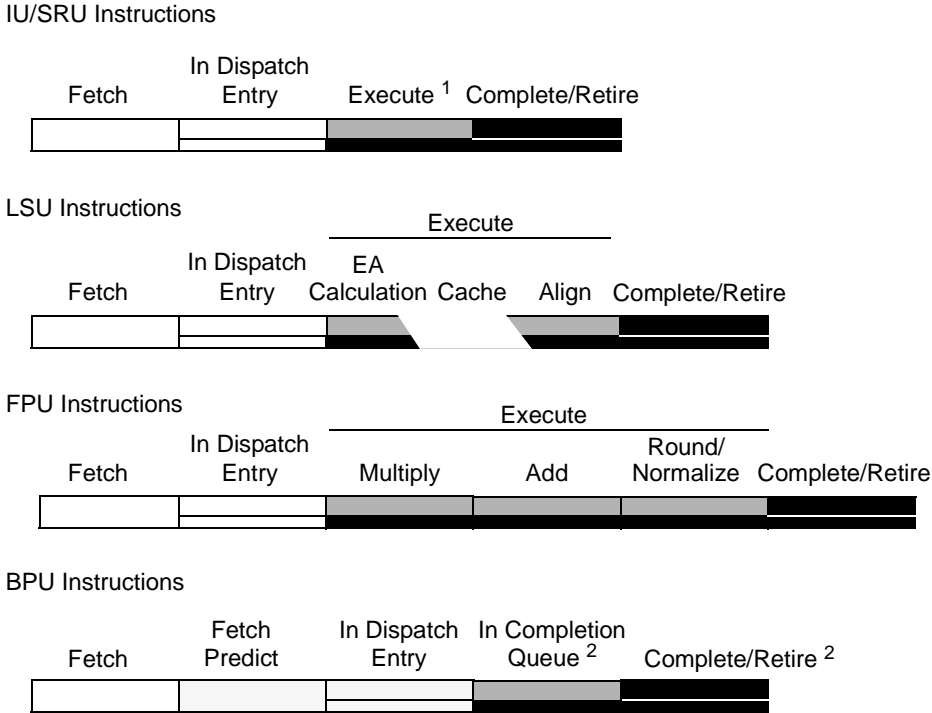
- 
In dispatch entry (IQ0/IQ1)—Instructions can be dispatched from IQ0 and IQ1. Because dispatch is instantaneous, it is perhaps more useful to describe it as an event that marks the point in time between the last cycle in the fetch stage and the first cycle in the execute stage.
- 
Execute—The operations specified by an instruction are being performed by the appropriate execution unit. The black stripe is a reminder that the instruction occupies an entry in the CQ, described in Figure 7-3.
- 
Complete—The instruction is in the CQ. In the final stage, the results of the executed instruction are written back and the instruction is retired. The CQ has five entries, CQ0–CQ4.
- 
In retirement entry—Completed instructions can be retired from CQ0 and CQ1. Like dispatch, retirement is an event that in this case occurs at the end of the final cycle of the complete stage.

Figure 7-3 shows the stages of G2 core execution units.



¹ Several integer instructions, such as multiply and divide instructions, require multiple cycles in the execute stage.
² Only those branch instructions that update the LR or CTR take an entry in the completion queue.

Figure 7-3. G2 Core Processor Pipeline Stages

7.3 Timing Considerations

The G2 core is a superscalar processor; as many as three instructions can be dispatched to the execution units (one branch instruction to the branch processing unit, and two instructions dispatched from the dispatch queue to the other execution units) during each clock cycle. Only one instruction can be dispatched to each execution unit.

Although instructions appear to the programmer to execute in program order, the G2 core improves performance by executing multiple instructions at a time, using hardware to manage dependencies. When an instruction is dispatched, the register file provides the source data to the execution unit. The register files and rename register have sufficient bandwidth to allow dispatch of two instructions per clock under most conditions.

The BPU decodes and executes branches immediately after they are fetched. When a conditional branch cannot be resolved due to a CR data dependency, the branch direction is predicted and execution continues from the predicted path. If the prediction is incorrect, the following steps are taken:

1. The instruction queue is purged and fetching continues from the correct path.
2. Any instructions ahead of the predicted branch in the CQ are allowed to complete.
3. Instructions after the mispredicted branch are purged.
4. Dispatching resumes from the correct path.

After an execution unit executes an instruction, it places resulting data into the appropriate GPR or FPR rename register. The results are then stored into the correct GPR or FPR during the write-back stage. If a subsequent instruction needs the result as a source operand, it is made available simultaneously to the appropriate execution unit, which allows a data-dependent instruction to be decoded and dispatched without waiting to read the data from the register file. Branch instructions that update either the LR or CTR write back their results in a similar fashion.

The following section describes this process in greater detail.

7.3.1 General Instruction Flow

As many as two instructions can be fetched into the instruction queue (IQ) in a single clock cycle. Instructions enter the IQ and are dispatched to the various execution units from the dispatch queue. The IQ is a six-entry queue, which together with the CQ is the backbone of the master pipeline for the microprocessor. The G2 core tries to keep the IQ full at all times.

The number of instructions requested in a clock cycle is determined by the number of vacant spaces in the IQ during the previous clock cycle. This is shown in the examples in this chapter. Although the IQ can accept as many as two new instructions in a single clock cycle and even if there are more than two spaces available on the current clock cycle, if only one IQ entry was vacant on the previous cycle, only one instruction is fetched. Typically,

instructions are fetched from the on-chip instruction cache. If the instruction request hits in the on-chip instruction cache, it can usually present the first two instructions of the new instruction stream in the next clock cycle, giving enough time for the next pair of instructions to be fetched from the cache with no idle cycles. Instructions not in the instruction cache are fetched from system memory.

Branch instructions that do not update the LR or CTR are removed from the instruction stream either by branch folding or removal of fall-through branch instructions, as described in Section 7.4.1.1, “Branch Folding.” Branch instructions that update the LR or CTR are treated as if they require dispatch (even though they are not dispatched to an execution unit in the process). They are assigned a position in the CQ to ensure that the CTR and LR are updated sequentially.

All other instructions are dispatched from IQ0 and IQ1. The dispatch rate depends on the availability of resources such as the execution units, rename registers, and CQ entries, and on the serializing behavior of some instructions. Instructions are dispatched in program order; an instruction in IQ1 can be dispatched at the same time as one in IQ0, but cannot be dispatched ahead of one in IQ0.

Instruction state and all information required for completion is kept in the five-entry, FIFO completion queue. A completion queue entry is allocated for each instruction when it is dispatched to an execute unit; if no entry is available, the dispatch unit stalls. A maximum of two instructions per cycle may be completed and retired from the completion queue, and the flow of instructions can stall when a longer-latency instruction reaches the last position in the completion queue. Store instructions and instructions executed by the FPU and SRU (with the exception of integer add and compare instructions) can only be retired from the last position in the completion queue. Subsequent instructions cannot be completed and retired until that longer-latency instruction completes and retires. Examples of this are shown in Section 7.3.2.2, “Cache Hit,” and Section 7.3.2.3, “Cache Miss.”

The rate of instruction completion is also affected by the ability to write instruction results from the rename registers to the architected registers. The G2 core can perform two write-back operations from the rename registers to the GPRs each clock cycle, but can perform only one write-back per cycle to the CR, FPR, LR, and CTR.

7.3.2 Instruction Fetch Timing

Instruction fetch latency depends on the fetch hits of the on-chip instruction cache. If no hit occurs, a memory transaction is required, in which case fetch latency is affected by bus traffic, bus clock speed, and memory translation. These conditions are discussed in the following sections.

7.3.2.1 Cache Arbitration

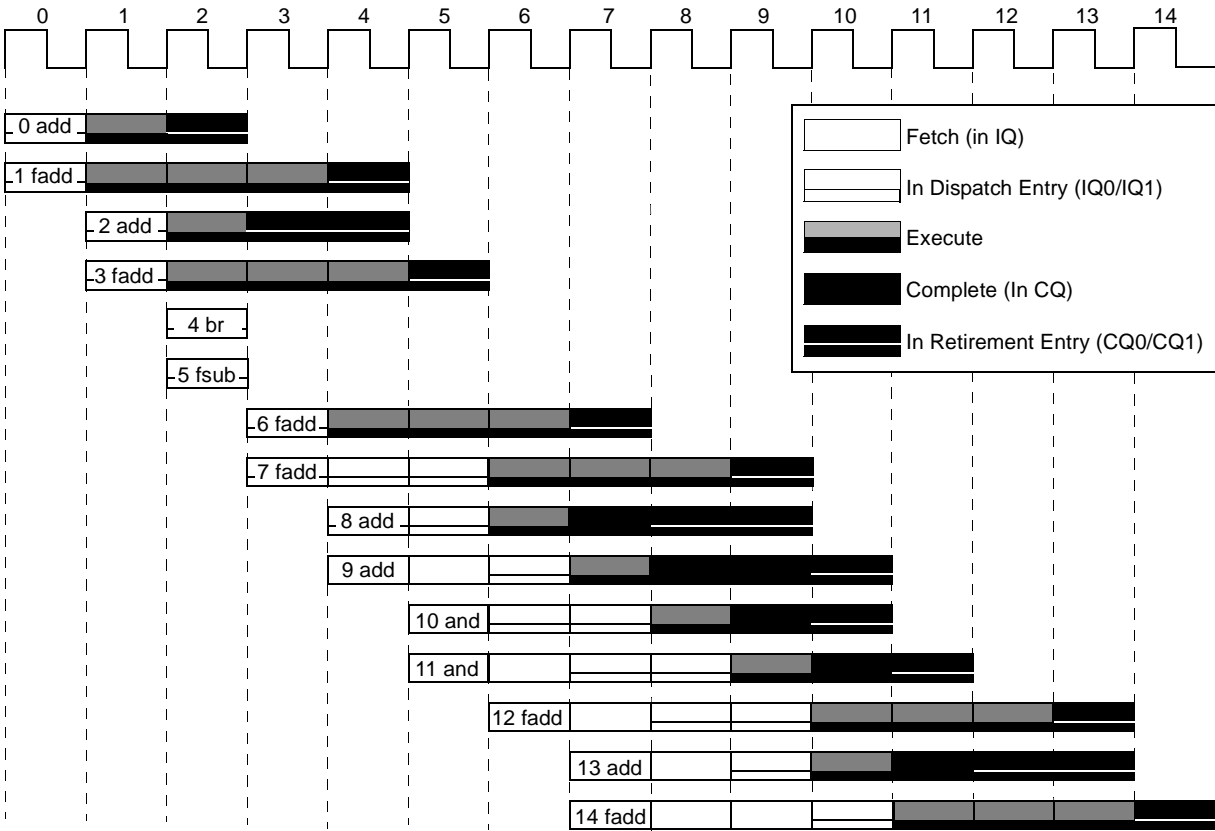
When the fetcher requests instructions from the cache, two things may happen. If the instruction cache is idle and the requested instructions are present, they are provided on the next clock cycle. However, if the instruction cache is busy due to a cache-line-reload operation, instructions cannot be fetched until that operation completes.

7.3.2.2 Cache Hit

An instruction fetch that hits the instruction cache takes only one clock cycle after the request for as many as two instructions to enter the IQ. Note that the cache is not blocked to internal accesses until a cache reload completes (hits under misses). The critical-double-word is written simultaneously to the cache and forwarded to the requesting unit, minimizing stalls due to load delays.

Figure 7-4 shows a simple example of instruction fetching that hits in the on-chip cache. This example uses a series of integer **add**, **and**, and double-precision floating-point add instructions to show how the number of instructions to be fetched is determined, how program order is maintained by the IQ and CQ, how instructions are dispatched and retired in pairs (maximum), and how the FPU pipeline functions. The following instruction sequence is examined:

```
0  add
1  fadd
2  add
3  fadd
4  br 6
5  fsub
6  fadd
7  fadd
8  add
9  add
10 and
11 and
12 fadd
13 add
14 fadd
15 .
16 .
17 .
```



Instruction Queue

						11		14						
						10	12	13	14					
					9	9	11	12	13	14				
1	3	5	7	8	8	10	11	12	13					
0	2	4	6	7	7	9	10	11	12	14				

Completion Queue

									11	13				
		3		6				9	10	12	14			
		2	3	3			8	8	9	11	13	14	14	
	1	1	2	2	6	7	7	8	8	10	12	13	13	
0	0	0	1	1	3	6	6	7	7	9	11	12	12	14

Figure 7-4. Instruction Timing—Cache Hit

The instruction timing for this example is described cycle-by-cycle as follows:

- 0. In cycle 0, instructions 0 and 1 are fetched from the instruction cache and are placed in the two entries in the instruction queue (IQ0 and IQ1), where they can be dispatched on the next clock cycle.

1. In cycle 1, instructions 0 and 1 are dispatched to the IU and FPU, respectively. Notice that for instructions to be dispatched, they must be assigned positions in the CQ. In this case, because the CQ is empty, instructions 0 and 1 take the two lowest CQ entries (CQ0 and CQ1). Instructions 2 and 3 are fetched from the instruction cache.
2. At least two IQ positions were available in the IQ in cycle 1, so in cycle 2, instructions 4 and 5 are fetched. Instruction 4 is a branch unconditional instruction that resolves immediately as taken. Because the branch is taken and does not update CTR or LR, it can be folded from the IQ. Instruction 0 completes, writes back its results, and vacates the CQ by the end of the clock cycle. Instruction 1 enters the second FPU execute stage, instruction 2 enters the single-stage IU, and instruction 3 is dispatched into the first FPU stage.
3. In cycle 3, target instructions 6 and 7 are fetched, replacing the folded **br** instruction 4 and instruction 5. Instruction 1 enters the last FPU execute stage, instruction 2 has executed but must remain in the CQ until instruction 1 completes. Note that it can make its results available to subsequent instructions, but cannot be removed from the CQ. Instruction 3 passes into the last FPU execute stage. Note that all three FPU stages are full. To allow for the potential need for denormalization, the dispatch logic prevents instruction 7 (**fadd**) from being dispatched in the next clock cycle.
4. In cycle 4, target instructions (8 and 9) are fetched. Instruction 1 completes in cycle 4, allowing instruction 2, which had finished executing in the previous clock cycle, to be removed from the CQ. Instruction 6 replaces instruction 3 in the first stage of the FPU. Also, as will be shown in cycle 5, a single-cycle stall occurs when the FPU pipeline is full.
5. In cycle 5, instruction 3 completes, instruction 6 continues through the FPU pipeline, and although the first stage of the FPU pipeline is free, instruction 7 cannot be dispatched because of the potential need for one of the previous floating-point instructions to require denormalization. Because instruction 7 cannot be dispatched neither can instruction 8. This dispatch stall causes the instruction queue to become full when instructions 10 and 11 are fetched.
6. In cycle 6, instruction 12 is fetched. Instruction 7 is dispatched to the first FPU stage, so instruction 8 can also be dispatched to the IU. Instructions 9 and 10 move to IQ0 and IQ1, but because instructions 9, 10, and 11 are integer instructions, only one instruction is dispatched in each of the next two clock cycles. Note that moving instruction 12 (**fadd**) up further in the program flow would improve dispatch throughput.
7. In cycle 7, instruction 6 completes, instruction 7 is in the second FPU execute stage, and although instruction 8 has executed, it must wait for instruction 7 to complete. Instruction 9 dispatches to the IU. Instructions 10 and 11 move down in the IQ. Fetching resumes with instructions 13 and 14.

8. In cycle 8, instruction 7 is in the third FPU execute stage. Instructions 8 and 9 have executed and they remain in the CQ until instruction 7 completes. Instruction 10 is dispatched to the IU.
9. In cycle 9, instruction 7 completes, allowing instruction 8 to complete. Because the CQ is full, instructions 12 and 13 cannot be dispatched.
10. In cycle 10, instructions 9 and 10 complete. Instruction 11 has executed but cannot exit the CQ from CQ2. Instructions 12 and 13 are dispatched to the FPU and IU, respectively. Instruction 14 drops into IQ0.
11. In cycle 11, instruction 11 completes and instruction 12 is in the second FPU execute stage. Instruction 13 has executed but must remain in the CQ until instruction 12 completes. Instruction 14 enters the first FPU execute stage.

7.3.2.3 Cache Miss

Figure 7-5 shows an instruction fetch that misses the on-chip cache and shows how that fetch affects the instruction dispatch. Note that a processor/bus clock ratio of 1:2 is used. The same instruction sequence is used as in Section 7.3.2.2, “Cache Hit.”

A cache miss extends the latency of the fetch stage, so in this example, the fetch stage represents not only the time the instruction spends in the IQ but also the time required for the instruction to be loaded from system memory, beginning in clock cycle 3.

During clock cycle 2, the target instruction for the **br** instruction is not in the instruction cache; therefore, a memory access must occur. During clock cycle 5, the address of the block of instructions is sent to the system bus. During clock cycle 9, two instructions (64 bits) are returned from memory on the first beat and are forwarded both to the cache and instruction fetcher.

7.3.3 Instruction Dispatch and Completion Considerations

Several factors affect the ability of the G2 core to dispatch instructions at a peak rate of two per cycle—the availability of the execution unit, destination rename registers, and completion queue, as well as the handling of completion-serialized instructions. Several of these limiting factors are illustrated in the previous instruction timing examples.

To reduce dispatch unit stalls due to instruction data dependencies, the G2 core provides a single-entry reservation station for the FPU, SRU, and each IU, and a two-entry reservation station for the LSU. If a data dependency keeps an instruction from starting execution, that instruction is dispatched to the reservation station associated with its execution unit (and the rename registers are assigned), thereby freeing the positions in the instruction queue so instructions can be dispatched to other execution units. Execution begins during the same clock cycle that the rename buffer is updated with the data the instruction is dependent on.

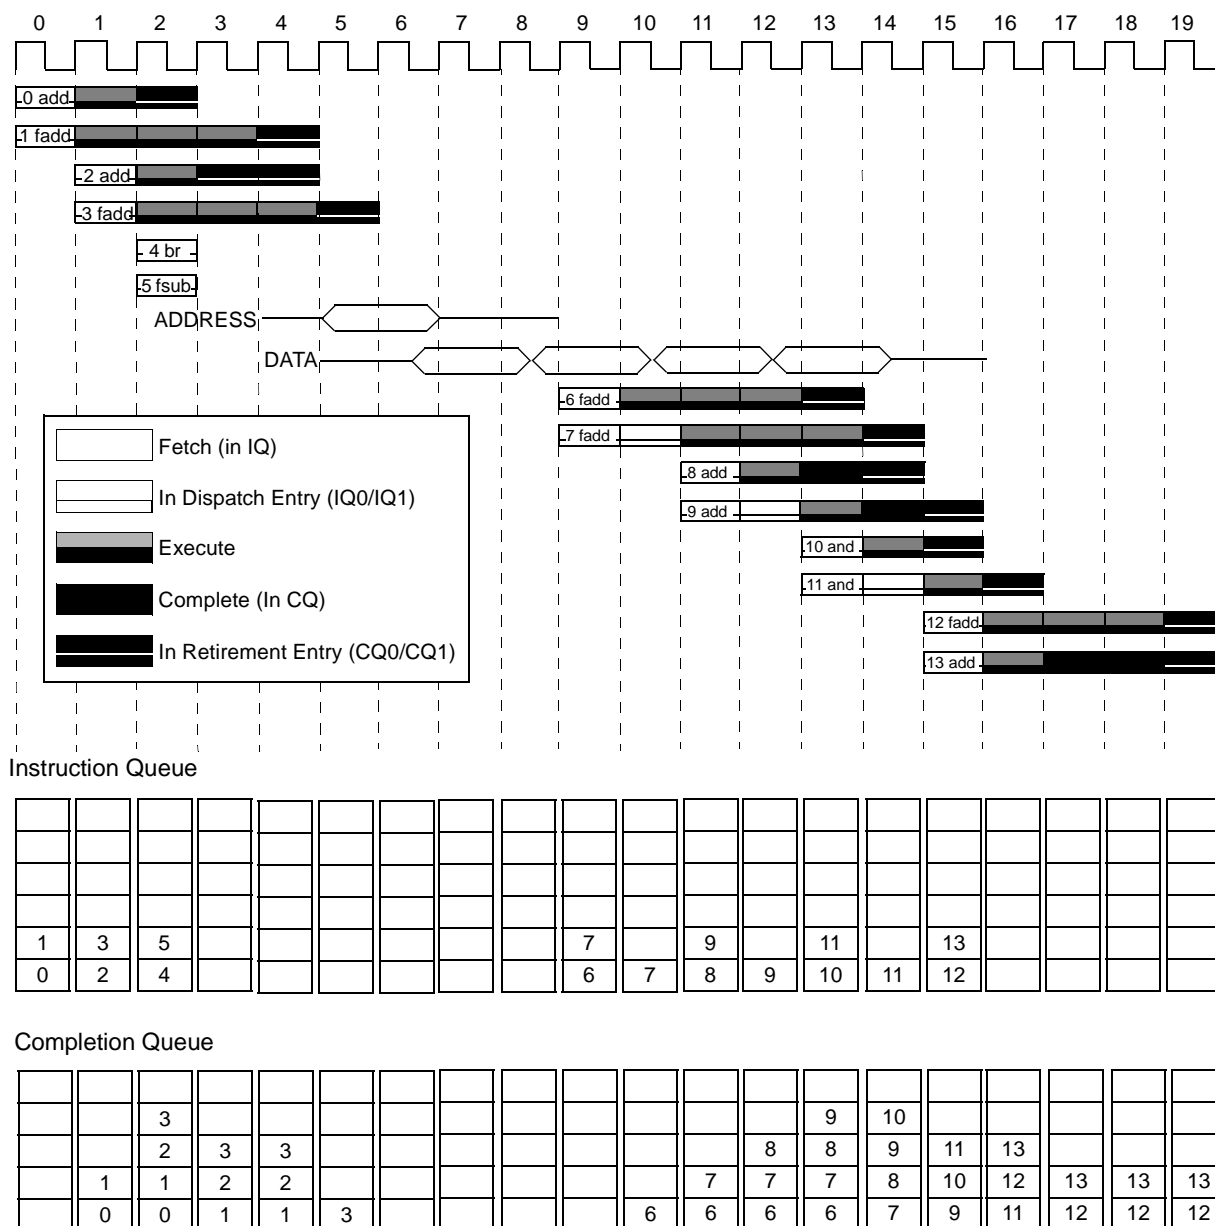


Figure 7-5. Instruction Timing—Cache Miss

If both instructions in IQ0 and IQ1 require the same execution unit, the instruction in IQ1 cannot be dispatched until the first instruction proceeds through the pipeline and provides the subsequent instruction with a vacancy in the requested execution unit.

The completion unit maintains program order after instructions are dispatched, guaranteeing in-order completion and a precise exception model. Completing an instruction implies committing execution results to the architected destination registers. In-order completion ensures the correct architectural state when the core must recover from a mispredicted branch or an exception.

The G2 core can execute instructions out-of-order, but in-order completion by the completion unit ensures a precise exception mechanism. Program-related exceptions are signaled when the instruction causing the exception reaches the last position in the completion queue. Prior instructions are allowed to complete before the exception is taken.

7.3.3.1 Rename Register Operation

To avoid contention for a given register file location, the G2 core provides rename registers for holding instruction results before the completion commits them to the architected register. There are five GPR rename registers, four FPR rename registers, and one each for the CR, LR, and CTR.

When an instruction dispatches to its execution unit, any required rename registers are allocated for the results of that instruction. If an instruction is dispatched to the reservation station associated with an execution unit due to a data dependency, the dispatcher also provides a tag to the execution unit identifying the rename register that forwards the required data at completion. When the source data reaches the rename register, execution can begin.

Instruction results are transferred from rename registers to architected registers when an instruction is retired from the CQ after any associated exceptions are handled and any predicted branch conditions preceding it in the CQ are resolved. If a branch prediction is incorrect, the instructions following the branch are flushed from the CQ and any results of those instructions are flushed from the rename registers.

7.3.3.2 Instruction Serialization

Although the G2 core can dispatch and complete two instructions per cycle, serializing instructions can be used to limit dispatch and completion to one instruction per cycle. Serialization falls into three categories—completion, dispatch, and refetch serialization, which are described as follows:

- Completion serialized instructions are held in the execution unit until all prior instructions in the completion unit have been retired. Completion serialization is used for instructions that access or modify a resource for which no rename register exists. Results from these instructions are not available or forwarded for subsequent instructions until the serializing instruction is retired. Instructions that are completion serialized are as follows:
 - Instructions (with the exception of integer add and compare instructions) executed by the system register unit (SRU)
 - Floating-point instructions that access or modify the FPSCR or CR (**mtfsb1**, **mcrfs**, **mtfsfi**, **mffs**, and **mtfsf**).
 - Instructions that manage caches and TLBs

- Instructions that directly access the GPRs (load and store multiple word and load and store string instructions)
- Instructions defined by the architecture to have synchronizing behavior
- Dispatch serialized inhibit the dispatching of subsequent instructions until the serializing instruction is retired. Dispatch serialization is used for instructions that access renamed resources used by the dispatcher, and for instructions requiring refetch serialization, including the following:
 - The load multiple instructions, **lmw**, **lswi**, and **lswx**.
 - The **mtspr**(XER) and **mcrrx** instructions
 - The synchronizing instructions, **sync**, **isync**, **mtmsr**, **rfi**, **rftci** (for the G2_LE core) and **sc**.
- Refetch serialized instructions inhibit dispatching of subsequent instructions and force the refetching of subsequent instructions after the serializing instructions are retired. The context synchronizing instruction, **isync**, is refetch serializing.

7.3.3.3 Execution Unit Considerations

As previously noted, the G2 core can dispatch and retire two instructions per clock cycle. The peak dispatch rate is affected by the availability of execution units on each clock cycle.

For an instruction to be dispatched, the required execution unit must be available. The dispatcher monitors the availability of all execution units and suspends instruction dispatch if the required execution unit is unavailable. An execution unit may not be available if it can accept and execute only one instruction per cycle or if an execution unit's pipeline becomes full, which may occur if instruction execution takes more clock cycles than the number of pipeline stages in the unit and additional instructions are dispatched to that unit to fill the remaining pipeline stages.

7.4 Execution Unit Timings

The following sections describe instruction timing considerations for each execution unit.

7.4.1 Branch Processing Unit Execution Timing

Flow control operations (conditional branches, unconditional branches, and traps) are typically expensive to execute in most machines because they disrupt normal flow in the instruction stream. When a change in program flow occurs, the IQ must be reloaded with the target instruction stream. During this time the execution units will be idle. However, previously dispatched instructions will continue to execute while the new instruction stream makes its way into the IQ.

Performance features such as branch folding and static branch prediction help minimize penalties associated with flow control operations. The timing for branch instruction execution is determined by many factors including the following:

- Whether the branch requires prediction
- Whether the branch is predicted as taken or not taken
- Whether the branch is taken
- Whether the target instruction stream is in the on-chip cache
- Whether the prediction is correct

7.4.1.1 Branch Folding

When a branch instruction is encountered by the fetcher, the BPU immediately tries to pull that instruction out of the instruction stream and resolve it. When the BPU removes the branch instruction from the stream, the subsequent instruction is shifted down to take the place of the removed branch instruction. This technique is called branch folding. Often, it eliminates the penalties of flow control instructions because instruction execution proceeds as though the branch were never there.

If the folded branch instruction changes program flow (the branch is said to be taken), the BPU immediately requests the instructions at the new target from the on-chip cache. In most cases, the new instructions arrive in the IQ before any bubbles are introduced into the execution units. If the folded branch does not change program flow (the branch is not taken), the branch instruction is already removed and execution continues as if there were never a branch in the original sequence.

When a conditional branch cannot be resolved due to a CR data dependency, the branch is executed by means of static branch prediction and instruction fetching proceeds down the predicted path. If the prediction is incorrect when the branch is resolved, the IQ and all subsequently executed instructions are purged, instructions executed before the predicted branch are allowed to complete, and instruction fetching resumes down the correct path.

There are several situations where instruction sequences create dependencies that prevent a branch instruction from being resolved immediately, thereby causing execution of the subsequent instruction stream based on the predicted outcome of the branch instruction. The instruction sequences, and the resulting action of the branch instruction is described as follows:

- An **mtspr**(LR) followed by a **bclr**—Fetching is stopped and the branch waits for the **mtspr** to execute.
- An **mtspr**(CTR) followed by a **bcctr**—Fetching is stopped and the branch waits for the **mtspr** to execute.
- An **mtspr**(CTR) followed by a **bc**(CTR)—Fetching is stopped and the branch waits for the **mtspr** to execute. (Note: Branch conditions can be a function of the CTR and

CR; if the CTR condition is sufficient to resolve the branch, then a CR-dependency is ignored.)

- A **bc**(CTR) followed by another **bc**(CTR)—Fetching is stopped, and the second branch waits for the first to be completed.
- A **bc**(CTR) followed by a **bcctr**—Fetching is stopped, and the **bcctr** waits for the first branch to be completed.
- A **branch**(LK = 1) followed by a **branch**(LK = 1)—Fetching is stopped, and the second branch waits for the first branch to be completed. (Note: a **bl** instruction does not have to wait for a **branch**(LK = 1) to complete.)
- A **bc**(based-on-CR) waiting for resolution due to a CR-dependency followed by a **bc**(based-on-CR)—Fetching is stopped and the second branch waits for the first CR-dependency to be resolved.

7.4.1.2 Static Branch Prediction

Static branch prediction allows software (for example, compilers) to give a hint to the machine hardware about the direction the branch is likely to take. When a branch instruction encounters a data dependency, the BPU waits for the required condition code to become available. Rather than stalling instruction dispatch until the source operand is ready, the G2 core predicts the likely path and instructions are fetched and executed along that path. When the branch operand becomes available, the branch is evaluated. If the prediction is correct, program flow continues along that path uninterrupted; otherwise, the processor backs up and program flow resumes along the correct path.

If the target address of the branch (link or count register) is modified by an instruction that appears before the branch instruction, the BPU waits until the target address is available.

The G2 core executes through one level of prediction. The processor may not predict a branch if a prior branch instruction is still unresolved.

The number of instructions that can be executed after branch prediction is limited by the fact that instructions in the predicted stream cannot update the register files or memory until the branch is resolved. That is, instructions may be dispatched and executed, but cannot reach the write-back stage in the completion unit, instead, it stalls in the completion queue. When CQ is full, no more instructions can be dispatched.

In the case of a misprediction, the G2 core is able to redirect the machine state rather effortlessly because the programing model has not been updated. When a branch is found to be mispredicted, all instructions that were dispatched subsequent to the predicted branch instruction are simply flushed from the completion queue, and their results flushed from the rename registers. No architected register state needs to be restored because no architected register state was modified by the instructions following the unresolved predicted branch.

7.4.1.2.1 Predicted Branch Timing Examples

Figure 7-6 shows how both taken and non-taken branches are handled and how the G2 core handles both correct and incorrect predictions. The example shows the timing for the following instruction sequence (note that the first **bc** instruction is correctly taken, whereas the second **bc** is incorrectly predicted):

```

0  add
1  add
2  bc
3  mulhw
4  bc T0
5  fadd
6  and
T0 add
T1 add
T2 add
T3 add
T4 and
T5 or

```

0. During clock cycle 0, instructions 0 and 1 are dispatched in the beginning of clock cycle 1.
1. In clock cycle 1, instructions 2 and 3 are fetched in the IQ. Instruction 2 is a branch instruction that updates the CTR and instruction 3 is a **mulhw** instruction on which instruction 4 depends. Instruction 0 enters the IU. Instruction 1 has a single-cycle stall.
2. In clock cycle 2, instructions 4 (a second **bc** instruction) and 5 are fetched. The second **bc** instruction is predicted as taken. It can be folded, but it cannot be resolved until instruction 3 writes back. Instruction 0 completes at the end of this cycle. Instruction 1 is dispatched to the IU. Instruction 2 takes entry in the CQ.
3. In clock cycle 3, target instruction T0 and T1 are fetched. Instructions 1 and 2 complete, instruction 4 has been folded, and instruction 5 has been flushed from the IQ. Instruction 3 is assigned to CQ2.
4. In clock cycle 4, target instructions T2 and T3 are fetched. IU instructions T0 and T1 have multiple stalls as one execution possible in a clock cycle. Instruction 3 is assigned to CQ0.
5. In clock cycle 5, instruction 3, on which the second branch instruction depended, writes back and the branch prediction is proven incorrect. Even though T0 is in CQ0, where it could be written back, it is not because the prediction was incorrect. All target instructions are flushed from their positions in the pipeline at the end of this clock cycle, as there are many results in the rename registers.

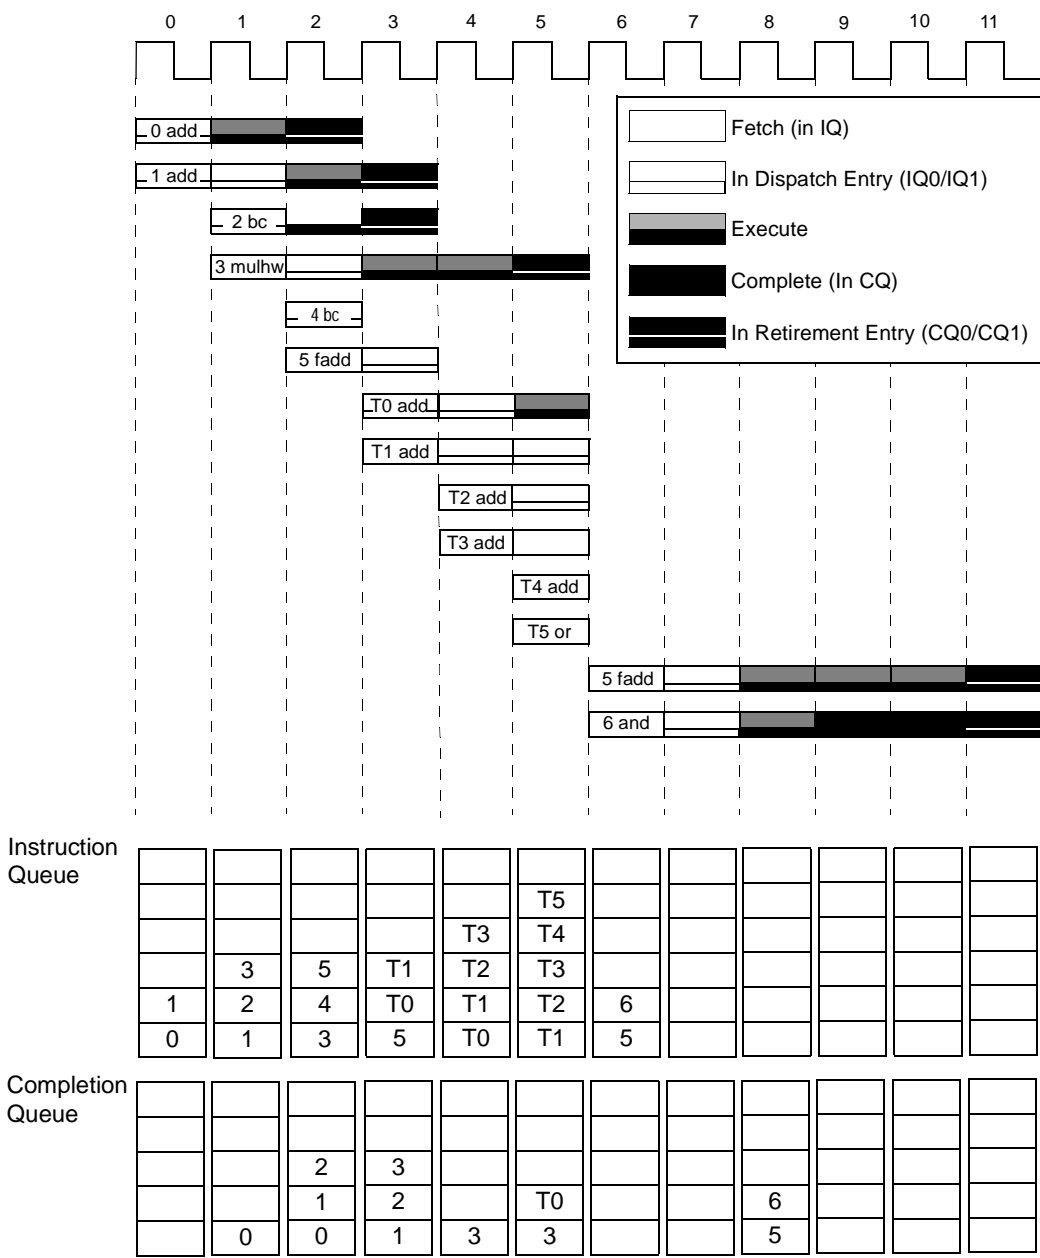


Figure 7-6. Branch Instruction Timing

After one clock cycle required to refetch the original instruction stream, instruction 5, the same instruction that was fetched in clock cycle 2, is brought back into the IQ from the instruction cache, along with one other.

7.4.2 Integer Unit Execution Timing

The integer unit executes all integer and bit-field computational instructions. Many of these instructions execute in a single clock cycle. The integer unit has one execute stage so when

a multiple-cycle integer instruction is executed, no other integer instructions can also begin to execute. See Table 7-4 for integer instruction execution timing.

7.4.3 Floating-Point Unit Execution Timing

The FPU on the G2 core executes all floating-point computational instructions. The LSU performs integer floating-point loads and stores. Execution of most floating-point instructions is pipelined within the FPU, allowing up to three instructions to be executing in the FPU concurrently. While most floating-point instructions execute with three- or four-cycle latency, and one- or two-cycle throughput, three instructions (**fdivs**, **fdiv**, and **fres**) execute with latencies of 18 to 33 cycles. The **fdivs**, **fdiv**, **fres**, **mtfsb0**, **mtfsb1**, **mtfsfi**, **mffs**, and **mtfsf** instructions block the floating-point unit pipeline until they complete execution, and thereby inhibit the dispatch of additional floating-point instructions. With the exception of the **mcrfs** instruction, all floating-point instructions will immediately forward their CR results to the BPU for fast branch resolution without waiting for the instruction to be retired by the completion unit, and the CR updated. See Table 7-5 for floating-point instruction execution timing.

7.4.4 Load/Store Unit Execution Timing

The LSU executes all floating-point and integer loads and stores. It also executes other instructions that address memory. The execution of most load and store instructions is pipelined. The LSU has two pipeline stages; the first is for effective address calculation and MMU translation, and the second is for accessing the physically addressed memory. Load and store instructions have a two-cycle latency and one-cycle throughput.

If operands are misaligned, additional latency may be required either for an alignment exception to be taken or for additional bus accesses. Load instructions that miss in the cache prevent subsequent cache accesses during the cache line refill. See Table 7-6 for load and store instruction execution timing.

7.4.5 System Register Unit Execution Timing

Most SRU instructions access or modify nonrenamed registers, or directly access renamed registers. They generally execute in a serial manner. Results from these instructions are not available or forwarded for use by subsequent instructions until the instruction completes and is retired. The SRU can also execute the integer instructions **addi**, **addis**, **add**, **addo**, **cmpi**, **cmp**, **cmpli**, and **cmpl** without serialization and in parallel with another integer instruction. Refer to Section 7.3.3.2, “Instruction Serialization,” for additional information on serializing instructions and Table 7-2, Table 7-3, and Table 7-4 for SRU instruction execution timing.

7.5 Memory Performance Considerations

Due to the G2 core instruction throughput of three instructions per clock cycle, lack of data bandwidth can become a performance bottleneck. For the G2 core to approach its potential performance levels, it must be able to read and write data quickly and efficiently. If there are many processors in a system environment, one processor may experience long memory latencies while another bus master (for example, a direct-memory access controller) is using the external bus.

To alleviate this possible contention, the G2 core provides three memory update modes—copy-back, write-through, and cache-inhibit. Each page of memory is specified to be in one of these modes. If a page is in copy-back mode, data being stored to that page is written only to the on-chip cache. If a page is in write-through mode, writes to that page update the on-chip cache on hits and always update main memory. If a page is cache-inhibited, data in that page will never be stored in the on-chip cache. All three of these modes of operation have advantages and disadvantages. A decision as to which mode to use depends on the system environment as well as the application.

The following sections describe how performance is impacted by each memory update mode. For details about the operation of the on-chip cache and the memory update modes, see Chapter 4, “Instruction and Data Cache Operation.”

7.5.1 Copy-Back Mode

When data is stored in a location marked as copy back, store operations for cacheable data do not necessarily cause an external bus cycle to update memory. Instead, memory updates only occur on modified line replacements, cache flushes, or when another processor attempts to access a specific address for which there is a corresponding modified cache entry. For this reason, copy-back mode may be preferred when external bus bandwidth is a potential bottleneck—for example, in a multiprocessor environment. Copy-back mode is also well suited for data that is closely coupled to a processor, such as local variables.

If more than one device uses data stored in a page marked as copy back, snooping must be enabled to allow copy-back operations and cache invalidations of modified data. The G2 core implements snooping hardware to prevent other devices from accessing invalid data. When bus snooping is enabled, depending on the device integration, the processor can monitor the transactions of the other devices. For example, if another device accesses a memory location and its memory-coherent (M) bit is set and the G2 core on-chip cache has a modified value for that address, the processor preempts the bus transaction and updates memory with the cache data. If the cache contents associated with the snooped address are unmodified, the G2 core invalidates the cache block. The other device can then attempt an access to the updated address. See Chapter 4, “Instruction and Data Cache Operation.”

Copy-back mode provides complete cache/memory coherency as well as maximizing available external bus bandwidth.

7.5.2 Write-Through Mode

Store operations to memory in write-through mode always update memory as well as the on-chip cache (on cache hits). Write-through mode is used when the data in the cache must always agree with external memory (for example, video memory), when shared (global) data may be used frequently, or when allocation of a cache line on a cache miss is undesirable. Automatic copy back of cached data is not performed if that data is from a memory page marked as write-through mode because valid cache data always agrees with memory.

Stores to memory that are in write-through mode may cause a decrease in performance. Each time a store is performed to memory in write-through mode, the bus is potentially busy for the extra clock cycles required to update memory; therefore, load operations that miss the on-chip cache must wait while the external store operation completes.

7.5.3 Cache-Inhibited Accesses

Data for a page marked cache-inhibited cannot be stored in the on-chip cache.

Areas of the memory map can be cache-inhibited by the operating system. If a cache-inhibited access hits in the on-chip cache, the corresponding cache line is invalidated. If the line is marked modified, it is copied back to memory before being invalidated.

In summary, the copy-back mode allows both load and store operations to use the on-chip cache. The write-through mode allows load operations to use the on-chip cache, but store operations cause a memory access and a cache update if the data is already in the cache. Lastly, the cache-inhibited mode causes memory access for both loads and stores.

7.6 Instruction Scheduling Guidelines

The performance of the G2 core can be improved by avoiding resource conflicts and promoting parallel utilization of execution units through efficient instruction scheduling. Instruction scheduling on the G2 core can be improved by observing the following guidelines:

- Implement good static branch prediction (setting of y bit in BO field).
- When branch prediction is uncertain, or an even probability, predict fall through.
- To reduce mispredictions, separate the instruction that sets CR bits from the branch instruction that evaluates them; separation by more than nine instructions ensures that the CR bits will be immediately available for evaluation.
- When branching conditionally to a location specified by count registers (CTRs) or link registers (LRs), or when branching conditionally based on the value in the count register, separate the **mtspr** instruction that initializes the CTR or LR from the

branch instruction performing the evaluation. Separation of the branch and **mtspr** instruction by more than nine instructions ensures the register values will be immediately available for use by the branch instruction.

- Schedule instructions such that they can dual dispatch.
- Schedule instructions to minimize stalls when an execution unit is busy.
- Avoid using serializing instructions.
- Schedule instructions to avoid dispatch stalls due to renamed resource limitations.
 - Only five instructions can be in execute-complete stage at any one time.
 - Only five GPR destinations can be in execute-complete-deallocate stage at any one time. Note that load with update address instructions use two destination registers.
 - Only four FPR destinations can be in execute-complete-deallocate stage at any one time.

7.6.1 Branch, Dispatch, and Completion Unit Resource Requirements

This section describes the specific resources required to avoid stalls during branch resolution, instruction dispatching, and instruction completion.

7.6.1.1 Branch Resolution Resource Requirements

The following is a list of branch instructions and the resources required to avoid stalling the fetch unit in the course of branch resolution:

- The **bclr** instruction requires LR availability.
- The **bcctr** instruction requires CTR availability.
- Branch and link instructions require shadow LR availability.
- The branch conditional on counter decrement and CR condition requires CTR availability or the CR condition must be false, and the G2 core cannot be executing instructions following an unresolved predicted branch when the branch is encountered by the BPU.
- The branch conditional on CR condition cannot be executed following an unresolved predicted branch instruction.

7.6.1.2 Dispatch Unit Resource Requirements

The following is a list of resources required to avoid stalls in the dispatch unit. Note that the two dispatch buffers, IQ0 and IQ1, are at the bottom of the instruction queue:

- Requirements for dispatching from IQ0 are as follows:
 - Needed execution unit available
 - Needed GPR rename registers available
 - Needed FPR rename registers available
 - Completion queue is not full
 - Instruction is dispatch serialized and completion buffer is empty
 - A dispatch serialized instruction is not currently being executed
- Requirements for dispatching from IQ1 are as follows:
 - Instruction in IQ0 must dispatch
 - Instruction dispatched by IQ0 is not dispatch serialized
 - Needed execution unit is available (after dispatch from IQ0)
 - Needed GPR rename registers are available (after dispatch from IQ0)
 - Needed FPR rename register is available (after dispatch from IQ0)
 - Completion queue is not full (after dispatch from IQ0)
 - Instruction dispatched from IQ1 is not dispatch serialized

7.6.1.3 Completion Unit Resource Requirements

The following is a list of resources required to avoid stalls in the completion unit; note that the two completion buffers are described as CQ0 and CQ1, where CQ0 is the entry at the end of the completion queue:

- Requirements for completing an instruction from CQ0 are as follows:
 - Instruction in CQ0 must be finished
 - Instruction in CQ0 must not follow an unresolved predicted branch
 - Instruction in CQ0 must not cause an exception
- Requirements for completing an instruction from CQ1 are as follows:
 - Instruction in CQ0 must complete in same cycle
 - Instruction in CQ1 must be finished
 - Instruction in CQ1 must not follow an unresolved predicted branch
 - Instruction in CQ1 must not cause an exception
 - Instruction in CQ1 must be an integer or load instruction
 - Number of CR updates from both CQ0 and CQ1 must not exceed one

- Number of GPR updates from both CQ0 and CQ1 must not exceed two
- Number of FPR updates from both CQ0 and CQ1 must not exceed one

7.7 Instruction Latency Summary

Table 7-1 through Table 7-6 list the latencies associated with each instruction executed by the G2 core. Note that the instruction latency tables contain no 64-bit architected instructions. These instructions will trap to an illegal instruction exception handler when encountered. Recall that the term latency is defined as the total time it takes to execute an instruction and make ready the results of that instruction.

Table 7-1 provides the latencies for the branch instructions.

Table 7-1. Branch Instructions

Mnemonic	Primary	Extended	Unit	Latency (in Cycles) ¹
bc[l][a]	16	—	BPU	1
b[l][a]	18	—	BPU	1
bclr[l]	19	016	BPU	1
bcctr[l]	19	528	BPU	1

¹ These operations may be folded for an effective cycle time of 0.

Table 7-2 provides the latencies for the system register instructions.

Table 7-2. System Register Instructions

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
sc	17	--1	SRU	3
rfi, rfc (G2_LE only)	19	050	SRU	3
isync	19	150	SRU	1&
mfmsr	31	083	SRU	1
mtmsr	31	146	SRU	2
mtsr	31	210	SRU	2
mtsrin	31	242	SRU	2
mfspr (not I/DBATs)	31	339	SRU	1
mfspr (DBATs)	31	339	SRU	3&
mfspr (IBATs)	31	339	SRU	3&
mtspr (not IBATs)	31	467	SRU	2 (XER-&)
mtspr (IBATs)	31	467	SRU	2&
mfsr	31	595	SRU	3&
sync	31	598	SRU	1&

Table 7-2. System Register Instructions (continued)

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
mfsrin	31	659	SRU	3&
eieio	31	854	SRU	1
mftb	31	371	SRU	1
mttb	31	467	SRU	1

Note: Cycle times marked with & require a variable number of cycles due to serialization.

Table 7-3 provides the latencies for the condition register logical instructions.

Table 7-3. Condition Register Logical Instructions

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
mcrf	19	000	SRU	1
crnor	19	033	SRU	1
crandc	19	129	SRU	1
crxor	19	193	SRU	1
crnand	19	225	SRU	1
crand	19	257	SRU	1
creqv	19	289	SRU	1
crorc	19	417	SRU	1
cror	19	449	SRU	1
mfcrr	31	019	SRU	1
mtcrf	31	144	SRU	1
mcrxr	31	512	SRU	1&

Note: Cycle times marked with & require a variable number of cycles due to serialization.

Table 7-4 provides the latencies for the integer instructions.

Table 7-4. Integer Instructions

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
twi	03	—	Integer	2
mulli	07	—	Integer	2,3
subfic	08	—	Integer	1
cmpli	10	—	Integer & SRU	1 [^]
cmpi	11	—	Integer & SRU	1 [^]
addic	12	—	Integer	1

Table 7-4. Integer Instructions (continued)

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
addic.	13	—	Integer	1
addi	14	—	Integer & SRU	1
addis	15	—	Integer & SRU	1
rlwimi[.]	20	—	Integer	1
rlwinm[.]	21	—	Integer	1
rlwnm[.]	23	—	Integer	1
ori	24	—	Integer	1
oris	25	—	Integer	1
xori	26	—	Integer	1
xoris	27	—	Integer	1
andi.	28	—	Integer	1
andis.	29	—	Integer	1
cmp	31	000	Integer & SRU	1 [^]
tw	31	004	Integer	2
subfc[o][.]	31	008	Integer	1
addc[o][.]	31	010	Integer	1
mulhwu[.]	31	011	Integer	2,3,4,5,6
slw[.]	31	024	Integer	1
cntlzw[.]	31	026	Integer	1
and[.]	31	028	Integer	1
cmpl	31	032	Integer & SRU	1 [^]
subf[.]	31	040	Integer	1
andc[.]	31	060	Integer	1
mulhw[.]	31	075	Integer	2,3,4,5
neg[o][.]	31	104	Integer	1
nor[.]	31	124	Integer	1
subfe[o][.]	31	136	Integer	1
adde[o][.]	31	138	Integer	1
subfze[o][.]	31	200	Integer	1
addze[o][.]	31	202	Integer	1
subfme[o][.]	31	232	Integer	1
addme[o][.]	31	234	Integer	1
mull[o][.]	31	235	Integer	2,3,4,5
add[o][.]	31	266	Integer & SRU ¹	1
eqv[.]	31	284	Integer	1

Table 7-4. Integer Instructions (continued)

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
xor [.]	31	316	Integer	1
orc [.]	31	412	Integer	1
or [.]	31	444	Integer	1
divwu [o][.]	31	459	Integer	20
nand [.]	31	476	Integer	1
divw [o][.]	31	491	Integer	20
srw [.]	31	536	Integer	1
sraw [.]	31	792	Integer	1
srawi [.]	31	824	Integer	1
extsh [.]	31	922	Integer	1
extsb [.]	31	954	Integer	1

Note: ^ indicates that the cycle time immediately forwards their CR results to the BPU for fast branch resolution.

¹ The SRU can only execute the **add** and **add[o]** instructions.

Table 7-5 provides the latencies for the floating-point instructions.

Table 7-5. Floating-Point Instructions

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
fdivs [.]	59	018	FPU	18 [^]
fsubs [.]	59	020	FPU	1-1-1 [^]
fadds [.]	59	021	FPU	1-1-1 [^]
fres [.]	59	024	FPU	18 [^]
fmls [.]	59	025	FPU	1-1-1 [^]
fmsubs [.]	59	028	FPU	1-1-1 [^]
fmadds [.]	59	029	FPU	1-1-1 [^]
fnmsubs [.]	59	030	FPU	1-1-1 [^]
fnmadds [.]	59	031	FPU	1-1-1 [^]
fcmpu	63	000	FPU	1-1-1 [^]
frsp [.]	63	012	FPU	1-1-1 [^]
fctiw [.]	63	014	FPU	1-1-1 [^]
fctiwz [.]	63	015	FPU	1-1-1 [^]
fdiv [.]	63	018	FPU	33 [^]
fsub [.]	63	020	FPU	1-1-1 [^]
fadd [.]	63	021	FPU	1-1-1 [^]

Table 7-5. Floating-Point Instructions (continued)

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
fsel [.]	63	023	FPU	1-1-1^
fmul [.]	63	025	FPU	2-1-1^
frsqste [.]	63	026	FPU	1-1-1^
fmsub [.]	63	028	FPU	2-1-1^
fmadd [.]	63	029	FPU	2-1-1^
fnmsub [.]	63	030	FPU	2-1-1^
fnmadd [.]	63	031	FPU	2-1-1^
fcmpo	63	032	FPU	1-1-1^
mtfsb1 [.]	63	038	FPU	1-1-1&^
fneg [.]	63	040	FPU	1-1-1^
mcrfs	63	064	FPU	1-1-1&
mtfsb0 [.]	63	070	FPU	1-1-1&^
fmr [.]	63	072	FPU	1-1-1^
mtfsfi [.]	63	134	FPU	1-1-1&^
fnabs [.]	63	136	FPU	1-1-1^
fabs [.]	63	264	FPU	1-1-1^
mffs [.]	63	583	FPU	1-1-1&^
mtfsf [.]	63	711	FPU	1-1-1&^

Notes: Cycle times marked with & require a variable number of cycles due to completion serialization.
 Cycle times marked with ^ immediately forward their CR results to the BPU for fast branch resolution.
 Cycle times marked with a - specify the number of clock cycles in each pipeline stage.
 Instructions with a single entry in the cycles column are not pipelined.

Table 7-6 provides latencies for the load and store instructions.

Table 7-6. Load and Store Instructions

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
lwarx	31	020	LSU	2:1
lwzx	31	023	LSU	2:1
dcbst	31	054	LSU	2/5&
lwzux	31	055	LSU	2:1
dcbf	31	086	LSU	2/5&
lbzx	31	087	LSU	2:1
lbzux	31	119	LSU	2:1
stwcx.	31	150	LSU	8

Table 7-6. Load and Store Instructions (continued)

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
stwx	31	151	LSU	2:1
stwux	31	183	LSU	2:1
stbx	31	215	LSU	2:1
dcbtst	31	246	LSU	2
stbux	31	247	LSU	2:1
dcbt	31	278	LSU	2
lhzx	31	279	LSU	2:1
tlbie	31	306	LSU	3&
eciwx	31	310	LSU	2:1
lhzux	31	311	LSU	2:1
lhax	31	343	LSU	2:1
lhaux	31	375	LSU	2:1
sthx	31	407	LSU	2:1
ecowx	31	438	LSU	2:1
sthux	31	439	LSU	2:1
dcbi ¹	31	470	LSU	2&
lswx	31	533	LSU	2 + n&
lwbrx	31	534	LSU	2:1
lfsx	31	535	LSU	2:1
tlbsync	31	566	LSU	2&
lfsux	31	567	LSU	2:1
lswi	31	597	LSU	2 + n&
lfdx	31	599	LSU	2:1
lfdux	31	631	LSU	2:1
stswx	31	661	LSU	1 + n&
stwbrx	31	662	LSU	2:1
stfsx	31	663	LSU	2:1
stfsux	31	695	LSU	2:1
stswi	31	725	LSU	1 + n&
stfdx	31	727	LSU	2:1
stfdux	31	759	LSU	2:1
lhbrx	31	790	LSU	2:1
sthbrx	31	918	LSU	2:1
tlbld	31	978	LSU	2&
icbi	31	982	LSU	3&

Table 7-6. Load and Store Instructions (continued)

Mnemonic	Primary	Extended	Unit	Latency (in Cycles)
stfiwx	31	983	LSU	2:1
tlbli	31	1010	LSU	3&
dcbz	31	1014	LSU	10&
lwz	32	—	LSU	2:1
lwzu	33	—	LSU	2:1
lbz	34	—	LSU	2:1
lbzu	35	—	LSU	2:1
stw	36	—	LSU	2:1
stwu	37	—	LSU	2:1
stb	38	—	LSU	2:1
stbu	39	—	LSU	2:1
lhz	40	—	LSU	2:1
lhzu	41	—	LSU	2:1
lha	42	—	LSU	2:1
lhau	43	—	LSU	2:1
sth	44	—	LSU	2:1
sthu	45	—	LSU	2:1
lmw	46	—	LSU	2 + n&
stmw	47	—	LSU	1 + n&
lfs	48	—	LSU	2:1
lfsu	49	—	LSU	2:1
lfd	50	—	LSU	2:1
lfdu	51	—	LSU	2:1
stfs	52	—	LSU	2:1
stfsu	53	—	LSU	2:1
stfd	54	—	LSU	2:1
stfdu	55	—	LSU	2:1

Notes: Cycle times marked with & require a variable number of cycles due to serialization.

Cycle times marked with a / specify hit and miss times for cache management instructions that require conditional bus activity.

Cycle times marked with a : specify cycles of total latency and throughput.

Load and store multiple and string instruction cycles are shown as a fixed number of cycles plus a variable number of cycles where *n* is the number of words accessed by the instruction.

¹ The **dcbi** instruction should never be used on the G2 core.

Chapter 8

Signal Descriptions

This chapter describes the signals of the G2 core that are candidates for being driven as external device signals. It contains a concise description of the individual signals, showing behavior when the signal is asserted and negated, which signals are input/output pairs with output enable signals, and which signals also have high-impedance control signals.

NOTE

A bar over a signal name indicates that the signal is active-low—for example, core_artry (address retry) and core_ts (transfer start). Active-low signals are referred to as asserted (active) when they are low and negated when they are high. Signals that are not active-low, such as core_ap[0:3] (address bus parity signals) and core_tt[0:4] (transfer type signals) are referred to as asserted when they are high and negated when they are low.

8.1 Signal Groupings

The G2 core 60x bus interface protocol signals are grouped as follows:

- Address arbitration signals—The G2 core uses these signals to arbitrate for 60x address bus mastership.
- Address transfer start signals—These signals indicate that a bus master has begun a transaction on the address bus of the 60x bus.
- Address transfer signals—These signals, consisting of the address bus, address parity, and address parity error signals, are used to transfer the address and to ensure the integrity of the transfer.
- Transfer attribute signals—These signals provide information about the type of transfer, such as the transfer size and whether the transaction is bursted, write-through, or cache-inhibited.
- Address transfer termination signals—These signals are used to acknowledge the end of the address phase of the transaction. They also indicate whether a condition exists that requires the address phase to be repeated.

- Data arbitration signals—The G2 core uses these signals to arbitrate for data bus mastership of the 60x data bus.
- Data transfer signals—These signals, consisting of the data bus, data parity, and data parity error signals, are used to transfer the data and to ensure the integrity of the transfer.
- Data transfer termination signals—Data termination signals are required after each data beat in a data transfer. In a single-beat transaction, the data termination signals also indicate the end of the tenure. In burst accesses, the data termination signals apply to individual beats and indicate the end of the tenure only after the final data beat. They also indicate whether a condition exists that requires the data phase to be repeated.
- Output enable signals—These output signals indicate that the corresponding outputs of the G2 core are driving, provided the corresponding high-impedance control signal is also asserted.
- High-impedance control signals—These input signals (static) enable the operation of the output-enable signals.
- Input enable signals—When these input signals are asserted, it indicates that they expect to receive valid data into the core.

In addition, there are many other signals on the G2 core that control and affect other aspects of the device, aside from the bus protocol. They are as follows:

- System status signals—These signals include the external interrupt signal, the critical interrupt signal (G2_LE only), checkstop signals, and both soft- and hard-reset signals. These signals are used to interrupt and, under various conditions, to reset the core.
- Reset configuration signals—These signals are sampled while `core_hreset` is asserted and they control certain modes of operation.
- JTAG/COP interface signals—The JTAG (IEEE 1149.1) interface and common on-chip processor (COP) unit provides a serial interface to the system for performing monitoring and boundary tests.
- Processor status—These signals include the memory reservation signal, machine quiesce control signals, time base enable signal, and `core_tlbisync` signal.
- Debug control—These signals are implemented to control debug features of the PowerPC architecture with respect to the G2 and G2_LE cores.
- Clock signals—These signals provide for system clock input and frequency control.
- Test interface signals—Signals like address matching, combinational matching, and watchpoint are used in the G2_LE for production testing.

8.2 Signal Configurations

This section provides various mappings of the G2 core signals.

8.2.1 Functional Groupings

Figure 8-1 shows how the G2 core signals are grouped by function.

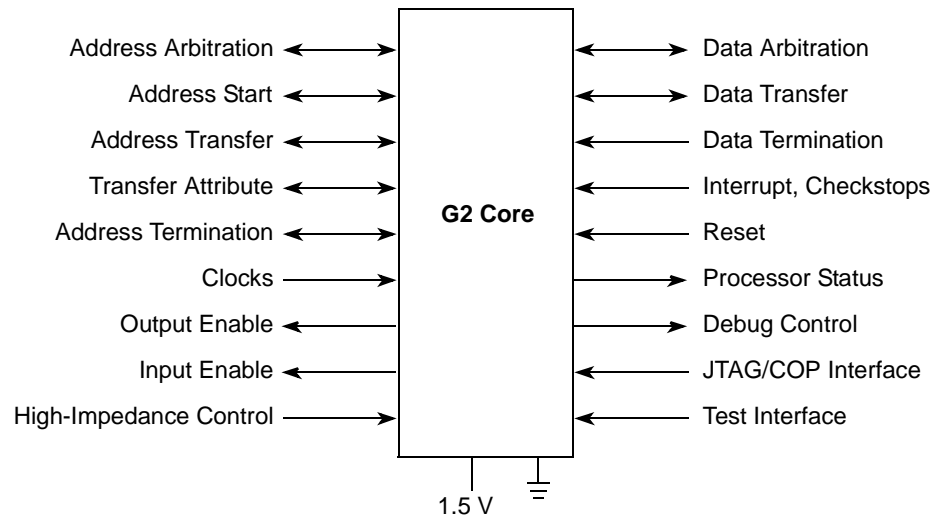


Figure 8-1. Functional Signal Groups

8.2.2 Input/Output Enable and High-Impedance Control Signals

The G2 core splits the bidirectional signals of the 60x bus into separate input and output signal pairs. In addition, high-impedance signals are included, allowing these input and output signals to be reconnected into a bidirectional signal elsewhere on the device. Table 8-1 maps the high-impedance control signals and the output-enable and input-enable indicators to their corresponding 60x bus signals. Each signal in the left column applies to all of the signals in the right column in the same row.

Table 8-1. Input/Output Enable and High-Impedance Signal Mappings

Input/Output Enable and High-Impedance Control Signals	Affected Signals
core_a_oe core_a_tre core_ap_ien	core_a_out[0:31] core_ap_out[0:3] core_ci core_cse[0:1] core_gbl_out core_tbst_out core_tc[0:1] core_tsiz[0:2] core_ts_out core_tt_out[0:4] core_wt
core_d_oe core_d_tre	core_dh_out[0:31] core_dl_out[0:31] core_dp_out[0:7]
core_dh_ien	core_dh_out[0:31]
core_dl_ien	core_dl_out[0:31]
core_dp_ien	core_dp_out[0:7]
core_abb_oe core_abb_tre	core_abb_out
core_dbb_oe core_dbb_tre	core_dbb_out
core_ape_oe core_ape_tre	core_ape
core_dpe_oe core_dpe_tre	core_dpe
core_artry_oe core_artry_tre	core_artry_out
core_ckstp_oe core_ckstp_tre	core_ckstp_out
core_outputs_oe	G2 core: core_qreq, core_br, core_rsrv, core_iabr Additional G2_LE core signals: core_iabr2, core_dabr, core_dabr2. If desired, the core_outputs_oe can be used to qualify the signals for JTAG, checkstop, and hreset states.
core_tdo_oe	core_tdo core_tdo is always driven regardless of the state of core_tdo_oe. core_tdo_oe is asserted when lssd_mode is asserted, lssd scanning, or JTAG scanning (valid core_tdo).

8.2.2.1 Unidirectional/Bidirectional Signals

Table 8-2 illustrates the conditions for setting up uni- or bidirectional signals external to the core, showing how the high-impedance control signal should be tied.

Table 8-2. Conditions for Unidirectional/Bidirectional Signals

Signal Type	core_XXX_tre (Input Signal)	core_XXX_oe (Output Signal)	core_XXX_ien (Output Signal)	core_XXX_in, core_XXX_out
Unidirectional	Negated (tie low)	Output enable signal is used when valid data is presented to the system level logic.	Input enable signal is used when valid data is required to be presented to the internal core logic.	Used as unidirectional signals to or from the system level logic interface.
Bidirectional	Asserted (tie high)	Typically not used or if it is used, it has similar conditions as for a unidirectional signal.	Typically not used or if it is used, it has similar conditions as for a unidirectional signal.	Input and output signals are wire-ORed together to generate bidirectional signal with value of high, low, or high-impedance appropriately.

8.2.2.2 Logic Gate Equivalent and Bidirectional Signals

All bidirectional signals from the 60x bus interface are implemented as separate input, output, output enable, and high-impedance enable signals, and in some cases, there is an input enable signal. Figure 8-2 shows an example of how these signals can be used to create a bidirectional signal outside of the core.

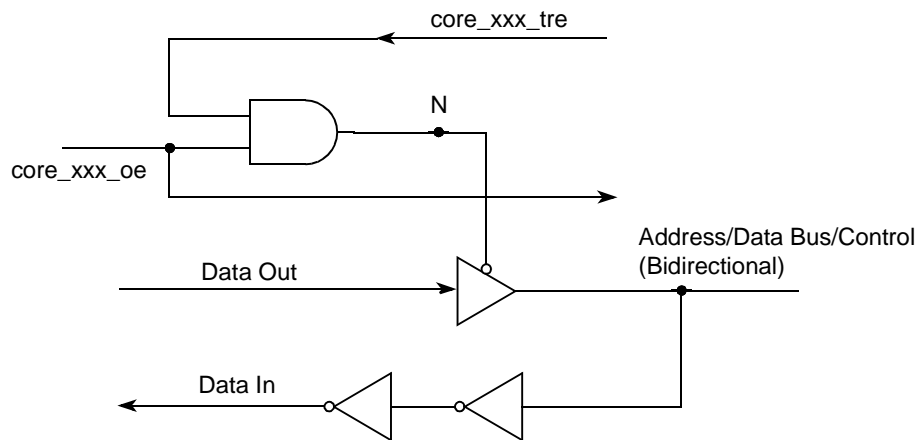


Figure 8-2. Logic Diagram for Bidirectional Signals

Table 8-3 represents the following conditions for a bidirectional signal created by wire-ORing the input and output signals from the core:

- If `core_XXX_tre = 1`, the `core_XXX_oe` signal controls the output on the address or data bus.
- If `core_XXX_tre = 0`, the data is driven on the bidirectional signal.

Table 8-3. Truth Table for Bidirectional Signals

core_xxx_tre	core_xxx_oe	Output on Node N	Address/Data Bus/Control
0	0	1	Drive output
0	1	1	Drive output
1	0	0	High Impedance
1	1	1	Drive output

8.2.3 Signal Summary

Table 8-4 provides alphabetically-ordered G2 core signals with related cross-reference that are relevant to the user. It details the signal name, signal grouping, number of signals, and whether the signal is an input or an output. It also lists which output enable, input enable and high-impedance control signal corresponds to the signal. Finally, the table provides a pointer to the section in this chapter where the signal function is described.

Table 8-4. G2 Core Signal Cross Reference

Signal (or Signal Pair)	Signal Name	Functional Grouping	Corresponding ien, oe, and tre	No. of Signals	I/O	Section No.
core_32bitmode	32-bit mode	Reset config.	—	1	I	8.3.10.3.1
core_a_in[0:31]	Address bus	Address transfer	—	32	I	8.3.3.1
core_a_out[0:31]			core_a_oe	32	O	
core_a_oe	Address bus output enable	Output enable	—	1	O	
core_a_tre	Address bus high-impedance enable	High-impedance control	—	1	I	
core_aack	Address acknowledge	Address termination	—	1	I	8.3.5.1
core_abb_in	Address bus busy	Address arbitration	—	1	I	8.3.1.3
core_abb_out			core_abb_oe	1	O	
core_abb_oe	abb output enable	Output enable	—	1	O	
core_abb_tre	abb high-impedance enable	High-impedance control	—	1	I	
core_ap_in[0:3]	Address bus parity	Address transfer	core_ap_ien	4	I	8.3.3.2
core_ap_out[0:3]			core_a_oe	4	O	
core_ap_ien	Address bus parity input enable	Input enable	—	1	O	
core_ape	Address parity error	Address transfer	core_ape_oe	1	O	8.3.3.3
core_ape_oe	ape output enable	Output enable	—	1	O	
core_ape_tre	ape high-impedance enable	High-impedance control	—	1	I	

Table 8-4. G2 Core Signal Cross Reference (continued)

Signal (or Signal Pair)	Signal Name	Functional Grouping	Corresponding ien, oe, and tre	No. of Signals	I/O	Section No.
$\overline{\text{core_artry_in}}$	Address retry	Address termination	—	1	I	8.3.5.2
$\overline{\text{core_artry_out}}$			core_artry_oe	1	O	
core_artry_oe	Address retry output enable	Output enable	—	1	O	
core_artry_tre	Address retry high-impedance enable	High-impedance control	—	1	I	
$\overline{\text{core_bg}}$	Bus grant	Address arbitration	—	1	I	8.3.1.2
$\overline{\text{core_br}}$	Bus request	Address arbitration	core_outputs_oe	1	O	8.3.1.1
$\overline{\text{core_ci}}$	Cache inhibit	Transfer attribute	core_a_oe	1	O	8.3.4.5
$\overline{\text{core_cint}}^1$	Critical interrupt	Interrupt, checkstop	—	1	I	8.3.9.2
core_clk_out	Test clock	Clocks	core_outputs_oe	1	O	8.3.15.2
$\overline{\text{core_ckstp_in}}$	Checkstop	Interrupt, checkstop	—	1	I	8.3.9.5
$\overline{\text{core_ckstp_out}}$			core_ckstp_oe	1	O	
core_ckstp_oe	Checkstop output enable	Output enable	—	1	O	
core_ckstp_tre	Checkstop high-impedance enable	High-impedance control	—	1	I	
$\text{core_cse}[0:1]$	Cache set entry	Transfer attribute	core_a_oe	2	O	8.3.4.8
core_d_oe	Data bus output enable	Output enable	core_d_tre	1	O	8.3.7.1.4
$\overline{\text{core_d_tre}}$	Data bus high-impedance enable	High-impedance control	core_d_oe	1	I	
$\overline{\text{core_dabr}}^1$	dabr1 watchpoint	Debug control	core_outputs_oe	1	O	8.3.14.3
$\overline{\text{core_dabr2}}^1$	dabr2 watchpoint		core_outputs_oe	1	O	8.3.14.4
$\overline{\text{core_dbb_in}}$	Data bus busy	Data arbitration	—	1	I	8.3.6.3
$\overline{\text{core_dbb_out}}$			core_dbb_oe	1	O	
core_dbb_oe	Data bus busy output enable	Output enable	—	1	O	
core_dbb_tre	Data bus busy high-impedance enable	High-impedance control	—	1	I	
$\overline{\text{core_dbg}}$	Data bus grant	Data arbitration	—	1	I	8.3.6.1
$\overline{\text{core_dbdis}}$	Data bus disable	Data transfer	—	1	I	8.3.7.4
$\overline{\text{core_dbwo}}$	Data bus write only	Data arbitration	—	1	I	8.3.6.2
$\text{core_dh_in}[0:31]$	Data bus high	Data transfer	core_dh_ien	32	I	8.3.7.1
$\text{core_dh_out}[0:31]$			core_d_oe	32	O	
core_dh_ien	dh input enable	Input enable	—	1	O	

Table 8-4. G2 Core Signal Cross Reference (continued)

Signal (or Signal Pair)	Signal Name	Functional Grouping	Corresponding ien, oe, and tre	No. of Signals	I/O	Section No.
core_disable	Disable	Test interface	—	1	I	8.3.13.1
core_dl_in[0:31]	Data bus low	Data transfer	core_dl_ien	1	O	8.3.7.1
core_dl_out[0:31]			—	32	O	
core_dl_ien	dl input enable	Input enable	—	1	O	
core_dp_in[0:7]	Data bus parity	Data transfer	core_dp_ien	8	I	8.3.7.2
core_dp_out[0:7]			core_d_oe	8	O	
core_dp_ien	dp input enable	Input enable	—	1	O	
core_dpe	Data parity error	Data transfer	core_dpe_oe	1	O	8.3.7.3
core_dpe_oe	dpe output enable	Output enable	—	1	O	
core_dpe_tre	dpe high-impedance enable	High-impedance control	—	1	I	
core_drtry	Data retry	Data termination	—	1	I	8.3.8.2
core_drtrymode	Data retry mode	Reset config.	—	1	I	8.3.10.3.4
core_gbl_in	Global	Transfer attribute	—	1	I	8.3.4.7
core_gbl_out			core_a_oe	1	O	
core_hreset	Hard reset	Reset	—	1	I	8.3.10.1
core_int	Interrupt	Interrupt, checkstop	—	1	I	8.3.9.1
core_iabr	IABR1 watchpoint	Debug control	core_outputs_oe	1	I	8.3.14.1
core_iabr ² ₁	IABR2 watchpoint		core_outputs_oe	1	I	8.3.14.2
core_l1_tstclk	LSSD test clocks	Test interface	—	1	I	8.3.13.2
core_l2_tstclk			—	1	I	
core_lssd_mode	LSSD test control signals		—	1	I	8.3.13.3
core_mcp	Machine check	Interrupt, checkstop	—	1	I	8.3.9.4
core_msrip	MSR IP	Reset config.	—	1	I	8.3.10.3.3
core_outputs_oe	Core outputs enable	Output enable	—	1	O	8.3.11.5.1
core_pll_cfg[0:4]	PLL configuration	Clocks	—	5	I	8.3.15.3
core_qack	Quiescent acknowledge	Processor status	—	1	I	8.3.11.1
core_qreq	Quiescent request		core_outputs_oe	1	O	8.3.11.2
core_redpinmode	Reduced pinout mode	Reset config.	—	1	I	8.3.10.3.2
core_rsrv	Reservation	Processor status	core_outputs_oe	1	O	8.3.11.3
core_smi	System management interrupt	Interrupt, checkstop	—	1	I	8.3.9.3

Table 8-4. G2 Core Signal Cross Reference (continued)

Signal (or Signal Pair)	Signal Name	Functional Grouping	Corresponding ien, oe, and tre	No. of Signals	I/O	Section No.
$\overline{\text{core_sreset}}$	Soft reset	Reset	—	1	I	8.3.10.2
$\text{core_svr}[0:31]$ ¹	System version register	Reset config.	—	32	I	8.3.10.3.6
core_sysclk	System clock	Clocks	—	1	I	8.3.15.1
$\overline{\text{core_ta}}$	Transfer acknowledge	Data termination	—	1	I	8.3.8.1
core_tap_en	Test access point enable	Test interface	—	1	I	8.3.12.6
core_tben	Time base enable	Processor status	—	1	I	8.3.11.4
$\overline{\text{core_tbst_in}}$	Transfer burst	Transfer attribute	—	1	I	8.3.4.3
$\overline{\text{core_tbst_out}}$			core_a_oe	1	O	
$\text{core_tc}[0:1]$	Transfer code	Transfer attribute	core_a_oe	2	O	8.3.4.4
core_tck	JTAG test clock	JTAG/COP interface	—	1	I	8.3.12.1
core_tdi	JTAG test data	JTAG/COP interface	—	1	I	8.3.12.2
core_tdo			core_tdo_oe	1	O	8.3.12.3
core_tdo_oe	tdo output enable	Output enable	—	1	O	8.3.12.3.1
$\overline{\text{core_tea}}$	Transfer error acknowledge	Data termination	—	1	I	8.3.8.3
$\overline{\text{core_tlbisync}}$	TLBI sync	Processor status	—	1	I	8.3.11.5
core_tle ¹	True little-endian mode	Interrupts, checkstops, reset	—	1	I	8.3.10.3.5
core_tlmsel	Test linking module select	Test interface	—	1	O	8.3.12.7
core_tms	JTAG test mode select	JTAG/COP interface	—	1	I	8.3.12.4
$\overline{\text{core_trst}}$	JTAG test reset	JTAG/COP interface	—	1	I	8.3.12.5
$\overline{\text{core_ts_in}}$	Transfer start	Address start	—	1	I	8.3.2.1
$\overline{\text{core_ts_out}}$			core_a_oe	1	O	
$\text{core_tsiz}[0:2]$	Transfer size	Transfer attribute	core_a_oe , core_abb_oe	3	O	8.3.4.2
$\overline{\text{core_tt_in}}[0:4]$	Transfer type	Transfer attribute	—	5	I	8.3.4.1
$\overline{\text{core_tt_out}}[0:4]$			core_a_oe	5	O	
$\overline{\text{core_wt}}$	Write-through	Transfer attribute	core_a_oe	1	O	8.3.4.6

¹ G2_LE only.

8.3 Signal Descriptions

This section describes individual G2 core signals, grouped according to Figure 8-1. Note that the following sections are intended to provide a quick summary of signal functions. Chapter 9, “Core Interface Operation,” describes many of these signals in greater detail, both with respect to how individual signals function and how groups of signals interact.

Figure 8-3 shows the G2 core signals groups in greater detail. However, it does not show both the input and output versions of the signals, their directions (input/output), and the associated input/output enable signals.

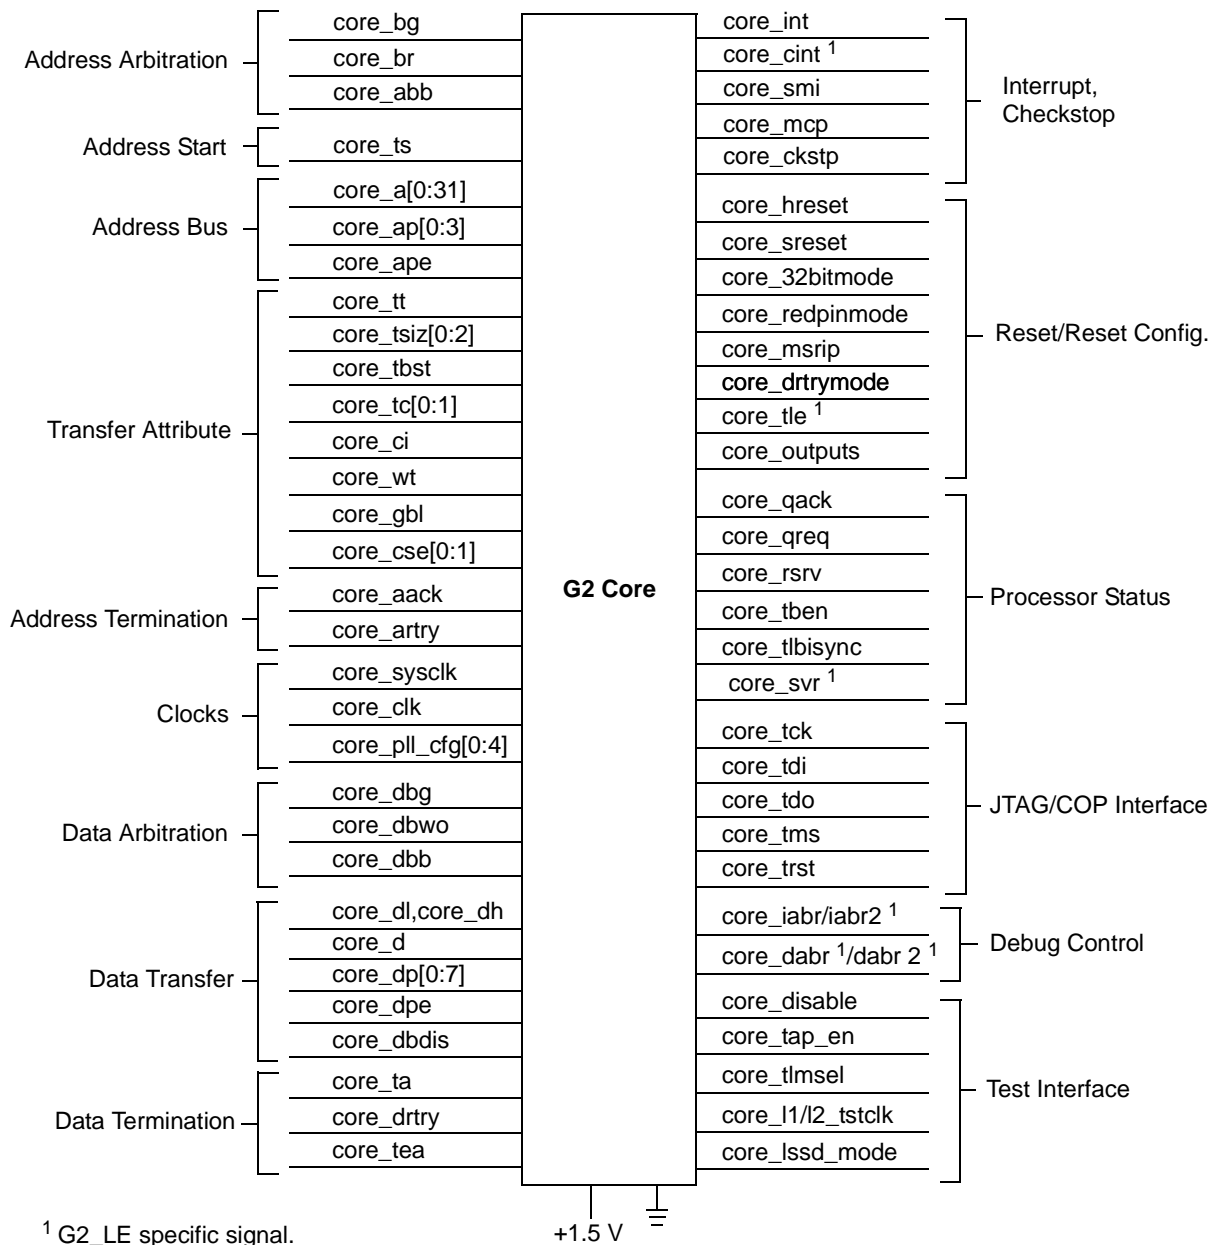


Figure 8-3. Detailed Signal Groups

8.3.1 Address Bus Arbitration Signals

The address arbitration signals are a collection of input and output signals that the G2 core uses to request the 60x address bus, recognize when the request is granted, and indicate when mastership is granted. For a detailed description of how these signals interact, see Section 9.3.1, “Address Bus Arbitration.”

8.3.1.1 Bus Request ($\overline{\text{core_br}}$)—Output

The $\overline{\text{core_br}}$ signal is an output on the G2 core. Following are the state meaning and timing comments for $\overline{\text{core_br}}$.

State Meaning Asserted—Indicates that the core is requesting mastership of the 60x address bus. Note that $\overline{\text{core_br}}$ may be asserted for one or more cycles, and then negated due to an internal cancellation of the bus request such as a load hit in the touch load buffer. See Section 9.3.1, “Address Bus Arbitration.”

Negated—Indicates that the core is not requesting the 60x address bus. The core may have no bus operation pending, it may be parked, or $\overline{\text{core_artry_in}}$ was asserted on the previous bus clock cycle.

Timing Comments Assertion—Occurs when the core is not parked and a bus transaction is needed. This may occur even if the two possible pipeline accesses have occurred. $\overline{\text{core_br}}$ is also asserted for one cycle during the execution of a **dcbz** instruction and during the execution of a load instruction that hits in the touch load buffer.

Negation—Occurs for at least one bus clock cycle after an accepted, qualified bus grant (see Section 8.3.1.3.1, “Address Bus Busy In ($\overline{\text{core_abb_in}}$,”), even if another transaction is pending. It is also negated for at least one bus clock cycle when the assertion of $\overline{\text{core_artry_in}}$ is detected on the bus.

8.3.1.2 Bus Grant ($\overline{\text{core_bg}}$)—Input

The $\overline{\text{core_bg}}$ signal is an input on the G2 core. A qualified bus grant occurs when either:

- $\overline{\text{core_bg}}$ is asserted, and $\overline{\text{core_abb_out}}$ and $\overline{\text{core_artry_out}}$ (after $\overline{\text{core_aack}}$) are negated, or
- $\overline{\text{core_bg}}$ is asserted, and $\overline{\text{core_abb_in}}$ and $\overline{\text{core_artry_in}}$ (after $\overline{\text{core_aack}}$) are negated.

$\overline{\text{core_abb}}$ and $\overline{\text{core_artry}}$ are both inputs and outputs on the G2 core and are driven by the core or other bus masters. If the core is parked, $\overline{\text{core_br}}$ need not be asserted for the qualified bus grant. Following are the state meaning and timing comments for $\overline{\text{core_bg}}$.

State Meaning	<p>Asserted—Indicates that the G2 core may, with the proper qualification, assume mastership of the 60x address bus. See Section 9.3.1, “Address Bus Arbitration.”</p> <p>Negated—Indicates that the core is not the next potential address bus master.</p>
Timing Comments	<p>Assertion—May occur at any time to indicate the core is free to use the address bus. After the core assumes bus mastership, it does not check for a qualified bus grant again until the cycle during which the address bus tenure is completed (<u>assuming it has another transaction to run</u>). The core <u>does not accept a core_bg in the cycles between the assertion of either core_ts_in or core_ts_out and core_aack</u>.</p> <p>Negation—May occur anytime to indicate the core cannot use the bus. The core may <u>still assume bus mastership on the bus clock cycle of the negation of core_bg because during the previous cycle core_bg indicated to the core that it was free to take mastership (if qualified)</u>.</p>

8.3.1.3 Address Bus Busy

There is both an address bus busy input and address bus busy output signal on the G2 core. The core also implements address bus busy output enable and address bus busy high-impedance enable signals.

8.3.1.3.1 Address Bus Busy In (core_abb_in)

Following are the state meaning and timing comments for core_abb_in.

State Meaning	<p>Asserted—Indicates that the address bus is in use. This condition effectively blocks the core from assuming address bus ownership, regardless of the <u>core_bg</u> input; see Section 9.3.1, “Address Bus Arbitration.”</p> <p>Negated—Indicates that the address bus is not owned by another bus master and that it is available to the core when accompanied by a qualified bus grant.</p>
Timing Comments	<p>Assertion—May occur when the core must be prevented from using the address bus (and the processor is not currently asserting <u>core_abb_out</u>).</p> <p>Negation—May occur whenever the core can use the address bus.</p>

8.3.1.3.2 Address Bus Busy Out (core_abb_out)

The core also implements address bus busy output enable and address bus busy high-impedance enable signals. core_abb_out acts as follows:

- If core_abb_tre is asserted, the output is in one of the following three states—high impedance, driven high, or driven low.
- If core_abb_tre is negated, the output is either driven to the high or low state. In this case, a valid value on core_abb_out exists when core_abb_oe is asserted.

Following are the state meaning and timing comments for core_abb_out.

State Meaning Asserted—Indicates that the core is the 60x address bus master. See Section 9.3.1, “Address Bus Arbitration.”

Negated—Indicates that the core is not using the address bus. If core_abb_out is negated during the bus clock cycle following a qualified bus grant, the core does not accept mastership, even if core_br is asserted. This can occur if a potential transaction is aborted internally before the transaction is started.

Timing Comments Assertion—Occurs on the bus clock cycle following a qualified core_bg that is accepted by the processor.

Negation—Occurs for a minimum of one-half bus clock cycle following the assertion of core_aack. If core_abb_out is negated during the bus clock cycle following a qualified bus grant, the core does not accept mastership, even if core_br is asserted.

High Impedance—Occurs one-half clock cycle after core_abb_out is negated, after the negation of core_abb_oe, if core_abb_tre is asserted. If core_abb_tre is negated, core_abb_out is always driven.

8.3.1.3.3 Address Bus Busy Output Enable (core_abb_oe)—Output

core_abb_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_abb_oe.

State Meaning Asserted—Indicates that the core is driving a valid core_abb_out.

Negation—Indicates one of the following two conditions:

If core_abb_tre is negated, negated core_abb_oe indicates that the core is not driving a valid core_abb_out value.

If core_abb_tre is asserted, negated core_abb_oe indicates that core_abb_out is in the high-impedance state.

Timing Comments Assertion—Occurs one clock cycle after an accepted qualified bus grant and remains asserted for the duration of the address tenure.

Negation— Remains asserted for a minimum of one-half processor cycle (dependent on the clock mode) and cycle starts after the assertion of `core_aack`, and then negates.

Note that negation of `core_abb_oe` may force `core_abb_out` to the high-impedance state, if `core_abb_tre` is asserted.

8.3.1.3.4 Address Bus Busy High-Impedance Enable (`core_abb_tre`)—Input

Following are the state meaning and timing comments for `core_abb_tre`. `core_abb_tre` is a high-impedance enable signal on the G2 core and can be used to create an external bidirectional `core_abb` signal. When the related input/output signals (`core_abb_in` and `core_abb_out`) are wire-ORed together, the resulting signal functions to a bidirectional 60x bus signal when `core_abb_tre` is asserted. See Section 8.2.2.2, “Logic Gate Equivalent and Bidirectional Signals,” for more information.

State Meaning Asserted—`core_abb_oe` controls whether `core_abb_out` is driven or forced to a high-impedance state.

Negated—Indicates that `core_abb_out` is always driven.

Timing Comments Assertion/Negation—Must be set up prior to the negation of `core_hreset` signal and remain stable during core operation. This is a static configuration.

8.3.2 Address Transfer Start Signals

Address transfer start signals are input and output signals that indicate that an address bus transfer has begun. For detailed information about how the transfer start signals interact with other signals, refer to Section 9.3.2, “Address Transfer.”

8.3.2.1 Transfer Start

There is both a transfer start input and transfer start output signal on the G2 core.

8.3.2.1.1 Transfer Start In (`core_ts_in`)

Following are the state meaning and timing comments for `core_ts_in`.

State Meaning Asserted—Indicates that another master has begun a bus transaction and that the address bus and transfer attribute signals are valid for snooping (see `core_gbl_in`).

Negated—Indicates that no bus transaction is occurring.

Timing Comments Assertion—May occur during the assertion of `core_abb_in`.

Negation—Must occur one bus clock cycle after `core_ts_in` is asserted.

8.3.2.1.2 Transfer Start Out (core_ts_out)

Following are the state meaning and timing comments for core_ts_out.

State Meaning Asserted—Indicates that the core has begun a memory bus transaction and that the address bus and transfer attribute signals are valid. When asserted with the appropriate core_tt[0:4] signals, it is also an implied data bus request for a memory transaction (unless it is an address-only operation).

Negated—Indicates that no bus transaction is occurring during normal operation.

Timing Comments Assertion—Coincides with the assertion of core_abb_out.

Negation—Occurs one bus clock cycle after core_ts_out is asserted.

High Impedance—Coincides with the negation of core_abb_out.

8.3.3 Address Transfer Signals

The address transfer signals are used to transmit the address and to generate and monitor parity for the 60x address bus transfer. For a detailed description of how these signals interact, refer to Section 9.3.2, “Address Transfer.”

8.3.3.1 Address Bus

The G2 core address bus consists of 32 input and 32 output signals along with output enable and high-impedance enable signals.

8.3.3.1.1 Address Bus In (core_a_in[0:31])

Following are the state meaning and timing comments for core_a_in[0:31].

State Meaning Asserted/Negated—Represents the physical address of a snoop operation.

Timing Comments Assertion/Negation—Must occur on the same bus clock cycle as the assertion of core_ts_in; is sampled by the core only on this cycle.

8.3.3.1.2 Address Bus Out (core_a_out[0:31])

The core also implements address bus output enable and address bus high-impedance enable signals. core_a_out[0:31] act as follows:

- If core_a_tre is asserted, the outputs are in one of the following three states—high impedance, driven high, or driven low.
- If core_a_tre is negated, the outputs are either driven to the high or low state. In this case, valid values on core_a_out[0:31] exist when core_a_oe is asserted.

Following are the state meaning and timing comments for `core_a_out[0:31]`.

State Meaning Asserted/Negated—Represents the physical address (real address) of the data to be transferred. On burst transfers, the address bus out signal presents the double-word-aligned address containing the critical code or data that missed the cache on a read operation, or the first double word of the cache line on a write operation. Note that the address output during burst operations is not incremented. See Section 9.3.2, “Address Transfer.”

Timing Comments Assertion/Negation—Occurs on the bus clock cycle after a qualified bus grant (coincides with assertion of `core_abb_out` and `core_ts_out`).

High Impedance—Occurs one bus clock cycle after `core_aack` is asserted, after the negation of `core_a_oe`, if `core_a_tre` is asserted. If `core_a_tre` is negated, `core_a_out[0:31]` are always driven.

8.3.3.1.3 Address Bus Output Enable (`core_a_oe`)—Output

`core_a_oe` is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for `core_a_oe`.

State Meaning Asserted—Indicates that the core is driving a valid `core_a_out[0:31]`.

Negated—Indicates one of the following two conditions:

If `core_a_tre` is negated, negated `core_a_oe` indicates that the core is not driving valid `core_a_out[0:31]` values.

If `core_a_tre` is asserted, negated `core_a_oe` indicates that `core_a_out[0:31]` are in the high-impedance state.

Timing Comments Assertion/Negation—Occurs on the bus clock cycle after a qualified bus grant (coincides with assertion of `core_abb_out` and `core_ts_out`).

Note that negation of `core_a_oe` may force `core_a_out[0:31]` to the high-impedance state, if `core_a_tre` is asserted.

8.3.3.1.4 Address Bus High-Impedance Enable (`core_a_tre`)—Input

Following are the state meaning and timing comments for `core_a_tre`. `core_a_tre` is a high-impedance enable signal on the G2 core and can be used to create an external bidirectional `core_a_out[0:31]` bus. When the related input/output signals (`core_a_in[0:31]` and `core_a_out[0:31]`) are wire-ORed together, the resulting bus functions similar to a bidirectional 60x address bus when `core_a_tre` is asserted. See Section 8.2.2.2, “Logic Gate Equivalent and Bidirectional Signals,” for more information.

State Meaning	<p>Asserted—core_a_oe controls whether core_a_out[0:31] are driven or forced to a high-impedance state.</p> <p>Negated—Indicates that core_a_out[0:31] are always driven.</p>
Timing Comments	<p><u>Assertion/Negation</u>—Must be set up prior to negation of the core_hreset signal and remain stable during core operation. This is a static configuration.</p>

8.3.3.2 Address Bus Parity

There are both address bus parity input and output signals reflecting 1 bit of odd-byte parity for each of the 4 bytes of address when a valid address is on the bus. The G2 core also implements an address bus parity input enable signal.

8.3.3.2.1 Address Bus Parity In (core_ap_in[0:3])

Following are the state meaning and timing comments for core_ap_in[0:3].

State Meaning	<p>Asserted/Negated—Represents odd parity for each of 4 bytes of the physical address for snooping operations. Detected even parity causes the processor to take a machine check exception or enter the checkstop state if address parity checking is enabled in the HID0 register; see Section 2.1.2.1, “Hardware Implementation Register 0 (HID0).” (See also the core_ape signal description.)</p>
Timing Comments	<p><u>Assertion/Negation</u>—The same as core_a_in[0:31].</p>

8.3.3.2.2 Address Bus Parity Input Enable (core_ap_ien)—Output

core_ap_ien is an input-enable indicator for its corresponding bus signals. Following are the state meaning and timing comments for core_ap_ien when core_a_tre is negated.

State Meaning	<p>Asserted—Indicates that the G2 core is receiving valid address parity.</p> <p>Negated—Indicates that the address parity input data is ignored.</p>
Timing Comments	<p><u>Assertion/Negation</u>—Valid values must be presented on core_ap_in[0:3] when core_ap_ien is asserted to the system logic.</p>

8.3.3.2.3 Address Bus Parity Out (core_ap_out[0:3])

Following are the state meaning and timing comments for core_ap_out[0:3].

State Meaning	<p>Asserted/Negated—Represents odd parity for each of 4 bytes of the physical address for a transaction. Odd parity means that an odd</p>
----------------------	---

number of bits, including the parity bit, are driven high. The signal assignments correspond to the following:

core_ap_out0	core_a_out[0:7]
core_ap_out1	core_a_out[8:15]
core_ap_out2	core_a_out[16:23]
core_ap_out3	core_a_out[24:31]

For more information, see Section 9.3.2.1, “Address Bus Parity.”

Timing Comments Assertion/Negation—The same as core_a_out[0:31].

High Impedance—The same as core_a_out[0:31].

8.3.3.3 Address Parity Error (core_ape)—Output

core_ape is an output signal on the G2 core. The core also implements address parity error output enable and address parity error high-impedance enable signals. core_ape acts as follows:

- If core_ape_tre is asserted, the output is in one of the following three states—high impedance, driven high, or driven low.
- If core_ape_tre is negated, the output is either driven to the high or low state. In this case, a valid value on core_ape exists when core_ape_oe is asserted.

When the corresponding high-impedance enable signal is negated, core_ape always drives to a valid logic state. The core_ape signal is not asserted if address parity checking is disabled (HID0[EBA] is cleared). For more information, see Section 9.3.2.1, “Address Bus Parity.” Following are the state meaning and timing comments for the core_ape signal on the G2 core.

State Meaning Asserted—Indicates that incorrect address bus parity has been detected by the core on a snoop (core_gbl_in is asserted).

Negated—Indicates that the core has not detected a parity error (even parity) on the address bus.

Timing Comments Assertion—Occurs on the second bus clock cycle after core_ts_in is asserted.

Negation/High Impedance—Occurs on the third bus clock cycle after core_ts_in is asserted, after the negation of core_ape_oe, if core_ape_tre is asserted. If core_ape_tre is negated, core_ape is always driven.

8.3.3.3.1 Address Parity Error Output Enable (core_ape_oe)—Output

core_ape_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_ape_oe.

State Meaning	<p>Asserted—Indicates that the G2 core is driving a valid $\overline{\text{core_ape}}$.</p> <p>Negated—Indicates one of the following two conditions:</p> <p>If core_ape_tre is negated, $\overline{\text{negated core_ape_oe}}$ indicates that the core is not driving a valid core_ape value.</p> <p>If core_ape_tre is asserted, $\overline{\text{negated core_ape_oe}}$ indicates that core_ape is in the high-impedance state.</p>
Timing Comments	<p>Assertion—core_ape_oe is asserted on the second bus clock after the assertion of core_ts_in.</p> <p>Negation—Occurs on the third bus clock cycle after $\overline{\text{core_ts_in}}$ is asserted.</p> <p>Note that negation of core_ape_oe may force $\overline{\text{core_ape}}$ to the high-impedance state, if core_ape_tre is asserted.</p>

8.3.3.3.2 Address Parity Error High-Impedance Enable (core_ape_tre)—Input

Following are the state meaning and timing comments for core_ape_tre . core_ape_tre is a high-impedance enable signal on the G2 core and can be used to create a three-statable version of core_ape externally. The resulting core_ape output signal functions similar to a bidirectional 60x bus signal when core_ape_tre is asserted.

State Meaning	<p>Asserted—core_ape_oe controls whether $\overline{\text{core_ape}}$ is driven or forced to a high-impedance state.</p> <p>Negated—Indicates that $\overline{\text{core_ape}}$ is always driven.</p>
Timing Comments	<p>Assertion/Negation—Must be set up prior to negation of the core_hreset signal and remain stable during core operation. This is a static configuration.</p>

8.3.4 Address Transfer Attribute Signals

The transfer attribute signals are a set of signals that further characterize the transfer—such as the size of the transfer, whether it is a read or write operation, and whether it is a burst or single-beat transfer. For a detailed description of how these signals interact, see Section 9.3.2, “Address Transfer.”

8.3.4.1 Transfer Type

The transfer type signals consist of five inputs and five outputs on the G2 core. For a complete description of the transfer type signals and for transfer type encodings, see Table 8-6.

8.3.4.1.1 Transfer Type In (core_tt_in[0:4])

Following are the state meaning and timing comments for core_tt_in[0:4].

State Meaning Asserted/Negated—Indicates the type of transfer in progress (see Table 8-5).

Timing Comments Assertion/Negation—The same as core_a_in[0:31].

Table 8-5 describes the 60x bus specification transfer encodings and the G2 core bus snoop response on an address hit.

Table 8-5. G2 Core Snoop Hit Response

60x Bus Specification Command	Transaction Type	core_tt_inx					G2 Core as Snooper; Action on Hit
		tt0	tt1	tt2	tt3	tt4	
Clean block	Address only	0	0	0	0	0	N/A
Flush block	Address only	0	0	1	0	0	N/A
sync	Address only	0	1	0	0	0	N/A
Kill block	Address only	0	1	1	0	0	Kill, cancel reservation
eieio	Address only	1	0	0	0	0	N/A
External control word write	Single-beat write	1	0	1	0	0	N/A
TLB invalidate	Address only	1	1	0	0	0	N/A
External control word read	Single-beat read	1	1	1	0	0	N/A
lwarx reservation set	Address only	0	0	0	0	1	N/A
Reserved	—	0	0	1	0	1	N/A
tlbsync	Address only	0	1	0	0	1	N/A
icbi	Address only	0	1	1	0	1	N/A
Reserved	—	1	X	X	0	1	N/A
Write-with-flush	Single-beat write or burst	0	0	0	1	0	Flush, cancel reservation
Write-with-kill	Single-beat write or burst	0	0	1	1	0	Kill, cancel reservation
Read	Single-beat read or burst	0	1	0	1	0	Clean or flush
Read-with-intent-to-modify	Burst	0	1	1	1	0	Flush
Write-with-flush-atomic	Single-beat write	1	0	0	1	0	Flush, cancel reservation
Reserved	N/A	1	0	1	1	0	N/A
Read-atomic	Single-beat read or burst	1	1	0	1	0	Clean or flush
Read-with-intent-to modify-atomic	Burst	1	1	1	1	0	Flush
Reserved	—	0	0	0	1	1	N/A
Reserved	—	0	0	1	1	1	N/A

Table 8-5. G2 Core Snoop Hit Response (continued)

60x Bus Specification Command	Transaction Type	core_tt_inx					G2 Core as Snooper; Action on Hit
		tt0	tt1	tt2	tt3	tt4	
Read-with-no-intent-to-cache	Single-beat read or burst	0	1	0	1	1	Clean
Reserved	—	0	1	1	1	1	N/A
Reserved	—	1	X	X	1	1	N/A

8.3.4.1.2 Transfer Type Out (core_tt_out[0:4])

Following are the state meaning and timing comments for core_tt_out[0:4].

State Meaning Asserted/Negated—Indicates the type of transfer in progress.

Timing Comments Assertion/Negation/High Impedance—The same as core_a_out[0:31].

Table 8-6 describes the transfer type encodings for the G2 core as a bus master.

Table 8-6. Transfer Type Encoding for the G2 Core as a Bus Master

G2 Core Bus Master Transaction	Transaction Source	core_tt_outx					60x Bus Specification Command	Transaction Type
		tt0	tt1	tt2	tt3	tt4		
N/A	N/A	0	0	0	0	0	Clean block	Address only
N/A	N/A	0	0	1	0	0	Flush block	Address only
N/A	N/A	0	1	0	0	0	sync	Address only
Address only	dcbz	0	1	1	0	0	Kill block	Address only
N/A	N/A	1	0	0	0	0	eieio	Address only
Single-beat write (nongbl)	ecowx	1	0	1	0	0	External control word write	Single-beat write
N/A	N/A	1	1	0	0	0	TLB invalidate	Address only
Single-beat read (nongbl)	eciwx	1	1	1	0	0	External control word read	Single-beat read
N/A	N/A	0	0	0	0	1	lwarx Reservation set	Address only
N/A	N/A	0	0	1	0	1	Reserved	—
N/A	N/A	0	1	0	0	1	tlbsync	Address only
N/A	N/A	0	1	1	0	1	icbi	Address only
N/A	N/A	1	X	X	0	1	Reserved	—
Single-beat write	Caching-inhibited or write-through store	0	0	0	1	0	Write-with-flush	Single-beat write or burst
Burst (nongbl)	Cast-out, or snoop copy back	0	0	1	1	0	Write-with-kill	Single-beat write or burst

Table 8-6. Transfer Type Encoding for the G2 Core as a Bus Master (continued)

G2 Core Bus Master Transaction	Transaction Source	core_tt_outx					60x Bus Specification Command	Transaction Type
		tt0	tt1	tt2	tt3	tt4		
Single-beat read	Caching-inhibited load or instruction fetch	0	1	0	1	0	Read	Single-beat read or burst
Burst	Load miss, store miss, or instruction fetch	0	1	1	1	0	Read-with-intent-to-modify	Burst
Single-beat write	stwcx.	1	0	0	1	0	Write-with-flush-atomic	Single-beat write
N/A	N/A	1	0	1	1	0	Reserved	N/A
Single-beat read	lwarx (caching-inhibited load)	1	1	0	1	0	Read-atomic	Single-beat read or burst
Burst	lwarx (load miss)	1	1	1	1	0	Read-with-intent-to-modify-atomic	Burst
N/A	N/A	0	0	0	1	1	Reserved	—
N/A	N/A	0	0	1	1	1	Reserved	—
N/A	N/A	0	1	0	1	1	Read-with-no-intent-to-cache	Single-beat read or burst
N/A	N/A	0	1	1	1	1	Reserved	—
N/A	N/A	1	X	X	1	1	Reserved	—

When `HID0[ABE]` is set, the G2 core performs address-only bus transactions with the encodings shown in Table 8-7.

Table 8-7. Implementation-Specific Transfer Type Encoding

Transaction Source	core_tt_out[0:4]x					Bus Command	Transaction Type
	tt0	tt1	tt2	tt3	tt4		
dcbst	0	0	0	0	0	Clean block	Address only
dcbf	0	0	1	0	0	Flush block	Address only
dcbz, dcbi ¹	0	1	1	0	0	Kill block	Address only

¹ The **dcbi** instruction should never be used on the G2 core.

8.3.4.2 Transfer Size (`core_tsiz[0:2]`)—Output

The `core_tsiz[0:2]` signals consist of three output signals on the G2 core. Following are the state meaning and timing comments for the `core_tsiz[0:2]` outputs.

State Meaning Asserted/Negated—For memory accesses, these signals along with `core_tbst_out`, indicate the data transfer size for the current bus operation, as shown in Table 8-8. Table 9-5 shows how the transfer

size signals are used with the address signals for aligned transfers. Table 9-6 shows how the transfer size signals are used with the address signals for misaligned transfers.

For external control instructions (**eciwx** and **ecowx**), **core_tsiz[0:2]** are used to output bits 29–31 of the external access register (EAR), which are used to form the resource ID (**core_tbst_out**||**core_tsiz[0:2]**).

Timing Comments Assertion/Negation—The same as **core_a_out[0:31]**.
High Impedance—The same as **core_a_out[0:31]**.

Table 8-8. Data Transfer Size

core_tbst_out	core_tsiz[0:2]	Transfer Size
Asserted	010	Burst (32 bytes)
Negated	000	8 bytes
Negated	001	1 byte
Negated	010	2 bytes
Negated	011	3 bytes
Negated	100	4 bytes
Negated	101	5 bytes
Negated	110	6 bytes
Negated	111	7 bytes

8.3.4.3 Transfer Burst

There is both a transfer burst input and transfer burst output signal on the G2 core.

8.3.4.3.1 Transfer Burst In (**core_tbst_in**)

Following are the state meaning and timing comments for **core_tbst_in**.

State Meaning Assertion/Negated—Used when snooping single-beat reads (read with no intent to cache) to indicate that a burst transfer is in progress.

Timing Comments Assertion/Negation—The same as **core_a_in[0:31]**.

8.3.4.3.2 Transfer Burst Out (**core_tbst_out**)

Following are the state meaning and timing comments for **core_tbst_out**.

State Meaning Asserted—Indicates that a burst transfer is in progress.

Negated—Indicates that a burst transfer is not in progress.

For external control instructions (**eciwx** and **ecowx**), $\overline{\text{core_tbst_out}}$ is used to output EAR[28], which is used to form the resource ID ($\overline{\text{core_tbst_out}}|\text{core_tsiz}[0:2]$).

Timing Comments Assertion/Negation—The same as $\text{core_a_out}[0:31]$.
High Impedance—The same as $\text{core_a_out}[0:31]$.

8.3.4.4 Transfer Code (**core_tc[0:1]**)—Output

The $\text{core_tc}[0:1]$ consists of two output signals on the G2 core. Following are the state meaning and timing comments for the $\text{core_tc}[0:1]$ outputs.

State Meaning Asserted/Negated—Represents a special encoding for the transfer in progress (see Table 8-9).

Timing Comments Assertion/Negation—The same as $\text{core_a_out}[0:31]$.
High Impedance—The same as $\text{core_a_out}[0:31]$.

Table 8-9. Encodings for $\text{core_tc}[0:1]$ Signals

core_tc(0:1)	Read	Write
0 0	Data transaction	Any write
0 1	Touch load	—
1 0	Instruction fetch	—
1 1	Reserved	—

8.3.4.5 Cache Inhibit (**core_ci**)—Output

Following are the state meaning and timing comments for the $\overline{\text{core_ci}}$ output.

State Meaning Asserted—Indicates that a single-beat transfer is not cached, reflecting the setting of the I bit for the block or page that contains the address of the current transaction.

Negated—Indicates that a burst transfer in progress will allocate a line in the G2 core data cache.

Timing Comments Assertion/Negation—The same as $\text{core_a_out}[0:31]$.
High Impedance—The same as $\text{core_a_out}[0:31]$.

8.3.4.6 Write-Through (**core_wt**)—Output

Following are the state meaning and timing comments for the $\overline{\text{core_wt}}$ output.

State Meaning Asserted—Indicates that a single-beat transaction is write-through, reflecting the value of the W bit for the block or page that contains the address of the current transaction.

Negated—Indicates that a transaction is not for memory area designated as write-through.

Timing Comments Assertion/Negation—The same as `core_a_out[0:31]`.
High Impedance—The same as `core_a_out[0:31]`.

8.3.4.7 Global Signals

There is both a global input and global output signal on the G2 core.

8.3.4.7.1 Global In (`core_gbl_in`)

Following are the state meaning and timing comments for `core_gbl_in`.

State Meaning Asserted—Indicates that a transaction must be snooped by the G2 core.
Negated—Indicates that a transaction is not to be snooped by the G2 core.

Timing Comments Assertion/Negation—The same as `core_a_in[0:31]`.

8.3.4.7.2 Global Out (`core_gbl_out`)

Following are the state meaning and timing comments for `core_gbl_out`.

State Meaning Asserted—Indicates that a transaction is global, reflecting the setting of the M bit for the block or page that contains the address of the current transaction (except in the case of copy-back operations and instruction fetches, which are nonglobal).
Negated—Indicates that a transaction is not global.

Timing Comments Assertion/Negation—The same as `core_a_out[0:31]`.
High Impedance—The same as `core_a_out[0:31]`.

8.3.4.8 Cache Set Entry (`core_cse[0:1]`)—Output

Following are the state meaning and timing comments for the `core_cse[0:1]` outputs.

State Meaning Asserted/Negated—Represents the cache replacement set element for the current transaction reloading into or writing out of the cache. Can be used with the address bus and the transfer attribute signals to externally track the state of each cache line in the G2 core cache. Note that `core_cse[0:1]` are not meaningful during data cache touch load operations.

Timing Comments Assertion/Negation—The same as `core_a_out[0:31]`.
High Impedance—The same as `core_a_out[0:31]`.

8.3.5 Address Transfer Termination Signals

The address transfer termination signals are used to indicate either that the address phase of the transaction has completed successfully or must be repeated, and when it should be terminated. For detailed information about how these signals interact, see Section 9.3.3, “Address Transfer Termination.”

8.3.5.1 Address Acknowledge (core_aack)—Input

Following are the state meaning and timing comments for the core_aack input.

State Meaning	<p>Asserted—Indicates that the address phase of a transaction is complete. Causes <u>core_a_oe</u> to <u>negate</u> on the next bus clock cycle. The G2 core also <u>samples core_artry_in</u> on the bus clock cycle simultaneous with <u>core_aack</u> and on the <u>bus cycle following</u> the assertion of <u>core_aack</u>. The assertion of <u>core_artry_in</u> on the bus clock cycle simultaneous with the assertion of <u>core_aack</u> is known as an early address retry.</p> <p>Negated—Indicates that the <u>address bus</u> and transfer attributes must remain driven when <u>core_abb_out</u> is asserted.</p>
Timing Comments	<p><u>Assertion</u>—May occur as early as the bus clock cycle after <u>core_ts_out</u> is asserted (<u>unless the G2 core is configured for 1:1 or 1.5:1 clock modes, when core_aack can be asserted no sooner than the second cycle following the assertion of core_ts_out</u>—one address wait state); assertion can be delayed to allow adequate address access time for slow devices. For example, if an implementation supports <u>slow snooping devices</u>, an external arbiter can postpone the assertion of <u>core_aack</u>.</p> <p><u>Negation</u>—Must occur one bus clock cycle after the assertion of <u>core_aack</u>.</p>

8.3.5.2 Address Retry

There is both an address retry input and address retry output signal on the G2 core. The core also implements address retry output enable and address retry high-impedance enable signals.

8.3.5.2.1 Address Retry In (core_artry_in)

Following are the state meaning and timing comments for core_artry_in.

State Meaning	<p>Asserted—If the G2 core is the address bus master, <u>core_artry_in</u> indicates that the <u>core must retry the preceding address tenure and immediately negate core_br</u> (if asserted). If the associated data</p>
----------------------	--

tenure has already started, the core also aborts the data tenure immediately, even if the burst data has been received. If the core is not the address bus master, this input indicates that the core should immediately negate core_br for one bus clock cycle following the assertion of core_artry_in by the snooping bus master to allow an opportunity for a copy-back operation to main memory. Note that the subsequent address presented on the address bus may not be the same one associated with the assertion of core_artry_in.

Negated—Indicates that the core does not need to retry the last address tenure.

Timing Comments Assertion—May occur as early as the second cycle following the assertion of core_ts_out, and must occur by the bus clock cycle immediately following the assertion of core_aack if an address retry is required.

Negation—Must occur during the second cycle after the assertion of core_aack.

8.3.5.2.2 Address Retry Out (core_artry_out)

The core also implements address retry output enable and address retry high-impedance enable signals. core_artry_out acts as follows:

- If core_artry_tre is asserted, the output is in one of the following three states—high impedance, driven high, or driven low.
- If core_artry_tre is negated, the output is either driven to the high or low state. In this case, a valid value on core_artry_out exists when core_artry_oe is asserted.

Following are the state meaning and timing comments for core_artry_out.

State Meaning Asserted—Indicates that the G2 core detects a condition in which a snooped address tenure must be retried. If the core needs to update memory as a result of the snoop that caused the retry, the core asserts core_br the second cycle after core_aack if core_artry_out is asserted.

Negated—Indicates that the core does not need the snooped address tenure to be retried.

Timing Comments Assertion—Asserted the third bus cycle following the assertion of core_ts_in if a retry is required and remains asserted until one cycle after the core_aack is asserted.

Negation—Occurs on the second bus cycle after the assertion of core_aack and remains asserted for a minimum of one-half bus cycle (depends on clock mode) before it is negated for one bus cycle.

High Impedance—Indicates that the core does not need the snooped address tenure to be retired. Occurs two bus cycles after core_aack is asserted, after the negation of core_artry_oe, if core_artry_tre is asserted. If core_artry_tre is negated, core_artry_out is always driven.

8.3.5.2.3 Address Retry Output Enable (core_artry_oe)—Output

core_artry_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_artry_oe.

State Meaning	<u>Asserted</u> —Indicates that the G2 core is driving a valid <u>core_artry_out</u> .
	<u>Negated</u> —Indicates one of the following two conditions: If <u>core_artry_tre</u> is negated, <u>negated core_artry_oe</u> indicates that the core is not driving a valid <u>core_artry_out</u> value. If <u>core_artry_tre</u> is asserted, <u>negated core_artry_oe</u> indicates that <u>core_artry_out</u> is in the high-impedance state.
Timing Comments	<u>Assertion</u> —Asserted the second bus cycle following the assertion of <u>core_ts_in</u> if a retry is required and it remains asserted until one bus cycle after <u>core_aack</u> is asserted.
	<u>Negation</u> —Occurs the second bus cycle after the assertion of <u>core_aack</u> and remains asserted for a minimum of one-half bus cycle (depends on clock mode) before it is negated for one bus cycle.
	Note that negation of <u>core_artry_oe</u> may force <u>core_artry_out</u> to the high-impedance state, if <u>core_artry_tre</u> is asserted.

8.3.5.2.4 Address Retry High-Impedance Enable (core_artry_tre)—Input

Following are the state meaning and timing comments for core_artry_tre. core_artry_tre is a high-impedance enable signal on the G2 core and can be used to create an external bidirectional core_artry signal. When the related input/output signals (core_artry_in and core_artry_out) are wire-ORed together, the resulting signal functions similar to a bidirectional 60x bus signal when core_artry_tre is asserted. See Section 8.2.2.2, “Logic Gate Equivalent and Bidirectional Signals,” for more information.

State Meaning	<u>Asserted</u> — <u>core_artry_oe</u> controls whether <u>core_artry_out</u> is driven or forced to a high-impedance state.
	<u>Negated</u> —Indicates <u>core_artry_out</u> is always driven.
Timing Comments	<u>Assertion/Negation</u> —Must be set up prior to negation of the <u>core_hreset</u> signal and remain stable during core operation. This is a static configuration.

8.3.6 Data Bus Arbitration Signals

Like the address bus arbitration signals, data bus arbitration signals maintain an orderly process for determining 60x data bus mastership. Note that there is no data bus arbitration signal equivalent to the address bus arbitration signal `core_br` (bus request) because, except for address-only transactions, `core_ts_out` implies data bus requests. For a detailed description on how these signals interact, see Section 9.4.1, “Data Bus Arbitration.”

One special signal, `core_dbwo`, allows the core to be configured dynamically to write data out of order with respect to read data. For detailed information about using `core_dbwo`, see Section 9.10, “Using `core_dbwo` (Data Bus Write Only).”

8.3.6.1 Data Bus Grant (`core_dbg`)—Input

Following are the state meaning and timing comments for the `core_dbg` input.

State Meaning Asserted—Indicates that the core may, with the proper qualification, assume mastership of the data bus. The core derives a qualified data bus grant when `core_dbg` is asserted and `core_dbb_out`, `core_drtry`, and `core_artry_out` are negated; that is, the data bus is not busy (`core_dbb_out` is negated), there is no outstanding attempt to retry the current data tenure (`core_drtry` is negated), and there is no outstanding attempt to perform an `core_artry_out` of the associated address tenure.

Negated—Indicates that the core must hold off its data tenures.

Timing Comments Assertion—May occur any time to indicate the core is free to take data bus mastership. It is not sampled until `core_ts_out` is asserted.

Negation—May occur at any time to indicate the core cannot assume data bus mastership.

8.3.6.2 Data Bus Write Only (`core_dbwo`)—Input

Following are the state meaning and timing comments for the `core_dbwo` input.

State Meaning Asserted—Indicates that the core may perform the data bus tenure for an outstanding write address even if a read address is pipelined before the write address. Refer to Section 9.10, “Using `core_dbwo` (Data Bus Write Only),” for detailed instructions on using `core_dbwo`.

Negated—Indicates that the core must perform the data bus tenures in the same order as the address tenures.

Timing Comments Assertion—Must occur no later than a qualified `core_dbg` for an outstanding write tenure. `core_dbwo` is only recognized by the core on the clock of a qualified `core_dbg`. If no write requests are pending,

the core ignores $\overline{\text{core_dbwo}}$ and assumes data bus ownership for the next pending read request.

Negation—May occur any time after a qualified $\overline{\text{core_dbg}}$ and before the next assertion of $\overline{\text{core_dbg}}$.

8.3.6.3 Data Bus Busy

There is both a data bus busy input and data bus busy output signal on the G2 core. Data bus busy output enable and high-impedance enable signals are also implemented on the G2 core.

8.3.6.3.1 Data Bus Busy In ($\overline{\text{core_dbb_in}}$)

Following are the state meaning and timing comments for $\overline{\text{core_dbb_in}}$.

State Meaning Asserted—Indicates that another device is the bus master.

Negated—Indicates that the data bus is free (with proper qualification, see $\overline{\text{core_dbg}}$) for use by the core.

Timing Comments Assertion—Must occur when the core must be prevented from using the data bus.

Negation—May occur whenever the data bus is available.

8.3.6.3.2 Data Bus Busy Out ($\overline{\text{core_dbb_out}}$)

The core also implements data bus busy output enable and data bus busy high-impedance enable signals. $\overline{\text{core_dbb_out}}$ acts as follows:

- If $\overline{\text{core_dbb_tre}}$ is asserted, the output is in one of the following three states—high impedance, driven high, or driven low.
- If $\overline{\text{core_dbb_tre}}$ is negated, the output is either driven to the high or low state. In this case, a valid value on $\overline{\text{core_dbb_out}}$ exists when $\overline{\text{core_dbb_oe}}$ is asserted.

Following are the state meaning and timing comments for $\overline{\text{core_dbb_out}}$.

State Meaning Asserted—Indicates that the core is the 60x data bus master. The G2 core always assumes data bus mastership if it needs the data bus and is given a qualified data bus grant (see $\overline{\text{core_dbg}}$).

Negated—Indicates that the core is not using the data bus.

Timing Comments Assertion—Occurs during the bus clock cycle following a qualified $\overline{\text{core_dbg}}$.

Negation—Occurs for a minimum of one-half bus clock cycle (dependent on clock mode) following the assertion of the final $\overline{\text{core_ta}}$.

High Impedance—Occurs after $\overline{\text{core_dbb_out}}$ is negated, after the negation of $\overline{\text{core_dbb_oe}}$, if core_dbb_tre is asserted. If core_dbb_tre is negated, core_dbb_out is always driven.

8.3.6.3.3 Data Bus Busy Output Enable (core_dbb_oe)—Output

core_dbb_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_dbb_oe .

State Meaning Asserted—Indicates that the core is driving a valid $\overline{\text{core_dbb_out}}$.

Negated—Indicates one of the following two conditions:

If core_dbb_tre is negated, $\overline{\text{negated core_dbb_oe}}$ indicates that the core is not driving a valid core_dbb_out value.

If core_dbb_tre is asserted, $\overline{\text{negated core_dbb_oe}}$ indicates that core_dbb_out is in the high-impedance state.

Timing Comments Assertion/Negation—Asserted after a qualified $\overline{\text{core_dbg}}$ is asserted. Remains asserted for a minimum of one-half bus clock cycle following the assertion of output signals core_ta , core_tea , or core_artry_out .

Note that negation of core_dbb_oe may force $\overline{\text{core_dbb_out}}$ to the high-impedance state, if core_dbb_tre is asserted.

8.3.6.3.4 Data Bus Busy High-Impedance Enable (core_dbb_tre)—Input

Following are the state meaning and timing comments for core_dbb_tre . core_dbb_tre is a high-impedance enable signal on the G2 core and can be used to create an external bidirectional $\overline{\text{core_dbb}}$ signal. When the related input/output signals ($\overline{\text{core_dbb_in}}$ and core_dbb_out) are wire-ORed together, the resulting signal functions similar to a bidirectional 60x bus signal when core_dbb_tre is asserted. See Section 8.2.2.2, “Logic Gate Equivalent and Bidirectional Signals,” for more information.

State Meaning Asserted— core_dbb_oe controls whether $\overline{\text{core_dbb_out}}$ is driven or forced to a high-impedance state.

Negated—Indicates that $\overline{\text{core_dbb_out}}$ is always driven.

Timing Comments Assertion/Negation—Must be set up prior to negation of the $\overline{\text{core_hreset}}$ signal and remain stable during core operation. This is a static configuration.

8.3.7 Data Transfer Signals

Like the address transfer signals, the data transfer signals are used to transmit data and to generate and monitor parity for the data transfer. For a detailed description of how the data transfer signals interact, see Section 9.4.3, “Data Transfer.”

8.3.7.1 Data Bus

The data bus consists of 64 input and 64 output signals on the G2 core. The data bus has two halves—data bus high (dh) and data bus low (dl). See Table 8-10 for the data bus lane assignments. The data bus is driven once for noncached transactions and four times for cache transactions (bursts). The dh and dl signals are split into input, output, and input enable signals on the G2 core.

Table 8-10. Data Bus Lane Assignments

Data Bus Signals	Byte Lane
dh[0:7]	0
dh[8:15]	1
dh[16:23]	2
dh[24:31]	3
dl[0:7]	4
dl[8:15]	5
dl[16:23]	6
dl[24:31]	7

8.3.7.1.1 Data Bus In (core_dh_in[0:31], core_dl_in[0:31])

Following are the state meaning and timing comments for core_dh_in[0:31] and core_dl_in[0:31].

State Meaning Asserted/Negated—Represents the state of data during a data read transaction.

Timing Comments Assertion/Negation—Data must be valid on the same bus clock cycle that core_ta is asserted.

8.3.7.1.2 Data Bus Input Enable (core_dh_ien, core_dl_ien)—Output

core_dh_ien and core_dl_ien are input enable indicators to their corresponding bus signals. Following are the state meaning and timing comments for core_dh_ien and core_dl_ien. Note that not all input signals have input enable signals.

State Meaning Asserted—Indicates that the G2 core is expecting valid data bus input.

Negated—Indicates that the received data bus input is ignored.

Timing Comments Assertion/Negation—Valid data must be present to data bus input signals when core_dh_ien or core_dl_ien is asserted to the system logic. These signals allow integrators to support either a bidirectional or unidirectional data bus interface.

8.3.7.1.3 Data Bus Out (core_dh_out[0:31], core_dl_out[0:31])—Output

The core also implements data bus output enable and data bus high-impedance enable signals. core_dh_out[0:31] and core_dl_out[0:31] act as follows:

- If core_d_tre is asserted, the outputs are in one of the following three states—high impedance, driven high, or driven low.
- If core_d_tre is negated, the outputs are either driven to the high or low state. In this case, valid values on core_dh_out[0:31] and core_dl_out[0:31] exist when core_d_oe is asserted.

Following are the state meaning and timing comments for core_dh_out[0:31] and core_dl_out[0:31].

State Meaning	Asserted/Negated—Represent the state of data during a data write. Byte lanes not selected for data transfer do not supply valid data.
Timing Comments	Assertion/Negation—Occurs one clock cycle after qualified data bus grant (coincides with core_dbb_out). High Impedance—Occurs on the bus clock cycle after core_ta is asserted, after the negation of core_d_oe, if core_d_tre is asserted. If core_d_tre is negated, core_dh_out[0:31] and core_dl_out[0:31] are always driven.

8.3.7.1.4 Data Bus Output Enable (core_d_oe)—Output

core_d_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_d_oe.

State Meaning	Asserted—Indicates that the core is driving a valid data and data parity. Negated—Indicates one of the following two conditions: If core_d_tre is negated, negated core_d_oe indicates that the core is not driving valid core_dh_out[0:31] and core_dl_out[0:31] values. If core_d_tre is asserted, negated core_d_oe indicates that core_dh_out[0:31] and core_dl_out[0:31] are in the high-impedance state.
Timing Comments	Assertion/Negation—Occurs one clock cycle after qualified data bus grant (coincides with core_dbb_out). Note that negation of core_d_oe may force the data bus and data attribute signals to the high-impedance state, if core_d_tre is asserted.

8.3.7.1.5 Data Bus High-Impedance Enable (core_d_tre)—Input

Following are the state meaning and timing comments for core_d_tre. core_d_tre is a high-impedance enable signal on the G2 core and can be used to create an external bidirectional data bus. When the related input/output signals are wire-ORed together, the resulting bus functions similar to a bidirectional 60x data bus when core_d_tre is asserted. See Section 8.2.2.2, “Logic Gate Equivalent and Bidirectional Signals,” for more information.

State Meaning Asserted—core_d_oe controls whether the data bus output signals are driven or forced to a high-impedance state.

Negated—Indicates that data bus signals are always driven.

Timing Comments Assertion/Negation—Must be set up prior to negation of the core_hreset signal and remain stable during core operation. This is a static configuration.

8.3.7.2 Data Bus Parity (DP[0:7])

There are eight data bus parity inputs and eight data bus parity output signals on the G2 core. The core also implements a data bus parity input enable signal. The byte assignments are listed in Table 8-11.

Table 8-11. Data Bus Parity Signal Assignments

Data Bus Parity Signal	Data Bus Byte Assignment
dp0	dh[0:7]
dp1	dh[8:15]
dp2	dh[16:23]
dp3	dh[24:31]
dp4	dl[0:7]
dp5	dl[8:15]
dp6	dl[16:23]
dp7	dl[24:31]

8.3.7.2.1 Data Bus Parity In (core_dp_in[0:7])

Following are the state meaning and timing comments for core_dp_in[0:7].

State Meaning Asserted/Negated—Should represent odd parity for each byte of read data. Parity is checked on all data byte lanes, regardless of the size of the transfer. Detected even parity causes a checkstop if data parity errors are enabled in the HID0 register. (See core_dpe.)

Timing Comments Assertion/Negation—The same as core_dl_in[0:31].

8.3.7.2.2 Data Bus Parity Input Enable (core_dp_ien)—Output

core_dp_ien is an input-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_dp_ien when core_dp_tre is negated.

State Meaning Asserted—Indicates that the G2 core is excepting valid data bus parity.

Negated—Indicates that the data bus parity input is ignored.

Timing Comments Assertion/Negation—Valid data must be presented to core_dp_in[0:7] when core_dp_ien is asserted to the system logic. These signals allow integrators to support either a bidirectional or unidirectional data bus parity interface.

8.3.7.2.3 Data Bus Parity Out (core_dp_out[0:7])

Following are the state meaning and timing comments for core_dp_out[0:7].

State Meaning Asserted/Negated—Represents odd parity for each of 8 bytes of data for write transactions. Odd parity means that an odd number of bits, including the parity bit, are driven high.

Timing Comments Assertion/Negation—The same as core_dl_out[0:31].

High Impedance—The same as core_dl_out[0:31].

8.3.7.3 Data Parity Error ($\overline{\text{core_dpe}}$)—Output

The $\overline{\text{core_dpe}}$ signal is an output signal (output-only) on the G2 core. The core also implements data parity error output enable and data parity error high-impedance enable signals. core_dpe acts as follows:

- If core_dpe_tre is asserted, the output is in one of the following three states—high impedance, driven high, or driven low.
- If core_dpe_tre is negated, the output is either driven to the high or low state. In this case, a valid value on core_dpe exists when core_dpe_oe is asserted.

Following are the state meaning and timing comments for the $\overline{\text{core_dpe}}$ output.

State Meaning Asserted—Indicates that incorrect data bus parity was detected during a read transaction when HID0[EBD] is enabled. Internally, the core can take a machine check interrupt or enter a checkstop state.

Negated—Indicates correct data bus parity on the data bus.

Timing Comments Assertion—Occurs on the second bus clock cycle after $\overline{\text{core_ta}}$ is asserted to the core, unless $\overline{\text{core_ta}}$ is canceled by an assertion of core_drtry or core_artry (in certain cases).

Negation/High Impedance—Occurs on the third bus clock cycle after core_ta is asserted, if core_dpe_tre is asserted. If core_dpe_tre is negated, core_dpe is always driven.

8.3.7.3.1 Data Parity Error Output Enable (core_dpe_oe)—Output

core_dpe_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_dpe_oe.

- State Meaning** Asserted—Indicates that the core is driving a valid core_dpe.
Negated—Indicates one of the following two conditions:
 If core_dpe_tre is negated, negated core_dpe_oe indicates that the core is not driving a valid core_dpe value.
 If core_dpe_tre is asserted, negated core_dpe_oe indicates that core_dpe is in the high-impedance state.
- Timing Comments** Assertion—Asserted on the second bus clock cycle after core_ta is asserted to detect the incorrect parity, unless core_ta is canceled by an assertion of core_drtry or core_artry (in certain cases).
Negation—Occurs on the third bus clock cycle after core_ta is asserted.
 Note that the negation of core_dpe_oe may force core_dpe to the high-impedance state, if core_dpe_tre is asserted.

8.3.7.3.2 Data Parity Error High-Impedance Enable (core_dpe_tre)—Input

Following are the state meaning and timing comments for core_dpe_tre. core_dpe_tre is a high-impedance enable signal on the G2 core and can be used to create a three-statable version of core_dpe externally. The resulting core_dpe signal functions similar to a 60x bus signal when core_dpe_tre is asserted.

- State Meaning** Asserted—core_dpe_oe controls whether the data parity error output signal is driven or forced to a high-impedance state.
Negated—Indicates that core_dpe is always driven.
- Timing Comments** Assertion/Negation—Must be set up prior to negation of the core_hreset signal and remain stable during core operation. This is a static configuration.

8.3.7.4 Data Bus Disable (core_dbdis)—Input

The core_dbdis signal is an input signal (input-only) on the G2 core. Following are the state meaning and timing comments for the core_dbdis input.

State Meaning	<p><u>Asserted</u>—Indicates (for a write transaction) that the core must release the data bus and the data bus parity signals to high impedance <u>during the following cycle</u>. The data tenure remains active, <u>core_dbb_out</u> remains driven, and the transfer termination signals are still monitored by the core.</p> <p><u>Negated</u>—Indicates the data bus should remain normally driven. <u>core_dbdis</u> is ignored during read transactions.</p>
Timing Comments	<p><u>Assertion/Negation</u>—May be asserted on any clock cycle when the core is driving, or will be driving the data bus; may remain asserted for multiple cycles.</p>

8.3.8 Data Transfer Termination Signals

Data termination signals are required after each data beat in a data transfer. Note that in a single-beat transaction, the data termination signals also indicate the end of the tenure. While in burst accesses, the data termination signals apply to individual beats and indicate the end of the tenure only after the final data beat.

For a detailed description of how these signals interact, see Section 9.4.4, “Data Transfer Termination.”

8.3.8.1 Transfer Acknowledge (core_ta)—Input

Following are the state meaning and timing comments for the core_ta input.

State Meaning	<p><u>Asserted</u>—Indicates that a single-beat data transfer completed successfully or that a <u>data beat</u> in a burst transfer completed successfully (unless <u>core_drtry</u> is asserted on the next bus clock cycle).</p> <p>Note that <u>core_ta</u> must be asserted for each data beat in a burst transaction, and must be asserted during assertion of <u>core_drtry</u>. For more information, see Section 9.4.4, “Data Transfer Termination.”</p> <p><u>Negated</u>—(During assertion of <u>core_dbb_out</u>) indicates that, until <u>core_ta</u> is asserted, the core must continue to drive the data for the current write or must wait to sample the data for reads.</p>
Timing Comments	<p><u>Assertion</u>—Must not occur before <u>core_aack</u> for the current transaction (if the address retry mechanism is to be used to prevent invalid data from being used by the processor); <u>otherwise</u>, assertion may occur at any time during the <u>assertion</u> of <u>core_dbb_out</u>. The system can withhold assertion of <u>core_ta</u> to indicate that the core should insert wait states to extend the duration of the data beat.</p>

Negation—Must occur after the bus clock cycle of the final (or only) data beat of the transfer. For a burst transfer, the system can assert core_ta for one bus clock cycle and then negate it to advance the burst transfer to the next beat and insert wait states during the next beat. (Note: When the core is configured for 1:1 clock mode and is performing a burst read into the data cache, the core requires one wait state between the assertion of core_ts and the first assertion of core_ta for that transaction. If no-DRTRY mode is also selected, the core requires two wait states for 1:1 clock mode, or one wait state for 1.5:1 clock mode.)

8.3.8.2 Data Retry (core_drtry)—Input

Following are the state meaning and timing comments for the core_drtry input.

State Meaning	<p>Asserted—Indicates that the core must invalidate the data from the previous read operation.</p> <p>Negated—Indicates that data <u>presented with core_ta</u> on the previous read operation is valid. Note that <u>core_drtry</u> is ignored for write transactions.</p>
Timing Comments	<p><u>Assertion</u>—Must occur during the <u>bus clock cycle</u> immediately after <u>core_ta</u> is asserted if a retry is required. <u>core_drtry</u> may be held asserted for multiple bus clock cycles. When <u>core_drtry</u> is <u>negated</u>, data must have been valid on the previous clock with <u>core_ta</u> asserted.</p> <p><u>Negation</u>—Must occur during the <u>bus clock cycle</u> after a valid data beat. This may occur several cycles after <u>core_dbb_out</u> is negated, effectively extending the data bus tenure.</p> <p><u>Start-Up</u>—<u>core_drtrymode</u> is sampled at the negation of <u>core_hreset</u>; if <u>core_drtrymode</u> is asserted, <u>no-DRTRY mode</u> is selected. If <u>core_drtrymode</u> is negated at start-up, <u>core_drtry</u> is enabled.</p>

8.3.8.3 Transfer Error Acknowledge (core_tea)—Input

Following are the state meaning and timing comments for the core_tea input.

State Meaning	<p>Asserted—Indicates that a bus error occurred. Causes a machine check exception (and possibly causes the processor to enter checkstop state if machine check enable bit is cleared (MSR[ME] = 0)). For more information, see Section 5.5.2.2, “Checkstop State (MSR[ME] = 0).” <u>Assertion terminates</u> the current transaction; that is, assertion of <u>core_ta</u> and <u>core_drtry</u> are ignored.</p>
----------------------	---

The assertion of $\overline{\text{core_tea}}$ causes the negation of $\overline{\text{core_dbb_out}}$ in the next clock cycle. However, data entering the GPR or the cache is not invalidated.

Negated—Indicates that no bus error was detected.

Timing Comments Assertion—May be asserted while $\overline{\text{core_dbb_out}}$ is asserted, and the cycle after $\overline{\text{core_ta}}$ during a read operation. $\overline{\text{core_tea}}$ should be asserted for one cycle only.

Negation— $\overline{\text{core_tea}}$ must be negated no later than the negation of $\overline{\text{core_dbb_out}}$.

8.3.9 Interrupt and Checkstop Signals

Most interrupt and checkstop signals are input signals that indicate when exceptions are received, when checkstop conditions have occurred, and when the core must be reset. The G2 core generates the output signal $\overline{\text{core_ckstp_out}}$ when it detects a checkstop condition. For further detailed description of these signals, see Section 9.7, “Interrupt, Checkstop, and Reset Signals.”

8.3.9.1 External Interrupt ($\overline{\text{core_int}}$)—Input

Following are the state meaning and timing comments for the $\overline{\text{core_int}}$ input.

State Meaning Asserted—The core initiates an interrupt exception if MSR[EE] is set; otherwise, the core ignores the interrupt. To guarantee that the core takes the external interrupt, $\overline{\text{core_int}}$ must be held asserted until the core takes the interrupt.

Negated—Indicates that normal operation should proceed. See Section 9.7.1, “External Interrupts.”

Timing Comments Assertion—May occur at any time and may be asserted asynchronously to the input clocks. The $\overline{\text{core_int}}$ input is level-sensitive.

Negation—Should not occur until the external interrupt exception is taken.

8.3.9.2 Critical Interrupt ($\overline{\text{core_cint}}$)—Input: G2_LE Core-Only

Following are the state meaning and timing comments for the $\overline{\text{core_cint}}$ input on the G2_LE core. See Section 5.5.10, “Critical Interrupt Exception (0x00A00)—G2_LE Only,” for more information.

State Meaning Asserted—The core initiates an interrupt exception if MSR[CE] is set; otherwise, the core ignores the interrupt. To guarantee that the

core takes the critical interrupt, $\overline{\text{core_cint}}$ must be held asserted until the core takes the interrupt.

Negated—Indicates that normal operation should proceed. See Section 9.7.1, “External Interrupts.”

Timing Comments Assertion—May occur at any time and may be asserted asynchronously to the input clocks. The $\overline{\text{core_cint}}$ input is level-sensitive.

Negation—Should not occur until the critical interrupt exception is taken.

8.3.9.3 System Management Interrupt ($\overline{\text{core_smi}}$)—Input

Following are the state meaning and timing comments for the $\overline{\text{core_smi}}$ input. See Section 5.5.17, “System Management Interrupt (0x01400),” for more information.

State Meaning Asserted—The core initiates a system management interrupt exception if MSR[EE] is set; otherwise, the core ignores the exception condition. The system must hold $\overline{\text{core_smi}}$ asserted until the exception is taken.

Negated—Indicates that normal operation should proceed. See Section 9.7.1, “External Interrupts.”

Timing Comments Assertion—May occur at any time and may be asserted asynchronously to the input clocks. The $\overline{\text{core_smi}}$ input is level-sensitive.

Negation—Should not occur until the interrupt exception is taken.

8.3.9.4 Machine Check Interrupt ($\overline{\text{core_mcp}}$)—Input

Following are the state meaning and timing comments for the $\overline{\text{core_mcp}}$ input.

State Meaning Asserted—The core initiates a machine check interrupt exception if MSR[ME] and HID0[EMCP] are set; if MSR[ME] is cleared and HID0[EMCP] is set, the core terminates operation by internally gating off all clocks, and releasing all outputs (except $\overline{\text{core_ckstp_out}}$) to the high-impedance state. If HID0[EMCP] is cleared, the core ignores the interrupt condition. $\overline{\text{core_mcp}}$ must be held asserted for at least two bus clock cycles.

Negated—Indicates that normal operation should proceed. See Section 9.7.1, “External Interrupts.”

Timing Comments Assertion—May occur at any time and may be asserted asynchronously to the input clocks. $\overline{\text{core_mcp}}$ is negative edge-sensitive.

Negation—May be negated two bus cycles after assertion.

8.3.9.5 Checkstop Signals

There is both an checkstop input and checkstop output signal on the G2 core. The core also implements checkstop output enable and checkstop high-impedance enable signals.

8.3.9.5.1 Checkstop Input (core_ckstp_in)

Following are the state meaning and timing comments for core_ckstp_in.

State Meaning Asserted—Indicates that the core must terminate operation by internally gating off all clocks, and releasing all outputs (except core_ckstp_out) to the high-impedance state. Once core_ckstp_in is asserted, it must remain asserted until the system has been reset.

Negated—Indicates that normal operation should proceed. See Section 9.7.2, “Checkstops.”

Timing Comments Assertion—May occur at any time and may be asserted asynchronously to the input clocks.

Negation—May occur anytime after core_ckstp_out is asserted.

8.3.9.5.2 Checkstop Output (core_ckstp_out)

The core_ckstp_out signal is output only on the G2 core. The core also implements checkstop output enable and checkstop high-impedance enable signals. core_ckstp_out acts as follows:

- If core_ckstp_tre is asserted, the output is in one of the following three states—high impedance, driven high, or driven low.
- If core_ckstp_tre is negated, the output is either driven to the high or low state. In this case, a valid value on core_ckstp_out exists when core_ckstp_oe is asserted.

Following are the state meaning and timing comments for the core_ckstp_out output.

State Meaning Asserted—Indicates that the core has detected a checkstop condition and has ceased operation.

Negated—Indicates that the core is operating normally. See Section 9.7.2, “Checkstops.”

Timing Comments Assertion—May occur at any time and is asserted asynchronously to core_sysclk.

Negation—Is negated upon assertion of core_hreset.

High Impedance—Occurs after the negation of core_ckstp_oe, if core_ckstp_tre is asserted. If core_ckstp_tre is negated, core_ckstp_out is always driven.

8.3.9.5.3 Checkstop Output Enable (core_ckstp_oe)—Output

core_ckstp_oe is an output-enable indicator to its corresponding bus signals. Following are the state meaning and timing comments for core_ckstp_oe.

State Meaning Asserted—Indicates that the core is driving a valid $\overline{\text{core_ckstp_out}}$.

Negated—Indicates one of the following two conditions:

If core_ckstp_tre is negated, $\overline{\text{negated core_ckstp_oe}}$ indicates that the core is not driving a valid core_ckstp_out value.

If core_ckstp_tre is asserted, $\overline{\text{negated core_ckstp_oe}}$ indicates that core_ckstp_out is in the high-impedance state.

Timing Comments Assertion/Negation—core_ckstp_oe is valid after $\overline{\text{core_ckstp_out}}$ is asserted (asynchronous to core_sysclk).

Note that negation of core_ckstp_oe may force $\overline{\text{core_ckstp_out}}$ to the high-impedance state, if core_artry_tre is asserted.

8.3.9.5.4 Checkstop High-Impedance Enable (core_ckstp_tre)—Input

core_ckstp_tre is a high-impedance enable signal on the G2 core and can be used to create a bidirectional core_ckstp signal. When the related input/output signals ($\overline{\text{core_ckstp_in}}$ and $\overline{\text{core_ckstp_out}}$) are wire-ORed together, the resulting signal functions similar to a bidirectional 60x bus signal when core_ckstp_tre is asserted. See Section 8.2.2.2, “Logic Gate Equivalent and Bidirectional Signals,” for more information. Following are the state meaning and timing comments for core_ckstp_tre.

State Meaning Asserted—core_ckstp_oe controls whether $\overline{\text{core_ckstp_out}}$ is driven or forced to a high-impedance state.

Negated—Indicates that $\overline{\text{core_ckstp_out}}$ is always driven.

Timing Comments Assertion/Negation—Must be set up prior to negation of the $\overline{\text{core_hreset}}$ signal and remain stable during core operation. This is a static configuration.

8.3.10 Reset Signals

There are two reset signals on the G2 core—hard reset ($\overline{\text{core_hreset}}$) and soft reset ($\overline{\text{core_sreset}}$). Additionally, there is a group of reset configuration signals. Descriptions of the reset signals are as follows.

8.3.10.1 Hard Reset ($\overline{\text{core_hreset}}$)—Input

The $\overline{\text{core_hreset}}$ input must be used at power-on to properly reset the core. The reset configuration signals are sampled at the negation of $\overline{\text{core_hreset}}$. Following are the state meaning and timing comments for the $\overline{\text{core_hreset}}$ input.

State Meaning	<p>Asserted—Initiates a hard reset operation. Causes a reset exception as described in Section 5.5.1.1, “Hard Reset and Power-On Reset.” Output drivers are released to <u>high impedance</u> within five clock cycles after the assertion of <code>core_hreset</code>.</p> <p>Negated—Indicates that normal operation should proceed. See Section 9.7.3, “Reset Inputs.” The reset configuration signals are also sampled at the negation of <code>core_hreset</code>.</p>
Timing Comments	<p>Assertion—May occur at any time and may be asserted asynchronously to the core input clock; must be held asserted for a minimum of 255 clock cycles after the PLL lock time has been met. Refer to the appropriate hardware specifications for further timing comments.</p> <p>Negation—May occur any time after the minimum reset pulse width has been met.</p>

8.3.10.2 Soft Reset (`core_sreset`)—Input

The `core_sreset` signal is input only. Following are the state meaning and timing comments for the `core_sreset` input.

State Meaning	<p>Asserted—Initiates processing for a reset exception as described in Section 5.5.1.2, “Soft Reset.”</p> <p>Negated—Indicates that normal operation should proceed. See Section 9.7.3, “Reset Inputs.”</p>
Timing Comments	<p>Assertion—May occur at any time and <u>may be asserted</u> asynchronously to the core input clock. <code>core_sreset</code> is negative edge-sensitive.</p> <p>Negation—May be negated two bus cycles after assertion.</p>

8.3.10.3 Reset Configuration Signals

There are five reset configuration signals on the G2 core that are sampled at the negation of `core_hreset`.

8.3.10.3.1 32-Bit Mode (`core_32bitmode`)—Input

Following are the state meaning and timing comments for the `core_32bitmode` input.

State Meaning	<p>Asserted—Causes the core to be configured for 32-bit mode operation. See Section 9.6.1, “32-Bit Data Bus Mode,” for more information on the differences between 32- and 64-bit mode operation.</p>
----------------------	---

Negated—Causes the core to be configured for 64-bit mode operation.

Timing Comments Assertion/Negation—This signal is sampled during assertion of core_hreset and must be stable five cycles before the negation of core_hreset as defined in the hardware specification.

8.3.10.3.2 Reduced Pinout Mode (core_redpinmode)—Input

Following are the state meaning and timing comments for the core_redpinmode input.

State Meaning Asserted—Causes the processor to be configured for reduced pinout mode operation at core_hreset.

Negated—Causes the core to be configured for normal pinout mode operation at core_hreset.

Timing Comments Assertion/Negation—This signal is sampled during assertion of core_hreset and must be stable five cycles before the negation of core_hreset as defined in the hardware specification.

8.3.10.3.3 MSR IP Bit Set Mode (core_msrip)—Input

Following are the state meaning and timing comments for the core_msrip input.

State Meaning Asserted—Causes MSR[IP] to be initialized to a one at core_hreset. This causes the reset vector to be fetched from address 0xFFFFn_nnnn. See Table 2-4 for more information on the operation of the core when MSR[IP] = 1.

Negated—Causes MSR[IP] to be initialized to zero. This causes the reset vector to be fetched from address 0x000n_nnnn.

Timing Comments Assertion/Negation—This signal is sampled during assertion of core_hreset and must be stable five cycles before the negation of core_hreset as defined in the hardware specification.

8.3.10.3.4 DRTRY Mode (core_drtrymode)—Input

Following are the state meaning and timing comments for the core_drtrymode input.

State Meaning Asserted—Causes the core to be configured for normal DRTRY mode operation at core_hreset.

Negated—Causes the core to be configured for no_DRTRY mode operation at core_hreset.

Timing Comments Assertion/Negation—This signal is sampled during assertion of core_hreset and must be stable five cycles before the negation of core_hreset as defined in the hardware specification.

8.3.10.3.5 True Little-Endian Mode (`core_tle`)—Input

Following are the state meaning and timing comments for the `core_tle` input on the G2_LE core.

State Meaning Asserted—Causes `MSR[LE]`, `MSR[ILE]`, and `HID2[LET]` to be initialized to ones at `core_hreset`. See Table 2-8 for more information on the operation of the core when `HID2[LET] = 1`.

Negated—Causes `MSR[LE]`, `MSR[ILE]`, and `HID2[LET]` to be initialized to zeros at `core_hreset`.

Timing Comments Assertion/Negation—This signal is sampled during assertion of `core_hreset` and must be stable five cycles before the negation of `core_hreset` as defined in the hardware specification.

8.3.10.3.6 System Version Register (`core_svr[0:31]`)—Input

Following are the state meaning for the `core_svr[0:31]` inputs on the G2_LE core.

State Meaning Asserted/Negated—Identify the system version and revision level of the system on a chip (SOC) level of integration. The value of these signals is loaded into the system version register (SVR) at the negation of `core_hreset`. For further detailed description of the associated register, see Section 2.1.2.12, “System Version Register (SVR)—G2_LE Only.”

Timing Comments Assertion/Negation—These signals are sampled during assertion of `core_hreset` and must be stable five cycles before the negation of `core_hreset` as defined in the hardware specification.

8.3.11 Processor Status Signals

Processor status signals indicate the state of the processor. This includes the memory reservation, machine quiesce control, time base enable, and `core_tlbisync` signals.

8.3.11.1 Quiescent Acknowledge (`core_qack`)—Input

Following are the state meaning and timing comments for the `core_qack` input.

State Meaning Asserted—Indicates that all bus activity that requires snooping has terminated or paused, and that the core may enter the quiescent (or low-power) state.

Negated—Indicates that the core may not enter a quiescent state and must continue snooping the bus.

Timing Comments Assertion/Negation—May occur on any cycle following the assertion of `core_qreq`, and must be held asserted for a minimum of one bus clock cycle.

8.3.11.2 Quiescent Request (`core_qreq`)—Output

Following are the state meaning and timing comments for the `core_qreq` signal.

State Meaning Asserted—Indicates that the core is requesting all bus activity normally required to be snooped to terminate or to pause so the core may enter the quiescent (low-power) state. Once the core enters a quiescent state, it no longer snoops bus activity.

Negated—Indicates that the core is not making a request to enter the quiescent state.

Timing Comments Assertion/Negation—May assert on any cycle. `core_qreq` remains asserted for the duration of the quiescent state.

8.3.11.3 Reservation (`core_rsrv`)—Output

Following are the state meaning and timing comments for the `core_rsrv` output.

State Meaning Asserted/Negated—Represents the state of the reservation coherency bit in the reservation address register that is used by the `lwarx` and `stwcx` instructions. See Section 9.8.1, “Support for the `lwarx/stwcx` Instruction Pair.”

Timing Comments Assertion/Negation—Occurs synchronously with respect to bus clock cycles. The execution of an `lwarx` instruction sets the internal reservation condition.

8.3.11.4 Time Base Enable (`core_tben`)—Input

Following are the state meanings and timing comments for the `core_tben` input.

State Meaning Asserted—Indicates that the time base should continue clocking. This input is essentially a count enable control for the time base counter.

Negated—Indicates that the time base should stop clocking.

Timing Comments Assertion/Negation—May occur on any cycle.

8.3.11.5 TLBI Sync (`core_tlbisync`)—Input

Following are the state meaning and timing comments for the `core_tlbisync` input.

State Meaning Asserted—Indicates that instruction execution should stop after execution of a `tlbsync` instruction.

Negated—Indicates that the instruction execution may continue or resume after the completion of a **tlbsync** instruction.

Timing Comments Assertion/Negation—May occur on any cycle.

8.3.11.5.1 Output Enable (core_outputs_oe)—Output

For the G2 core, core_outputs_oe is associated with the $\overline{\text{core_qreq}}$, $\overline{\text{core_br}}$, $\overline{\text{core_rsrv}}$, and $\overline{\text{core_iabr}}$ signals; for the G2_LE core, core_outputs_oe is also associated with $\overline{\text{core_iabr2}}$, $\overline{\text{core_dabr}}$, and $\overline{\text{core_dabr2}}$. core_outputs_oe does not control any of these signals from the core. The signals listed above are always driven in normal operation. Following are the state meaning and timing comments for the core_outputs_oe output signal. Note that no high-impedance signal is associated with core_outputs_oe.

State Meaning Asserted—Indicates that the associated output signals are always driving valid data.

Negated—Indicates that the associated output signals are not driving valid data (does not occur in normal operation).

Timing Comments Assertion/Negation—In normal operation $\overline{\text{core_outputs_oe}}$ is asserted on the third clock cycle after $\overline{\text{core_hreset}}$ is negated.

8.3.12 COP/Scan Interface

The G2 core has extensive on-chip test capability including the following:

- Built-in instruction and data cache self-test (BIST)
- Debug control/observation (COP)
- Boundary scan (IEEE 1149.1 compliant interface)
- LSSD test control

The BIST hardware is not used as part of the power-on reset (POR) sequence. The COP and boundary scan logic are not used under typical operating conditions.

Detailed descriptions of the G2 core test functions is beyond the scope of this document; however, sufficient information has been provided to allow the system designer to disable the test functions that would impede normal operation.

The COP/scan interface is shown in Figure 8-4. For more information, see Section 9.9, “IEEE 1149.1-Compliant Interface.”

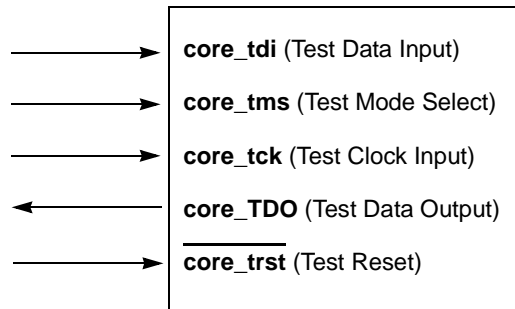


Figure 8-4. IEEE 1149.1-Compliant Boundary Scan Interface

These signals are not used during normal operation. **core_tms**, **core_tdi**, and **core_trst** have internal pull-up resistors provided; **core_tck** does not. For normal operation, **core_tms** and **core_tdi** may be left unconnected, and **core_tck** must be set high or low. The **core_trst** signal must be asserted sometime during power-up for JTAG logic initialization. Note that if **core_trst** is tied low, unnecessary power is consumed.

8.3.12.1 JTAG Test Clock (**core_tck**)—Input

The JTAG test clock (**core_tck**) signal is an input on the G2 core. Following are the state meaning and timing comments for the **core_tck** input signal.

State Meaning Asserted/Negated—This input should be driven by a free-running clock signal. Input signals to the test access port are clocked in on the rising edge of **core_tck**. Changes to the test access port output signals occur on the falling edge of **core_tck**. The test logic allows **core_tck** to be stopped.

Timing Comments Assertion/Negation—**core_tck** should not be used during normal operation and always must be set to either a high or low logic state.

8.3.12.2 JTAG Test Data Input (**core_tdi**)—Input

Following is the state meaning and timing comments for the **core_tdi** input signal.

State Meaning Asserted/Negated—The value presented on this signal on the rising edge of **core_tck** is clocked into the selected JTAG test instruction or data register.

Timing Comments Assertion/Negation—**core_tdi** should not be used during normal operation and always must be set to a high or low logic state. Note that this input contains an internal pull-up resistor to ensure that an unterminated input appears as a high signal level to the test logic.

8.3.12.3 JTAG Test Data Output (core_tdo)—Output

The JTAG test data output signal is an output on the G2 core. Following are the state meaning and timing comments for the core_tdo output signal.

State Meaning Asserted/Negated—The contents of the selected internal instruction or data register are shifted out onto this signal on the falling edge of core_tck. The core_tdo signal remains in a high-impedance state except when scanning of data is in progress.

Timing Comments Assertion/Negation—core_tdo should not be used for normal operation and is only valid when core_tdo_oe is asserted.

8.3.12.3.1 JTAG Test Data Output Enable (core_tdo_oe)—Output

The JTAG test data output enable signal is an output on the G2 core. Following are the state meaning and timing comments for the core_tdo_oe output signal.

State Meaning Asserted—Indicates that the G2 core is driving a valid $\overline{\text{core_tdo}}$ during the shiftDR or shiftID state of the TAP controller.

Negated—Indicates that the core is not driving a valid $\overline{\text{core_tdo}}$ value.

Timing Comments Assertion/Negation—The core_tdo signal is always driven, regardless of the state of core_tdo_oe. Also, core_tdo is always driven when core_lssd_mode is asserted and scanned.

8.3.12.4 JTAG Test Mode Select (core_tms)—Input

The test mode select (core_tms) signal is an input on the G2 core. Following are the state meaning for the core_tms input signal.

State Meaning Asserted/Negated—This signal is decoded by the internal JTAG TAP controller to distinguish the primary operation of the test support circuitry.

Timing Comments Assertion/Negation—core_tms should not be used during normal operation and always must be set to either a high or low logic state. Note that this input contains an internal pull-up resistor to ensure that an unterminated input appears as a high signal level to the test logic.

8.3.12.5 JTAG Test Reset (core_trst)—Input

The test reset ($\overline{\text{core_trst}}$) signal is an input on the G2 core. Following are the state meaning and timing comments for the core_trst input signal.

State Meaning Asserted—This input causes asynchronous initialization of the internal JTAG test access port controller. Note that the signal must be

asserted during the assertion of core_hreset in order to properly initialize the JTAG test access port. The core_trst signal must be asserted to properly initialize the boundary scan chain. This may be accomplished by connecting it to core_hreset, using logic to OR any external JTAG core_trst drivers.

Negated—Indicates normal operation.

Timing Comments Assertion/Negation—This input contains an internal pull-up resistor to ensure that an unterminated input appears as a high signal level (negated) to the test logic.

8.3.12.6 TLM TAP Enable (core_tap_en)—Input

The test linking module test access point enable (core_tap_en) signal is an input on the G2 core. Following are the state meaning and timing comments for the core_tap_en input signal.

State Meaning Asserted—Indicates that the test access point controller of the G2 core is in normal mode of operation which is controlled by core_tms.
Negated—Indicates that the core test access point enables the TLM function.

Timing Comments Assertion/Negation—This input signal should be either driven at low logic state during normal operation or not connected to TLM logic.

8.3.12.7 Test Linking Module Select (core_tlmselect)—Output

The test linking module select (core_tlmselect) signal is an output on the G2 core. Following are the state meaning and timing comments for core_tlmselect.

State Meaning Asserted—Indicates that the core test access point controller selects the TLM register by issuing a TLM instruction.
Negated—Indicates normal operation.

Timing Comments Assertion/Negation—The test access point controller transitions to run/test/ideal state until the core_tlmselect signal is re-enabled.

8.3.13 Test Interface

Test interface signals like LSSD test clock or test control signals are used in the G2 core for production testing. core_l1_tstclk, core_l2_tstclk, and core_lssd_mode are the test clock and test control signals.

8.3.13.1 Disable (core_disable)—Input

The disable (core_disable) signal is an input on the G2 core. Following are the state meaning and timing comments for core_disable.

State Meaning Asserted—All output signals are negated or forced to a high-impedance state. The core enters a sleep mode, and instruction fetching and dispatching are disabled.

Negated—The G2 core is in normal operating mode.

Timing Comments Assertion/Negation—The core_disable signal should be asserted or negated when core_hreset is asserted and should remain asserted or negated until core_hreset is negated.

8.3.13.2 LSSD Test Clock (core_l1_tstclk, core_l2_tstclk)—Input

The LSSD test clock signals are inputs on the G2 core. Following are the state meaning and timing comments for the core_l1_tstclk and core_l2_tstclk input signals.

State Meaning Asserted—Indicates the high phase of the test clock.

Negated—Indicates the low phase of the test clock.

Timing Comments Assertion/Negation—core_l1_tstclk or core_l2_tstclk are driven during normal operating mode and clocked during LSSD test mode.

8.3.13.3 LSSD Test Control (core_lssd_mode)—Input

The LSSD test control (core_lssd_mode) signal is an input on the G2 core. Following are the state meaning and timing comments for the core_lssd_mode input signal.

State Meaning Asserted—Indicates that the core is in LSSD mode for manufacturing tests where core_pll_cfg[0:4] is set to 0x00011 to bypass the core_sysclk. In LSSD mode core_l1_tstclk and core_l2_tstclk control clocking instead of core_sysclk.

Negated—The G2 core is in normal operating mode. In normal operating mode core_l1_tstclk and core_l2_tstclk are tied to high state. The setting of core_pll_cfg[0:4] is changed through the setting of core-bus frequency ratio where the core clock frequency is the multiple of core_sysclk frequency.

Timing Comments Assertion/Negation—The system must negate core_lssd_mode and must keep it stable in normal operation.

8.3.14 Debug Control Signals

This section describes the signals that are implemented to control debug features such as address matching, combinational matching, and watchpoint of the PowerPC architecture

with respect to the G2 and G2_LE cores. The control signals— $\overline{\text{core_iabr}}$, $\overline{\text{core_iabr2}}$, $\overline{\text{core_dabr}}$, and $\overline{\text{core_dabr2}}$ are watchpoint/breakpoint indicator signals.

8.3.14.1 Instruction Address Breakpoint Register Watchpoint (core_iabr)—Output

The instruction address breakpoint register ($\overline{\text{core_iabr}}$) signal is an output on the G2 core. See Section 2.1.2.14, “Instruction Address Breakpoint Registers (IABR and IABR2),” for more information. Following are the state meaning and timing comments for the $\overline{\text{core_iabr}}$ input signal.

State Meaning Asserted—Indicates that the IABR register has matched with the instruction address breakpoint condition set in the IBCR. See Section 2.1.2.14.1, “Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only,” for more information.

Negated—Indicates that IABR has not matched or IBCR has disabled the breakpoint.

Timing Comments Assertion/Negation—Occurs synchronously with respect to bus clock cycles.

8.3.14.2 Instruction Address Breakpoint Register Watchpoint (core_iabr2)—Output

The instruction address breakpoint register ($\overline{\text{core_iabr2}}$) signal is an output on the G2_LE core. See Section 2.1.2.14, “Instruction Address Breakpoint Registers (IABR and IABR2),” for more information. Following are the state meaning and timing comments for the $\overline{\text{core_iabr2}}$ input signal.

State Meaning Asserted—Indicates that the IABR2 register has matched with the instruction address breakpoint condition set in the IBCR. See Section 2.1.2.14.1, “Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only,” for more information.

Negation—Indicates that IABR2 has not matched or IBCR has disabled the breakpoint.

Timing Comments Assertion/Negation—Occurs synchronously with respect to bus clock cycles.

8.3.14.3 Data Address Breakpoint Register Watchpoint (core_dabr)—Output

The data address breakpoint register ($\overline{\text{core_dabr}}$) signal is an output on the G2_LE core. See Section 2.1.2.15, “Data Address Breakpoint Register (DABR and DABR2)—G2_LE Only,” for more information. Following is the state meaning and timing comments for the $\overline{\text{core_dabr}}$ input signal.

State Meaning	<p>Asserted—Indicates that the DABR register has matched with the data address breakpoint condition set in the DBCR. See Section 2.1.2.15.1, “Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only,” for more information.</p> <p>Negation—Indicates that DABR has not matched or DBCR has disabled the breakpoint.</p>
Timing Comments	<p>Assertion/Negation—Occurs synchronously with respect to bus clock cycles.</p>

8.3.14.4 Data Address Breakpoint Register Watchpoint (core_dabr2)—Output

The data address breakpoint register ($\overline{\text{core_dabr2}}$) signal is an output on the G2_LE core. See Section 2.1.2.15, “Data Address Breakpoint Register (DABR and DABR2)—G2_LE Only,” for more information. Following is the state meaning and timing comments for the $\overline{\text{core_dabr2}}$ input signal.

State Meaning	<p>Asserted—Indicates that the DABR2 register has matched the data breakpoint condition set in the DBCR. See Section 2.1.2.15.1, “Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only,” for more information.</p> <p>Negation—Indicates that DABR2 has not matched or DBCR has disabled the breakpoint.</p>
Timing Comments	<p>Assertion/Negation—Occurs synchronously with respect to bus clock cycles.</p>

8.3.15 Clock Signals

The clock signal inputs of the G2 core determine the system clock frequency and provide a flexible clocking scheme that allow the processor to operate at an integer multiple of the system clock frequency.

Refer to the appropriate hardware specifications for exact timing relationships of the clock signals.

8.3.15.1 System Clock (core_sysclk)—Input

The core requires a single system clock (core_sysclk) input. This input sets the frequency of operation for the bus interface. Internally, the core uses a phase-locked loop (PLL) circuit to generate a master clock for all of the CPU circuitry (including the bus interface circuitry) which is phase-locked to the core_sysclk input. The master clock may be set to an integer or half-integer multiple of the SYSCLK frequency allowing the CPU core to operate at an

equal or greater frequency than the bus interface. The hardware specification lists available frequency multipliers.

State Meaning Asserted/Negated—The `core_sysclk` input is the primary clock input for the core, and represents the bus clock frequency for `core_sysclk` bus operation. Internally, the core may be operating at an integer or half-integer multiple of the bus clock frequency.

Timing Comments Duty cycle—Refer to the appropriate hardware specifications for timing comments.

Note: `core_sysclk` is used as the frequency reference for the internal PLL clock generator, and must not be suspended or varied during normal operation to ensure proper PLL operation.

8.3.15.2 Test Clock Output (`core_clk_out`)

The G2 core provides the `core_clk_out` signal for test purposes. It allows the monitoring of the processor and bus clock frequencies. The frequency of `core_clk_out` is determined by the configuration of `HID0[SBCLK,ECLK]`, as shown in Table 8-12. Note that `core_clk_out` is driven at the processor frequency during the assertion of `core_hreset`; when `core_hreset` is negated, `core_clk_out` enters the default high-impedance state.

Table 8-12. `core_clk_out` Signal Configuration

<code>HID0[SBCLK]</code>	<code>HID0[ECLK]</code>	<code>core_clk_out</code> State
0	0	Bus clock frequency
0	1	Core/processor clock frequency
1	0	Bus clock frequency
1	1	Core/processor clock frequency

Following are the state meaning and timing comments for `core_clk_out`.

State Meaning Asserted/Negated—Provides PLL clock output for PLL testing and monitoring. The `core_clk_out` signal clocks at either the processor clock frequency, bus clock frequency, or half-bus clock frequency if enabled by the appropriate `HID0` bits; the default state of `core_clk_out` is high impedance. `core_clk_out` is provided only for testing.

Timing Comments Assertion/Negation—Refer to the appropriate hardware specifications for timing comments.

8.3.15.3 PLL Configuration (core_pll_cfg[0:4])—Input

The PLL is configured by core_pll_cfg[0:4]. For a given core_sysclk (bus) frequency, the PLL configuration signals set the internal CPU frequency of operation. Table 8-13 shows the PLL configuration options.

Following are the state meaning and timing comments for the core_pll_cfg[0:4] input.

State Meaning Asserted/Negated— Configures the operation of the PLL and the internal processor clock frequency. Settings are based on the desired bus, VCO divider, and internal frequency of operation.

Timing Comments Assertion/Negation—Must remain stable during operation; should only be changed during the assertion of core_hreset or during sleep mode.

Table 8-13. Core PLL Configuration

pll_cfg[0:4]	Bus-to-Core Multiplier	VCO Divider
0x02	1x	8
0x01	1x	4
0x0C	1.5x	8
0x00	1.5x	4
0x18	1.5x	2
0x05	2x	4
0x04	2x	2
0x11	2.5x	4
0x06	2.5x	2
0x10	3x	4
0x08	3x	2
0x0E	3.5x	2
0x0A	4x	2
0x07	4.5x	2
0x0B	5x	2
0x09	5.5x	2
0x0D	6x	2
0x12,	6.5x	2
0x14	7x	2
0x16	7.5x	2
0x1C	8x	2

Table 8-13. Core PLL Configuration (continued)

pll_cfg[0:4]	Bus-to-Core Multiplier	VCO Divider
0x03, 0x13	PLL off or bypassed ¹	PLL off—SYSCLK drives core clocks directly, 1x bus-to-core defaulted
0x0F, 0x1F	PLL off ¹	PLL off—no core clocking occurs
0x15, 0x17, 0x19, 0x1A, 0x1B, 0x1D, 0x1E	Reserved	These decodings are reserved for future use and should not be used.

¹ When PLL off or bypassed, the AC timing for the core interface is undefined.

See hardware specification for more details on setup and hold time.

Chapter 9

Core Interface Operation

This chapter describes the 60x bus interface of the G2 core and its operation. It shows how the core signals, defined in Chapter 8, “Signal Descriptions,” interact to perform address and data transfers. For a detailed discussion about the 60x bus interface, multiple bus masters, and memory coherency, refer to the *PowerPC Microprocessor Family: The Bus Interface for 32-Bit Microprocessors*.

9.1 Overview

The core interface prioritizes requests for bus operations from the instruction and data caches and performs bus operations following the 60x bus protocol. It includes address register queues, prioritization logic, and the bus control unit. The core interface latches snoop addresses for snooping in the data cache and address register queues, snoops for direct-store reply operations and reservations controlled by the Load Word and Reserve Indexed (**lwarx**) and Store Word Conditional Indexed (**stwcx.**) instructions, and maintains the touch load address for the data cache. The interface allows one level of pipelining; that is, with certain restrictions described in subsequent sections, there can be as many as two outstanding transactions at any given time. Accesses are prioritized with load operations preceding store operations.

Instructions are automatically fetched from the memory system into the instruction unit where they are dispatched to the execution units or forwarded to the branch processing unit at a peak rate of three instructions per clock (see Section 7.3, “Timing Considerations”). Conversely, load and store instructions explicitly specify the movement of operands to and from the general-purpose and floating-point registers (GPRs and FPRs) and the memory system.

When the G2 core encounters an instruction or data access, it calculates the logical address (effective address) and uses the low-order address bits to check for a hit in the on-chip, 16-Kbyte instruction or data caches. During cache lookup, the instruction and data memory management units (MMUs) use the higher-order address bits to calculate the virtual address, allowing them to calculate the physical address (real address). The physical address bits are then compared with the corresponding cache tag bits to determine if a cache hit occurred. If the access misses in the corresponding cache, the physical address is used to access system memory.

In addition to loads, stores, and instruction fetches, the core performs software table search operations following TLB misses, cache cast-out operations when least recently used (LRU) cache lines are written to memory after a cache miss, and cache-line snoop push-out operations when a modified cache line experiences a snoop hit from another bus master.

Figure 9-1 shows the address path from the execution units and instruction fetcher, through the translation logic to the caches and system interface logic.

The core uses separate address and data buses and a variety of control and status signals for performing reads and writes. The address bus is 32 bits wide and the data bus can be configured to be 32 or 64 bits wide on reset. The interface is synchronous—all core inputs are sampled at and all outputs are driven from the rising edge of the bus clock. The bus can run at the full processor-clock frequency or at an integer division of the processor-clock speed. The implementation of the internal voltage of the G2 core is process dependent; all I/O signals for the device depends on the system level requirement. Note that the G2 core has no direct external I/O connection.

9.1.1 Operation of the Instruction and Data Caches

The G2 core contains independent instruction and data caches. Each cache is a physically-addressed, 16-Kbyte cache with four-way set-associativity. Both caches consist of 128 sets of four 8-word cache lines.

Because the on-chip data cache is a write-back primary cache, the predominant type of transaction is burst-read memory operations, followed by burst-write memory operations, and single-beat (noncacheable or write-through) memory read and write operations. Additionally, there can be address-only operations, variants of the burst and single-beat operations (such as, global memory operations that are snooped and atomic memory operations), and address retry activity (such as, when a snooped read access hits a modified line in the cache).

Because data cache tags are single-ported, simultaneous load or store and snoop accesses cause resource contention. Snoop accesses have the highest priority and are given first access to the tags, unless the snoop access coincides with a tag write; in this case, the snoop is retried and must re-arbitrate for cache access. Loads or stores deferred due to snoop accesses are performed during the clock cycle following the snoop.

The core supports a three-state coherency protocol (MEI) that is a subset of the MESI (modified/exclusive/ shared/invalid) four-state protocol and operates coherently in systems that contain four-state caches. With the exception of the **dcbz** instruction, the core does not broadcast cache control instructions. The cache control instructions are intended for the management of the local cache but not for other caches in the system.

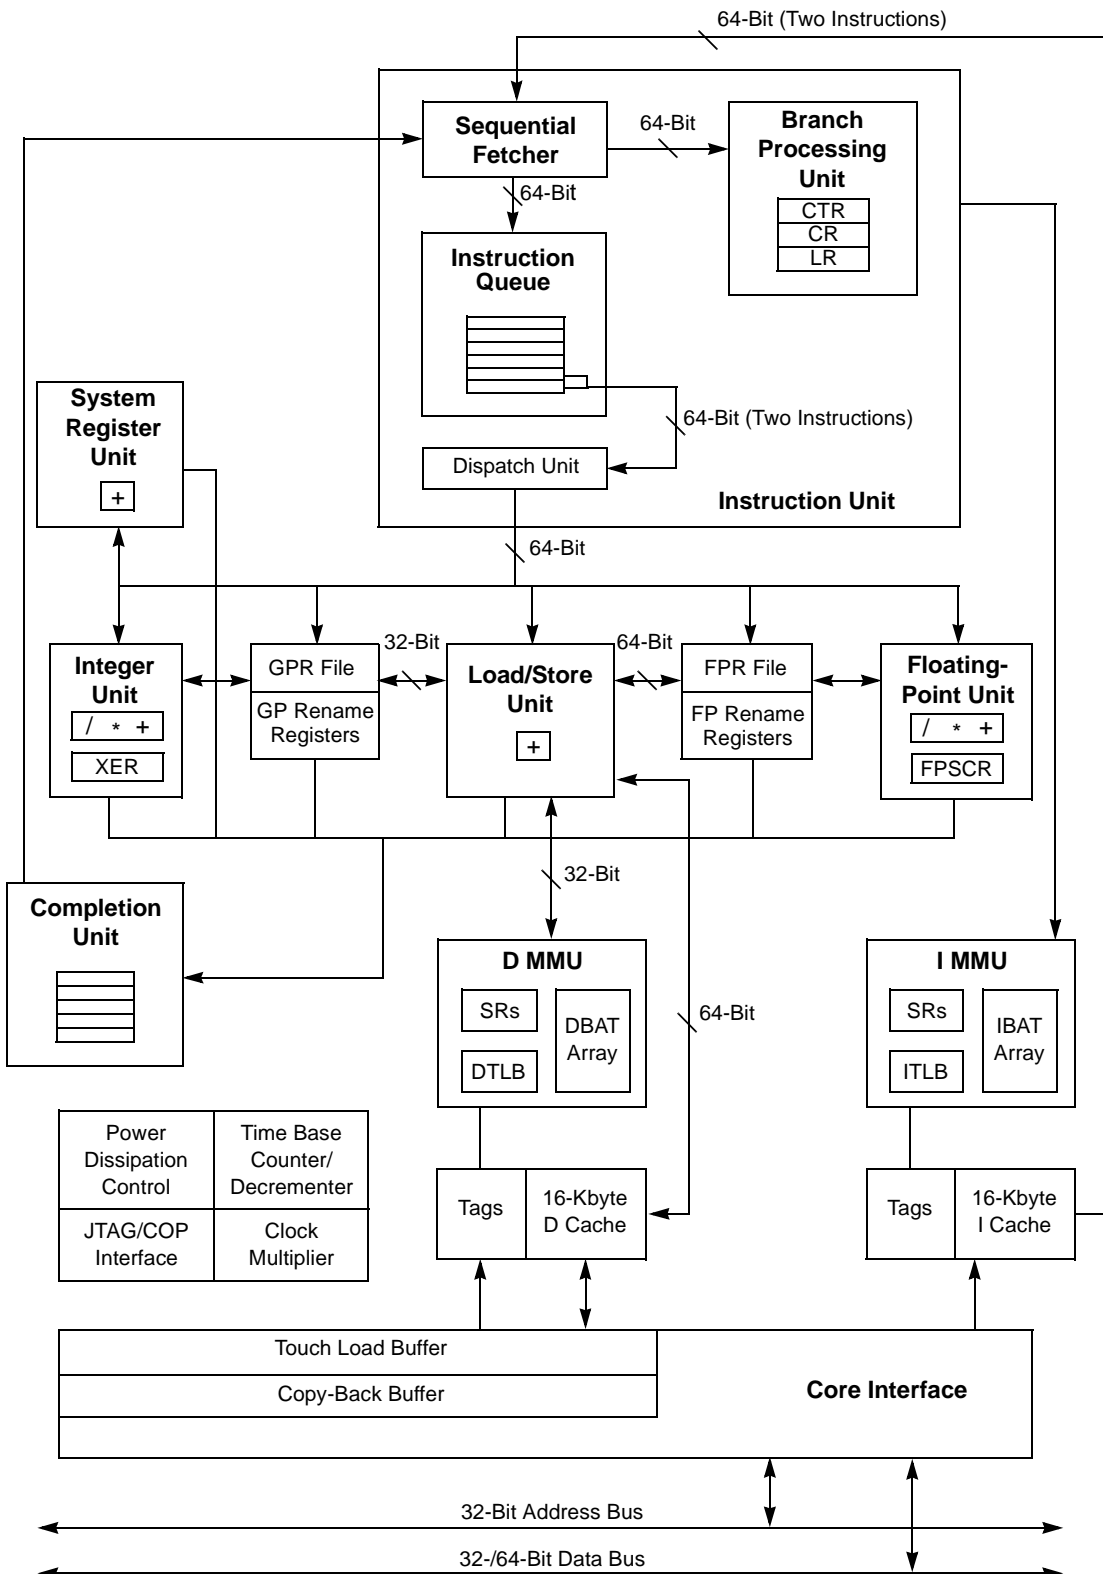


Figure 9-1. G2 Core Block Diagram

Cache lines in the core are loaded in four beats of 64 bits each (or eight 32-bit beats when operating in 32-bit bus mode). The burst load is performed as a critical-double-word-first operation. The cache that is being loaded is blocked to internal accesses until the load completes (that is, no hits under misses). The critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, minimizing stalls due to load delays.

Cache lines are selected for replacement based on an LRU algorithm. Each time a cache line is accessed, it is tagged as the most recently used line of the set. When a miss occurs, if all lines in the set are marked as valid, the LRU line is replaced with the new data. When data to be replaced is in the modified state, the modified data is written into a write-back buffer while the missed data is being read from memory. When the load completes, the core pushes the replaced line from the write-back buffer to main memory in a burst write operation.

9.1.2 Operation of the System Interface

Memory accesses can occur in single-beat (1 to 8 bytes) and four-beat (32 bytes) burst data transfers when the core is configured with a 64-bit data bus (`core_32bitmode` signal is negated at reset). When the core is in the optional 32-bit data bus mode (`core_32bitmode` signal is asserted at reset), memory accesses can occur in single-beat (1 to 4 bytes), two-beat (8 bytes), and eight-beat (32 bytes) bursts. The address and data buses are independent for memory accesses to support pipelining and split transactions. The core can pipeline as many as two transactions and has limited support for out-of-order split-bus transactions.

Access to the 60x bus interface is granted through an arbitration mechanism external to the core that allows devices to compete for bus mastership. This arbitration mechanism is flexible, allowing the core to be integrated into systems that implement various fairness and bus-parking procedures to avoid arbitration overhead.

Typically, memory accesses are weakly ordered—sequences of operations, including load/store string and multiple instructions, do not necessarily complete in the order they begin—maximizing the bus efficiency without sacrificing data coherency. The core allows load operations to precede store operations (except when a dependency exists). In addition, the core can be configured to reorder high-priority store operations ahead of lower-priority store operations. Because the processor can dynamically optimize run-time ordering of load/store traffic, overall performance is improved.

Note that the Synchronize (**sync**) instruction can be used to enforce strong ordering.

The following sections describe how the G2 core interface operates, providing detailed timing diagrams that show how the signals interact. A collection of more general timing diagrams are included as examples of typical bus operations.

9.1.3 Optional 32-Bit Data Bus Mode

The G2 core supports an optional 32-bit data bus mode, which differs from the 64-bit data bus mode only in the byte lanes involved in data transfers and the number of data beats performed. A data tenure in the 32-bit data bus mode takes one, two, or eight beats depending on the transfer size and the cache mode for the address. For additional information, see Section 9.6.1, “32-Bit Data Bus Mode.”

9.1.4 Direct-Store Accesses

The G2 core does not support the extended transfer protocol for accesses to the direct-store storage space. If SR[T] is set, the memory access is a direct-store access. An attempt to access to a direct-store segment results in a DSI exception.

9.2 Memory Access Protocol

Figure 9-2 shows that the address and data tenures are distinct from one another and that both consist of three phases—arbitration, transfer, and termination. Address and data tenures are independent (indicated in Figure 9-2 by the fact that the data tenure begins before the address tenure ends), which allows split-bus transactions to be implemented at the system level in multiprocessor systems. Figure 9-2 shows a data transfer that consists of a single-beat transfer of as many as 64 bits. Four-beat burst transfers of 32-byte cache lines require data transfer termination signals for each beat of data.

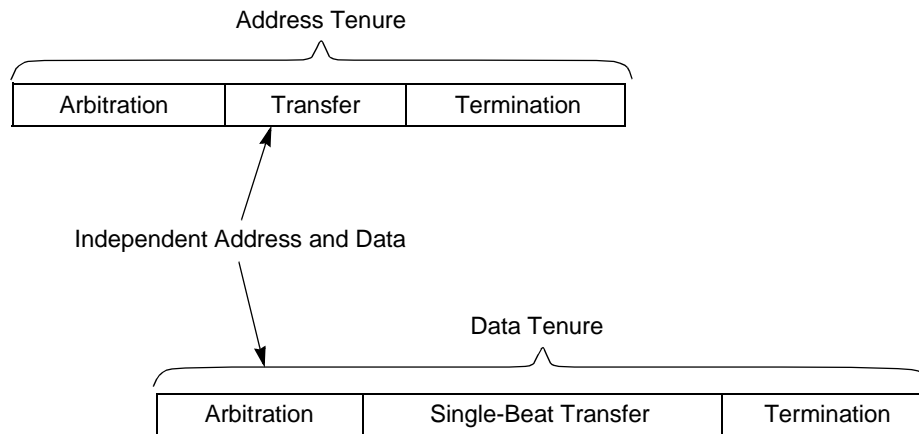


Figure 9-2. Overlapping Tenures on the Bus for a Single-Beat Transfer

The basic functions of the address and data tenures are as follows:

- Address tenure
 - Arbitration: During arbitration, address bus arbitration signals are used to gain address bus mastership.
 - Transfer: After the core is the address bus master, it transfers the address on the address bus. The address signals and the transfer attribute signals control the address transfer. The address parity and address parity error signals ensure the integrity of the address transfer.
 - Termination: After the address transfer, the system signals that the address tenure is completed or that it must be repeated.
- Data tenure
 - Arbitration: To begin the data tenure, the core arbitrates for data bus mastership.
 - Transfer: After the core is the data bus master, it samples the data bus for read operations or drives the data bus for write operations. The data parity and data parity error signals ensure the integrity of the transfer.
 - Termination: Data termination signals are required after each beat. Note that in a single-beat transaction, the data termination signals also indicate the end of the tenure, while in burst accesses, the data termination signals apply to individual beats and indicate the end of the tenure only after the final data beat.

The core generates an address-only bus transfer during the execution of the **dcbz** instruction, and uses only the address bus with no data transfer involved. Additionally, the core retry capability provides an efficient snooping protocol for systems with multiple memory systems (including caches) that must remain coherent.

9.2.1 Arbitration Signals

Arbitration for both address and data bus mastership is performed by a central, external arbiter and, minimally, by the arbitration signals shown in Section 8.3.1, “Address Bus Arbitration Signals.” Most arbiter implementations require additional signals to coordinate bus master/slave/snooping activities. Note that two arbitration signals—address bus busy (`core_abbx`) and data bus busy (`core_dbbx`) are both inputs and outputs on the G2 core. These signals are inputs unless the MPC603e has mastership of one or both of the respective buses; they must be connected high through pull-up resistors so that they remain negated when no devices have control of the buses.

The following list describes the address arbitration signals:

- `core_br` (bus request)—Assertion indicates that the core is requesting mastership of the address bus.
- `core_bg` (bus grant)—Assertion indicates that the core may, with the proper qualification, assume mastership of the address bus. A qualified bus grant occurs

when $\overline{\text{core_bg}}$ is asserted and $\overline{\text{core_abb_in}}$ and $\overline{\text{core_artry_in}}$ (after $\overline{\text{core_aack}}$) are not asserted.

If the G2 core is parked, $\overline{\text{core_br}}$ need not be asserted for the qualified bus grant.

- $\overline{\text{core_abb_out}}$ (address bus busy)—Assertion by the core indicates that the core is the address bus master.

The following list describes the data arbitration signals:

- $\overline{\text{core_dbg}}$ (data bus grant)—Indicates that the core may, with the proper qualification, assume mastership of the data bus. A qualified data bus grant occurs when $\overline{\text{core_dbg}}$ is asserted while $\overline{\text{core_dbb_in}}$, $\overline{\text{core_drtry}}$, and $\overline{\text{core_artry_in}}$ are negated; that is, the data bus is not busy ($\overline{\text{core_dbb_in}}$ is negated), there is no outstanding attempt to retry the current data tenure ($\overline{\text{core_drtry}}$ is negated), and there is no outstanding attempt to perform an $\overline{\text{core_artry_in}}$ of the associated address tenure.

$\overline{\text{core_dbbx}}$ is driven by the current bus master, $\overline{\text{core_drtry}}$ is driven only by the system, and $\overline{\text{core_artry}}$ is driven from the bus, but only for the address tenure associated with the current data tenure (that is, not from another address tenure).

- $\overline{\text{core_dbwo}}$ (data bus write only)—Assertion indicates that the core may perform the data bus tenure for an outstanding write address even if a read address is pipelined before the write address. If $\overline{\text{core_dbwo}}$ is asserted, the core assumes data bus mastership for a pending data bus write operation; the core takes the data bus for a pending read operation if this input is asserted along with $\overline{\text{core_dbg}}$ and no write is pending. Care must be taken with $\overline{\text{core_dbwo}}$ to ensure the desired write is queued (for example, a cache-line snoop push-out operation).
- $\overline{\text{core_dbb_out}}$ (data bus busy)—Assertion by the core indicates that the core is the data bus master. The core always assumes data bus mastership if it needs the bus and is given a qualified data bus grant (see $\overline{\text{core_dbg}}$).

For more detailed information on the arbitration signals, refer to Section 8.3.1, “Address Bus Arbitration Signals,” and Section 8.3.6, “Data Bus Arbitration Signals.”

9.2.2 Address Pipelining and Split-Bus Transactions

The 60x bus protocol provides independent address and data bus capability to support pipelined and split-bus transaction system organizations. Address pipelining allows the address tenure of a new bus transaction to begin before the data tenure of the current transaction has finished. Split-bus transaction capability allows other bus activity to occur (either from the same master or from different masters) between the address and data tenures of a transaction.

While this capability does not inherently reduce memory latency, support for address pipelining and split-bus transactions can greatly improve effective bus/memory throughput. For this reason, these techniques are most effective in shared-memory multiprocessor

implementations where bus bandwidth is an important measurement of system performance.

External arbitration is required in systems in which multiple devices must compete for the system bus. The design of the external arbiter affects pipelining by regulating address bus grant (core_bg), data bus grant (core_dbg), and address acknowledge (core_aack) signals. For example, a one-level pipeline is enabled by asserting core_aack to the current address bus master and granting mastership of the address bus to the next requesting master before the current data bus tenure has completed. Two address tenures can occur before the current data bus tenure completes.

The core can pipeline its own transactions to a depth of one level (intraprocessor pipelining); however, the 60x bus protocol does not constrain the maximum number of levels of pipelining that can occur on the bus between multiple masters (interprocessor pipelining). The external arbiter must control the pipeline depth and synchronization between masters and slaves.

In a pipelined implementation, data bus tenures are kept in strict order with respect to address tenures. However, external hardware can further decouple the address and data buses, allowing the data tenures to occur out of order with respect to the address tenures. This requires some form of system tag to associate the out-of-order data transaction with the proper originating address transaction (not defined for the G2 core interface). Individual bus requests and data bus grants from each processor can be used by the system to implement tags to support interprocessor, out-of-order transactions.

The G2 core supports a limited intraprocessor out-of-order, split-transaction capability via the data bus write only (core_dbwo) signal. For more information concerning the use of core_dbwo, see Section 9.10, “Using core_dbwo (Data Bus Write Only).”

9.2.3 Timing Diagram Conventions

Table 9-1 shows the conventions used in the timing diagrams.

This is a synchronous interface—all core input signals are sampled and output signals are driven on the rising edge of the bus clock cycle.

Table 9-1. Timing Diagram Legend

Feature	Example	Description
Grey	<u>core_artry_in</u>	Core input while the core is the bus master
Bold overbar	<u>core_br</u>	Core output while the core is the bus master
Plain	Data	Core input or output while the core is the bus master
+	ADDR+	Core output (grouped: here, address plus attributes)
Plain overbar	<u>qual_bg</u>	Internal core signal inaccessible to the user, but used in diagrams to clarify operations

Table 9-1. Timing Diagram Legend (continued)

Feature	Example	Description
Curled arrow		Dependency
Zig-zag		Indication that some clocks may have been skipped
Unshaded		A valid output or input signal or bus that can be in any of the possible states indicated
Shaded		Core nonsampled input or indeterminately driven output among the possible states indicated.
Dot		Signal with sample point
Dot on dotted vertical line		A sampled condition (dot on high or low state) with multiple dependencies
Dotted signal		Timing for a signal had it been asserted

9.3 Address Bus Tenure

This section describes the address bus arbitration, transfer, and termination phases.

9.3.1 Address Bus Arbitration

When the core needs access to the 60x bus and it is not parked ($\overline{\text{core_bg}}$ is negated), it asserts bus request (core_br) until it is granted mastership of the bus and the bus is available (see Figure 9-3). The external arbiter must grant master-elect status to the potential master by asserting the bus grant ($\overline{\text{core_bg}}$) signal. The core requesting the bus determines that the bus is available when the core_abb_in signal is negated. When the address bus is not busy ($\overline{\text{core_abb_in}}$ is negated), $\overline{\text{core_bg}}$ is asserted, and the address retry ($\overline{\text{core_artry}}$) input is negated. This is referred to as a qualified bus grant. The core assumes address bus mastership by asserting core_abb_out when it receives a qualified bus grant.

External arbiters must allow only one device at a time to be the address bus master. In implementations where no other device can be a master, $\overline{\text{core_bg}}$ can be grounded (always asserted) to continually grant mastership of the address bus to the core.

If the core asserts core_br before the external arbiter asserts $\overline{\text{core_bg}}$, the core is considered to be unparked, as shown in Figure 9-3. Figure 9-4 shows the parked case, where a qualified bus grant exists on the clock edge following a need_{bus} condition. Note that the bus clock cycle required for arbitration is eliminated if the core is parked, reducing overall memory

latency for a transaction. The core always negates $\overline{\text{core_abb_out}}$ for at least one bus clock cycle after core_aack is asserted, even if it is parked and has another transaction pending.

Typically, bus parking is provided to the device that was the most recent bus master. However, system designers may choose other schemes, such as providing unrequested bus grants in situations where it is easy to correctly predict the next device requesting bus mastership.

When the core receives a qualified bus grant, it assumes address bus mastership by asserting $\overline{\text{core_abb_out}}$ and negating core_br . Meanwhile, the core drives the address for the requested access onto the address bus and asserts core_ts_out to indicate the start of a new transaction.

Note that the core may assert $\overline{\text{core_br}}$ without using the bus after it receives the qualified bus grant when external bus arbitration logic is designed. For example, in a system using bus snooping, if the core asserts $\overline{\text{core_br}}$ to perform a replacement copy-back operation, another device can invalidate that line before the core is granted mastership of the bus. In that case, once the core is granted the bus, it no longer needs to perform the copy-back operation; therefore, the core does not assert $\overline{\text{core_abb_out}}$ and does not use the bus for the copy-back operation. Note that the core asserts $\overline{\text{core_br}}$ for at least 1 clock cycle in these instances.

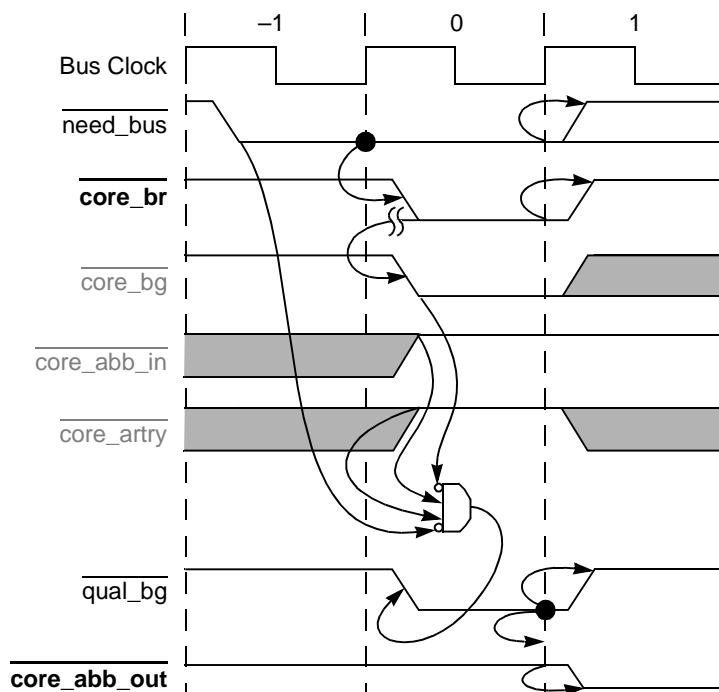


Figure 9-3. Address Bus Arbitration

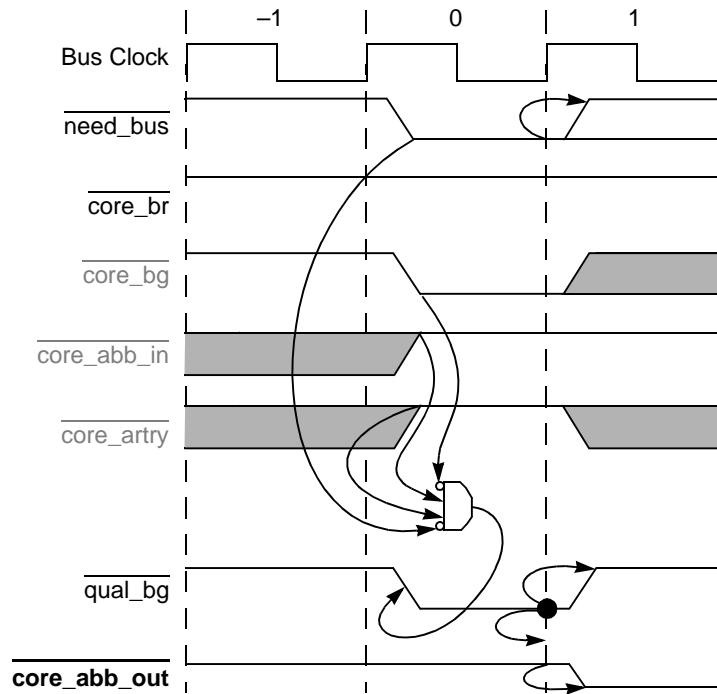


Figure 9-4. Address Bus Arbitration Showing Bus Parking

9.3.2 Address Transfer

During the address transfer, the physical address and all attributes of the transaction are transferred from the bus master to slave devices. Snooping logic may monitor the transfer to enforce cache coherency; see description of bus snooping in Section 9.3.3, “Address Transfer Termination.”

The signals used in the address transfer include the following signal groups:

- Address transfer start signal: Transfer start ($\overline{\text{core_ts_out}}$)
- Address transfer signals: Address bus ($\text{core_a_out}[0:31]$), address parity ($\text{core_ap_out}[0:3]$), and address parity error (core_ape).
- Address transfer attribute signals: Transfer type ($\text{core_tt_out}[0:4]$), transfer code ($\text{core_tc}[0:1]$), transfer size ($\text{core_tsiz}[0:2]$), transfer burst (core_tbst), cache inhibit (core_ci), write-through (core_wt), global (core_gbl_out), and cache set element ($\text{core_cse}[0:1]$).

Figure 9-5 shows that the timing for all of these signals, except $\overline{\text{core_ts_out}}$ and core_ape , is identical. All of the address transfer and address transfer attribute signals are combined into the ADDR+ grouping in Figure 9-5. The $\overline{\text{core_ts_out}}$ signal indicates that the core has begun an address transfer and that the address and transfer attributes are valid (within the context of a synchronous bus). The core always asserts $\overline{\text{core_ts_out}}$ coincident with core_abb_out . As an input, core_ts_in need not coincide with the assertion of core_abb_in

on the bus (that is, $\overline{\text{core_ts_in}}$ can be asserted with, or on, a subsequent clock cycle after $\overline{\text{core_abb_in}}$ is asserted; the core tracks this transaction correctly).

In Figure 9-5, the address transfer occurs during bus clock cycles 1 and 2 (arbitration occurs in bus clock cycle 0 and the address transfer is terminated in bus clock 3). In this diagram, the address bus termination input, $\overline{\text{core_aack}}$, is asserted to the core on the bus clock following assertion of $\overline{\text{core_ts_out}}$ (as shown by the dependency line). This is the minimum duration of the address transfer for the core; the duration can be extended by delaying the assertion of $\overline{\text{core_aack}}$ for one or more bus clocks.

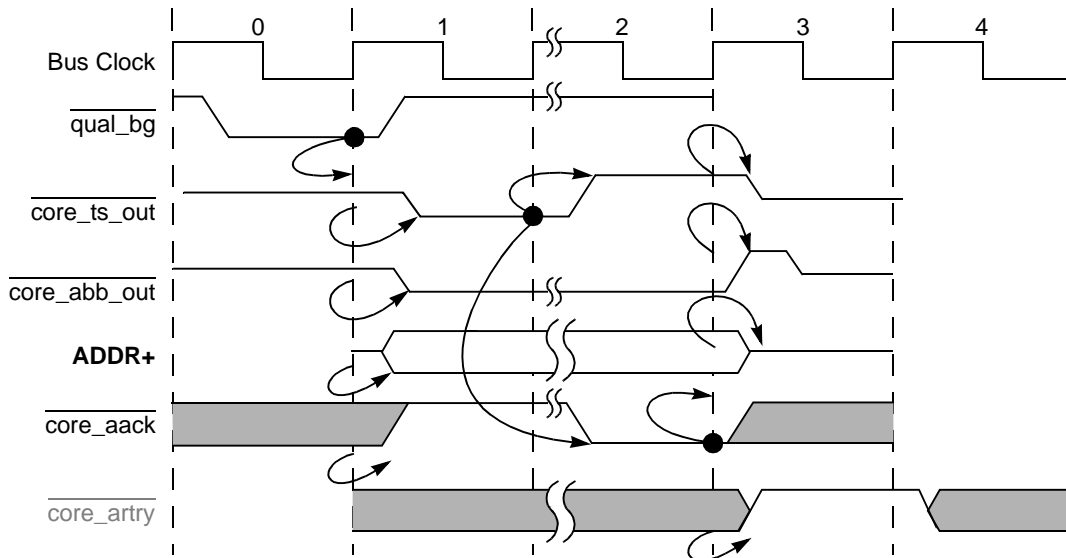


Figure 9-5. Address Bus Transfer

9.3.2.1 Address Bus Parity

The core always generates 1 bit of correct odd-byte parity for each of the 4 bytes of address when a valid address is on the bus. The calculated values are placed on the $\text{core_ap_out}[0:3]$ outputs when the core is the address bus master. If the core is not the master, and $\overline{\text{core_ts_in}}$ and $\overline{\text{core_gbl_in}}$ are asserted together (qualified condition for snooping memory operations), the calculated values are compared with the $\text{core_ap_in}[0:3]$ inputs. If there is an error and address parity checking is enabled (HID0[EBA] is set), the core_ape output is asserted. An address bus parity error causes a checkstop condition if MSR[ME] is cleared. For more information about checkstop conditions, see Chapter 5, “Exceptions.”

9.3.2.2 Address Transfer Attribute Signals

The transfer attribute signals include several encoded signals such as the transfer type (both $\text{core_tt_in}[0:4]$, $\text{core_tt_out}[0:4]$) signals, transfer burst (core_tbst_out) signal, transfer size ($\text{core_tsize}[0:2]$) signals, and transfer code ($\text{core_tc}[0:1]$) signals. Section 8.3.4, “Address

Transfer Attribute Signals,” describes the encodings for the address transfer attribute signals.

9.3.2.2.1 Transfer Type (`core_tt_in[0:4]`, `core_tt_out[0:4]`) Signals

Snooping logic should fully decode the transfer type input signals if the `core_gbl_in` signal is asserted. Slave devices can sometimes use the individual transfer type signals without fully decoding the group. For a complete description of the encoding for the transfer type signals, refer to Table 8-5 and Table 8-6.

9.3.2.2.2 Transfer Size (`core_tsiz[0:2]`) Signals

The transfer size signals (`core_tsiz[0:2]`) indicate the size of the requested data transfer as shown in Table 9-2. These signals may be used along with `core_tbst_out` and `core_a_out[29:31]` to determine which portion of the data bus contains valid data for a write transaction or which portion of the bus should contain valid data for a read transaction. Note that for a burst transaction (as indicated by the assertion of `core_tbst_out`), `core_tsiz[0:2]` are always set to 0b010. Therefore, if `core_tbst_out` is asserted, the memory system should transfer a total of 8 words (32 bytes), regardless of the `core_tsiz[0:2]` encoding.

Table 9-2. Transfer Size Signal Encodings

<code>core_tbst_out</code>	<code>core_tsiz[0:2]</code>			Transfer Size	32-Bit Bus Mode	64-Bit Bus Mode
	<code>core_tsiz0</code>	<code>core_tsiz1</code>	<code>core_tsiz2</code>			
Asserted	0	1	0	8-word burst	8 beats	4 beats
Negated	0	0	0	8 bytes	2 beats	1 beat
Negated	0	0	1	1 byte	1 beat	1 beat
Negated	0	1	0	2 bytes	1 beat	1 beat
Negated	0	1	1	3 bytes	1 beat	1 beat
Negated	1	0	0	4 bytes	1 beat	1 beat
Negated	1	0	1	5 bytes (N/A)	N/A	N/A
Negated	1	1	0	6 bytes (N/A)	N/A	N/A
Negated	1	1	1	7 bytes (N/A)	N/A	N/A

The basic coherency size of the bus is defined to be 32 bytes (corresponding to one cache line). Data transfers that cross an aligned, 32-byte boundary either must present a new address onto the bus at that boundary (for coherency consideration) or must operate as noncoherent data with respect to the core. The core never generates a bus transaction with a transfer size of 5, 6, or 7 bytes.

9.3.2.3 Burst Ordering During Data Transfers

During burst data transfer operations, 32 bytes of data (one cache line) are transferred to or from the cache in order. Burst write transfers are always performed zero-double-word-first, but because burst reads are performed critical-double-word-first, a burst read transfer may not start with the first double word of the cache line, and the cache line fill may wrap around the end of the cache line. This section describes burst ordering for the core when operating in either the 64- or 32-bit bus mode.

Table 9-3 describes the burst ordering when the core is configured with a 64-bit data bus.

Table 9-3. Burst Ordering—64-Bit Bus

Data Transfer	For Starting Address <code>core_a_outx</code> :			
	<code>a_out[27:28] = 00</code>	<code>a_out[27:28] = 01</code>	<code>a_out[27:28] = 10</code>	<code>a_out[27:28] = 11</code>
First data beat	DW0	DW1	DW2	DW3
Second data beat	DW1	DW2	DW3	DW0
Third data beat	DW2	DW3	DW0	DW1
Fourth data beat	DW3	DW0	DW1	DW2

Note: `core_a_out[29:31]` are always 0b000 for burst transfers by the core.

Table 9-4 describes the burst ordering when the core is configured with a 32-bit bus.

Table 9-4. Burst Ordering—32-Bit Bus

Data Transfer	For Starting Address <code>core_a_outx</code> :			
	<code>a_out[27:28] = 00</code>	<code>a_out[27:28] = 01</code>	<code>a_out[27:28] = 10</code>	<code>a_out[27:28] = 11</code>
First data beat	DW0-Upper_word	DW1-Upper_word	DW2-Upper_word	DW3-Upper_word
Second data beat	DW0-Lower_word	DW1-Lower_word	DW2-Lower_word	DW3-Lower_word
Third data beat	DW1-Upper_word	DW2-Upper_word	DW3-Upper_word	DW0-Upper_word
Fourth data beat	DW1-Lower_word	DW2-Lower_word	DW3-Lower_word	DW0-Lower_word
Fifth data beat	DW2-Upper_word	DW3-Upper_word	DW0-Upper_word	DW1-Upper_word
Sixth data beat	DW2-Lower_word	DW3-Lower_word	DW0-Lower_word	DW1-Lower_word
Seventh data beat	DW3-Upper_word	DW0-Upper_word	DW1-Upper_word	DW2-Upper_word
Eighth data beat	DW3-Lower_word	DW0-Lower_word	DW1-Lower_word	DW2-Lower_word

Note: `core_a_out[29:31]` are always 0b000 for burst transfers by the core.

9.3.2.4 Effect of Alignment in Data Transfers (64-Bit Bus)

Table 9-5 lists the aligned transfers that can occur on the 60x bus when configured with a 64-bit width. In these transfers data is aligned to an address that is an integer multiple of the

size of the data. For example, Table 9-5 shows that 1-byte data is always aligned; however, for a 4-byte word to be aligned, it must be at on an address that is a multiple of 4.

Table 9-5. Aligned Data Transfers (64-Bit Bus)

Transfer Size	tsiz0	tsiz1	tsiz2	core_a_out [29:31]	Data Bus Byte Lanes ¹								
					0	1	2	3	4	5	6	7	
Byte	0	0	1	0 0 0	A	—	—	—	—	—	—	—	—
	0	0	1	0 0 1	—	A	—	—	—	—	—	—	—
	0	0	1	0 1 0	—	—	A	—	—	—	—	—	—
	0	0	1	0 1 1	—	—	—	A	—	—	—	—	—
	0	0	1	1 0 0	—	—	—	—	A	—	—	—	—
	0	0	1	1 0 1	—	—	—	—	—	A	—	—	—
	0	0	1	1 1 0	—	—	—	—	—	—	A	—	—
	0	0	1	1 1 1	—	—	—	—	—	—	—	—	A
Half word	0	1	0	0 0 0	A	A	—	—	—	—	—	—	—
	0	1	0	0 1 0	—	—	A	A	—	—	—	—	—
	0	1	0	1 0 0	—	—	—	—	A	A	—	—	—
	0	1	0	1 1 0	—	—	—	—	—	—	A	A	—
Word	1	0	0	0 0 0	A	A	A	A	—	—	—	—	—
	1	0	0	1 0 0	—	—	—	—	A	A	A	A	—
Double word	0	0	0	0 0 0	A	A	A	A	A	A	A	A	A

¹ A: These entries indicate the byte portions of the requested operand that are read or written during that bus transaction.
 —: These entries are not required and are ignored during read transactions and are driven with undefined data during all write transactions.

The G2 core supports misaligned memory operations, although their use may substantially degrade performance. Misaligned memory transfers address memory that is not aligned to the size of the data being transferred (such as, a word read of an odd byte address). Although most of these operations hit in the primary cache (or generate burst memory operations if they miss), the core interface supports misaligned transfers within a word (32-bit aligned) boundary, as shown in Table 9-6. Note that the 4-byte transfer in Table 9-6 is only one example of misalignment. As long as the attempted transfer does not cross a word boundary, the core can transfer the data on the misaligned address (for example, a half-word read from an odd byte-aligned address). An attempt to address data that crosses a word boundary requires two bus transfers to access the data. Note that an attempt to load or store a floating-point operand that is not word-aligned results in a floating-point alignment exception. For more information, refer to Section 5.5.6, “Alignment Exception (0x00600).”

Table 9-6. Misaligned Data Transfers (4-Byte Examples)

Transfer Size (4 Bytes)		tsiz[0:2]	core_a_out [29:31]	Data Bus Byte Lanes							
				0	1	2	3	4	5	6	7
Aligned		1 0 0	0 0 0	A	A	A	A	—	—	—	—
Misaligned:	First access	0 1 1	0 0 1		A	A	A	—	—	—	—
	Second access	0 0 1	1 0 0	—	—	—	—	A	—	—	—
Misaligned:	First access	0 1 0	0 1 0	—	—	A	A	—	—	—	—
	Second access	0 1 0	1 0 0	—	—	—	—	A	A	—	—
Misaligned:	First access	0 0 1	0 1 1	—	—	—	A	—	—	—	—
	Second access	0 1 1	1 0 0	—	—	—	—	A	A	A	—
Aligned		1 0 0	1 0 0	—	—	—	—	A	A	A	A
Misaligned:	First access	0 1 1	1 0 1	—	—	—	—	—	A	A	A
	Second access	0 0 1	0 0 0	A	—	—	—	—	—	—	—
Misaligned:	First access	0 1 0	1 1 0	—	—	—	—	—	—	A	A
	Second access	0 1 0	0 0 0	A	A	—	—	—	—	—	—
Misaligned:	First access	0 0 1	1 1 1	—	—	—	—	—	—	—	A
	Second access	0 1 1	0 0 0	A	A	A	—	—	—	—	—

Notes:

- A: Byte lane used.
- : Byte lane not used.

Due to the performance degradations associated with misaligned memory operations, they are best avoided. Address translation logic can also generate substantial exception overhead when the load/store multiple and load/store string instructions access misaligned data, another reason to avoid using these instructions. It is strongly recommended that software attempt to align code and data where possible.

9.3.2.5 Effect of Alignment in Data Transfers (32-Bit Bus)

The aligned data transfer cases for 32-bit data bus mode are shown in Table 9-7. All of the transfers require a single data beat (if caching-inhibited or write-through) except for double-word cases which require two data beats. The double-word case is only generated by the core for load or store double operations to/from the floating-point GPRs. All caching-inhibited instruction fetches are performed as word operations.

Table 9-7. Aligned Data Transfers (32-Bit Bus Mode)

Transfer Size	tsiz0	tsiz1	tsiz2	core_a_out [29:31]	Data Bus Byte Lanes							
					0	1	2	3	4	5	6	7
Byte	0	0	1	0 0 0	A	—	—	—	x	x	x	x
	0	0	1	0 0 1	—	A	x	—	x	x	x	x
	0	0	1	0 1 0	—	—	A	—	x	x	x	x
	0	0	1	0 1 1	—	—	—	A	x	x	x	x
	0	0	1	1 0 0	A	—	—	—	x	x	x	x
	0	0	1	1 0 1	—	A	—	—	x	x	x	x
	0	0	1	1 1 0	—	—	A	—	x	x	x	x
	0	0	1	1 1 1	—	—	—	A	x	x	x	x
Half word	0	1	0	0 0 0	A	A	—	—	x	x	x	x
	0	1	0	0 1 0	—	—	A	A	x	x	x	x
	0	1	0	1 0 0	A	A	—	—	x	x	x	x
	0	1	0	1 1 0	—	—	A	A	x	x	x	x
Word	1	0	0	0 0 0	A	A	A	A	x	x	x	x
	1	0	0	1 0 0	A	A	A	A	x	x	x	x
Double word	0	0	0	0 0 0	A	A	A	A	x	x	x	x
Second beat	0	0	0	0 0 0	A	A	A	A	x	x	x	x

Notes:

- A: Byte lane used.
- : Byte lane not used.
- x: Byte lane not used in 32-bit bus mode.

Misaligned data transfers when the core is configured with a 32-bit data bus operate in the same way as when configured with a 64-bit data bus, with the exception that only the core_dh_out[0:31] or core_dh_in[0:31] data bus is used. See Table 9-8 for an example of a 4-byte misaligned transfer starting at each possible byte address within a double word.

Table 9-8. Misaligned 32-Bit Data Bus Transfer (4-Byte Examples)

Transfer Size (4 Bytes)		core_tsiz [0:2]	core_a-out [29:31]	Data Bus Byte Lanes							
				0	1	2	3	4	5	6	7
Aligned		1 0 0	0 0 0	A	A	A	A	x	x	x	x
Misaligned:	First access	0 1 1	0 0 1		A	A	A	x	x	x	x
	Second access	0 0 1	1 0 0	A	—	—	—	x	x	x	x
Misaligned:	First access	0 1 0	0 1 0	—	—	A	A	x	x	x	x
	Second access	0 1 0	1 0 0	A	A	—	x	x	x	x	x
Misaligned:	First access	0 0 1	0 1 1	—	—	—	A	x	x	x	x
	Second access	0 1 1	1 0 0	A	A	A	—	x	x	x	x
Aligned		1 0 0	1 0 0	A	A	A	A	x	x	x	x
Misaligned:	First access	0 1 1	1 0 1	—	A	A	A	x	x	x	x
	Second access	0 0 1	0 0 0	A	—	—	—	x	x	x	x
Misaligned:	First access	0 1 0	1 1 0	—	—	A	A	x	x	x	x
	Second access	0 1 0	0 0 0	A	A	—	—	x	x	x	x
Misaligned:	First access	0 0 1	1 1 1	—	—	—	A	x	x	x	x
	Second access	0 1 1	0 0 0	A	A	A	—	x	x	x	x

Notes:

- A: Byte lane used.
- : Byte lane not used.
- x: Byte lane not used in 32-bit bus mode.

9.3.2.5.1 Alignment of External Control Instructions

The size of the data transfer associated with the **eciwx** and **ecowx** instructions is always 4 bytes. However, if either of these instructions is misaligned and crosses any word boundary, the core generates two bus operations, each smaller than 4 bytes. For the first bus operation, bits `core_a_out[29:31]` equal bits 29–31 of the effective address of the instruction, which is 0b101, 0b110, or 0b111. The size associated with the first bus operation will be 3, 2, or 1 bytes, respectively. For the second bus operation, bits `core_a_out[29:31]` equal 0b000 and the size associated with the operation is 1, 2, or 3 bytes, respectively. For both operations, `core_tbst_out` and `core_tsiz[0:2]` are redefined to specify the resource ID (RID). The resource ID is copied from bits 28–31 of the EAR. For **eciwx/ecowx** operations, EAR[28] is set if `core_tbst_out` is high. The size of the second bus operation cannot be deduced from the operation itself; the system must determine how many bytes were transferred on the first bus operation to determine the size of the second operation.

Furthermore, the two bus operations associated with such a misaligned external control instruction are not atomic. That is, the core may initiate other types of memory operations

between the two transfers. Also, the two bus operations associated with a misaligned **ecowx** may be interrupted by an **eciwx** bus operation, and vice versa. The core guarantees that the two operations associated with a misaligned **ecowx** cannot be interrupted by another **ecowx** operation.

Because a misaligned external control address is considered a programming error, the system may choose to assert $\overline{\text{core_tea}}$ or otherwise cause an exception when a misaligned external control bus operation occurs.

9.3.2.6 Transfer Code (core_tc[0:1]) Signals

The `core_tc0` and `core_tc1` signals provide supplemental information about the corresponding address. Note that the `core_tcx` signals can be used with both `core_tt_in[0:4]` and `core_tt_out[0:4]`, and both `core_tbst_in` and `core_tbst_out` signals to further define the current transaction.

Table 9-9 shows the encodings of the `core_tc[0:1]` signals.

Table 9-9. Transfer Code Encoding

core_tc[0:1]	Read	Write
0 0	Data transaction	Any write
0 1	Touch load	N/A
1 0	Instruction fetch	N/A
1 1	(Reserved)	N/A

9.3.3 Address Transfer Termination

The address tenure of a bus operation is terminated when completed with the assertion of $\overline{\text{core_aack}}$, or retried with the assertion of $\overline{\text{core_artry_in}}$. The G2 core does not terminate the address transfer until the $\overline{\text{core_aack}}$ (address acknowledge) input is asserted; therefore, the system can extend the address transfer phase by delaying the assertion of $\overline{\text{core_aack}}$ to the G2 core. $\overline{\text{core_aack}}$ can be asserted as early as the bus clock cycle following $\overline{\text{core_ts_in}}$ (see Figure 9-6), which allows a minimum address tenure of two bus cycles. However, when the core clock is configured for 1:1 or 1.5:1 processor core-to-bus clock mode, the $\overline{\text{core_artry_out}}$ snoop response cannot be determined in the minimum allowed address tenure period. Thus, in a system with two or more G2 cores using 1:1 or 1.5:1 clock mode, $\overline{\text{core_aack}}$ must not be asserted until the third clock of the address tenure (one address wait state) to allow the snooping G2 cores an opportunity to assert $\overline{\text{core_artry_in}}$ on the bus. For other clock configurations (2:1, 2.5:1, 3:1, 3.5:1, and 4:1), the $\overline{\text{core_artry_out}}$ snoop response can be determined in the minimum address tenure period, and $\overline{\text{core_aack}}$ may be asserted as early as the second bus clock of the address tenure. As shown in Figure 9-6, these signals are asserted for one bus clock cycle, three-stated for half of the next bus clock

cycle, driven high until the following bus cycle, and finally three-stated. Note that `core_aack` must be asserted for only one bus clock cycle.

The address transfer can be terminated with the requirement to retry if `core_artry_in` is asserted anytime during the address tenure and through the cycle following `core_aack`. The assertion causes the entire transaction (address and data tenure) to be rerun. As a snooping device, the G2 core asserts `core_artry_out` for a snooped transaction that hits modified data in the data cache that must be written back to memory, or if the snooped transaction could not be serviced. As a bus master, the core responds to an assertion of `core_artry_out` by aborting the bus transaction and re-requesting the bus. Note that after recognizing an assertion of `core_artry_out` and aborting the transaction in progress, the G2 core is not guaranteed to run the same transaction the next time it is granted the bus due to internal reordering of load and store operations.

If an address retry is required, the `core_artry_in` response is asserted by a bus snooping device as early as the second cycle after the assertion of `core_ts_out` (or until the third cycle following `core_ts_out` if 1:1 or 1.5:1 processor core to bus clock ratio is selected). Once asserted, `core_artry_in` must remain asserted through the cycle after the assertion of `core_aack`. The assertion of `core_artry_in` during the cycle after the assertion of `core_aack` is referred to as a qualified `core_artry_in`. An earlier assertion of `core_artry_in` during the address tenure is referred to as an early `core_artry_in`.

As a bus master, the G2 core recognizes either an early or qualified `core_artry_in` and prevents the data tenure associated with the retried address tenure. If the data tenure has already begun, the core aborts and terminates the data tenure immediately even if the burst data has been received. If the assertion of `core_artry_in` is received up to or on the bus cycle following the first (or only) assertion of `core_ta` for the data tenure, the core ignores the first data beat, and if it is a load operation, does not forward data internally to the cache and execution units. If `core_artry_in` is asserted after the first (or only) assertion of `core_ta`, improper operation of the bus interface may result.

During the clock of a qualified `core_artry_in`, the G2 core also determines if it should negate `core_br` and ignore `core_bg` on the following cycle. On the following cycle, only the snooping master that asserted `core_artry_in` and needs to perform a snoop copy-back operation is allowed to assert `core_br`. This guarantees the snooping master an opportunity to request and be granted the bus before the just-retried master can restart its transaction. Note that a nonclocked bus arbiter may detect the assertion of address bus request by the bus master that asserted `core_artry_in`, and return a qualified bus grant one cycle earlier than shown in Figure 9-6.

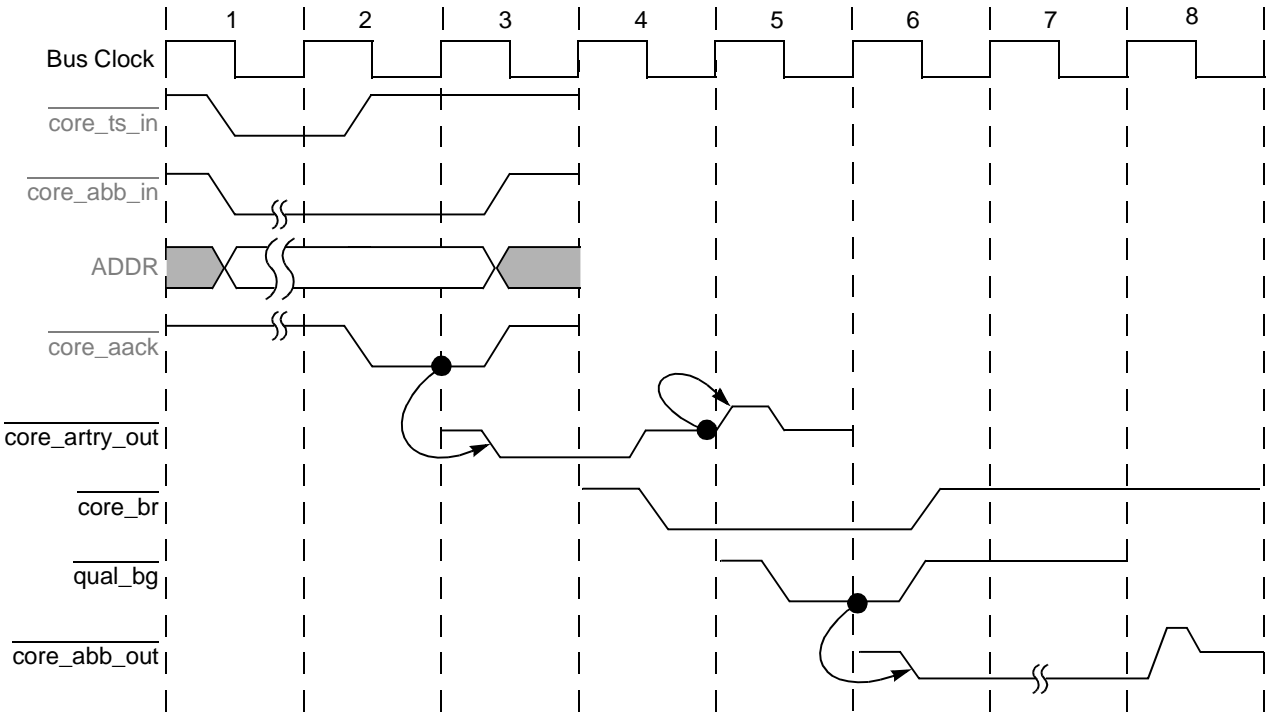


Figure 9-6. Snoop Address Cycle with `core_artry_out`

9.4 Data Bus Tenure

This section describes the data bus arbitration, transfer, and termination phases defined by the G2 core memory access protocol. The phases of the data tenure are identical to those of the address tenure, underscoring the symmetry in the control of the two buses.

9.4.1 Data Bus Arbitration

Data bus arbitration uses the data arbitration signal group—`core_dbg`, `core_dbwo`, and both `core_dbb` signals. Additionally, the combination of `core_ts_out` and `TT[0:4]` provides information about the data bus request to external logic.

The `core_ts_out` signal is an implied data bus request from the core; the arbiter must qualify `core_ts_out` with the transfer type `core_tt_out` encodings to determine if the current address transfer is an address-only operation, which does not require a data bus transfer (see Figure 9-6). If the data bus is needed, the arbiter grants data bus mastership by asserting the `core_dbg` input to the core. As with the address bus arbitration phase, the G2 core must qualify the `core_dbg` input with a number of input signals before assuming bus mastership, as shown in Figure 9-7.

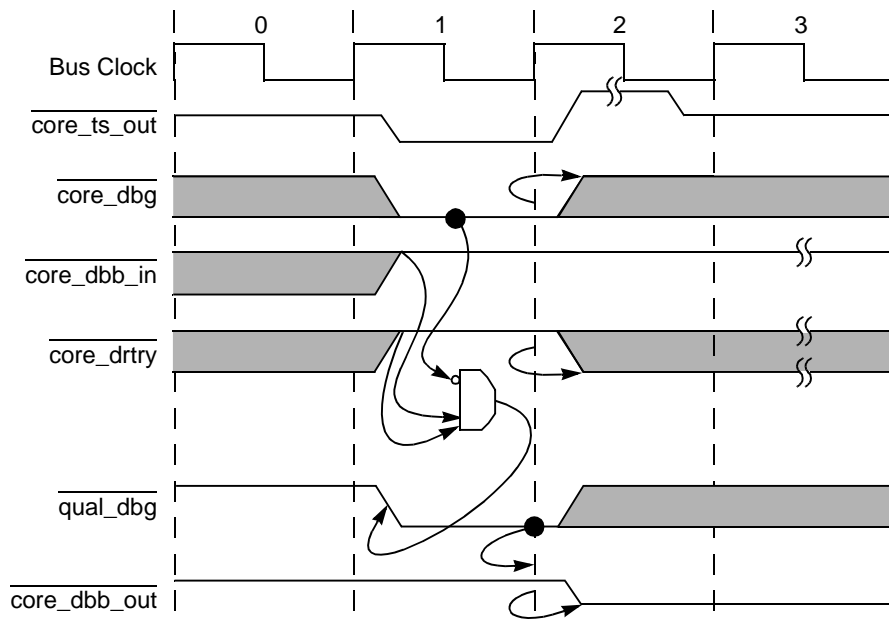


Figure 9-7. Data Bus Arbitration

A qualified data bus grant can be expressed as the following:

Qualified data bus grant = core_dbg asserted while core_dbb_out, core_drtry, and core_artry_out (associated with the data bus operation) are negated.

When a data tenure overlaps with its associated address tenure, a qualified core_artry_out assertion coincident with a data bus grant signal does not result in data bus mastership (core_dbb_out is not asserted). Otherwise, the G2 core always asserts core_dbb_out on the bus clock cycle after recognition of a qualified data bus grant. Because the core can pipeline transactions, there may be an outstanding data bus transaction when a new address transaction is retried. In this case, the core becomes the data bus master to complete the previous transaction.

9.4.1.1 Using the core_dbb_out Signal

The core_dbb_out signal should be connected between masters if data tenure scheduling is left to the masters. Optionally, the memory system can control data tenure scheduling directly with core_dbg. However, it is possible to ignore the core_dbb_out signal in the system if the core_dbb_out input is not used as the final data bus allocation control between data bus masters, and if the memory system can track the start and end of the data tenure. If core_dbb_out is not used to signal the end of a data tenure, core_dbg is only asserted to the next bus master the cycle before the cycle that the next bus master may actually begin its data tenure, rather than asserting it earlier (usually during another master's data tenure) and allowing the negation of core_dbb_out to be the final gating signal for a qualified data bus grant. Even if core_dbb_out is ignored in the system, the G2 core always recognizes its own assertion of core_dbb_out and requires one cycle after data tenure completion to

negate its own `core_dbb_out` before recognizing a qualified data bus grant for another data tenure. If `DBB` is ignored in the system, it must still be connected to a pull-up resistor on the G2 core to ensure proper operation.

9.4.2 Data Bus Write Only

As a result of address pipelining, the core may have up to two data tenures queued to perform when it receives a qualified `core_dbg`. Generally, the data tenures should be performed in strict order (the same order) as their address tenures were performed. The core, however, also supports a limited out-of-order capability with the data bus write only (`core_dbwo`) input. When recognized on the clock of a qualified `core_dbg`, `core_dbwo` may direct the core to perform the next pending data write tenure even if a pending read tenure would have normally been performed first. For more information on the operation of `core_dbwo`, refer to Section 9.10, “Using `core-dbwo` (Data Bus Write Only).”

If the G2 core has any data tenures to perform, it always accepts data bus mastership to perform a data tenure when it recognizes a qualified `core_dbg`. If `core_dbwo` is asserted with a qualified `core_dbg` and no write tenure is queued to run, the G2 core still takes mastership of the data bus to perform the next pending read data tenure.

Generally, `core_dbwo` should only be used to allow a copy-back operation (burst write) to occur before a pending read operation. If `core_dbwo` is used for single-beat write operations, it may negate the effect of the `eiio` instruction by allowing a write operation to precede a program-scheduled read operation.

9.4.3 Data Transfer

The data transfer signals include both input and output signals of `core_dh[0:31]`, `core_dl[0:31]`, `core_dp[0:7]`, and only output signal of `core_dpe`. For memory accesses, both input and output signals of `core_dh` and `core_dl` form a 64-bit data path for read and write operations.

The G2 core transfers data in either single- or four-beat burst transfers when configured with a 64-bit data bus; when configured with a 32-bit data bus, the G2 core performs one-, two-, and eight-beat data transfers. Single-beat operations can transfer from 1 to 8 bytes at a time and can be misaligned; see Section 9.3.2.4, “Effect of Alignment in Data Transfers (64-Bit Bus).” Burst operations always transfer eight words and are aligned on eight-word address boundaries. Burst transfers can achieve significantly higher bus throughput than single-beat operations.

The type of transaction initiated by the G2 core depends on whether the code or data is cacheable and, for store operations, whether the cache is considered in write-back or write-through mode, which software controls on either a page or block basis. Burst transfers support cacheable operations only; that is, memory structures must be marked as

cacheable (and write-back for data store operations) in the respective page or block descriptor to take advantage of burst transfers.

The core output $\overline{\text{core_tbst_out}}$ indicates to the system whether the current transaction is a single- or four-beat transfer (except during **eciwx/ecowx** transactions, when it signals the state of EAR[28]). A burst transfer has an assumed address order. For load or store operations that miss in the cache (and are marked as cacheable and, for stores, write-back), the G2 core uses the double-word-aligned address associated with the critical code or data that initiated the transaction. This minimizes latency by allowing the critical code or data to be forwarded to the processor before the rest of the cache line is filled. For all other burst operations, however, the cache line is transferred beginning with the eight-word-aligned data.

The G2 core does not directly support dynamic interfacing to subsystems with less than a 64-bit data path. It does, however, provide a static 32-bit data bus mode; for more information, see Section 9.1.3, “Optional 32-Bit Data Bus Mode.”

9.4.4 Data Transfer Termination

Four signals are used to terminate data bus transactions— $\overline{\text{core_ta}}$, $\overline{\text{core_drtry}}$ (data retry), core_tea (transfer error acknowledge), and core_artry_in . The core_ta signal indicates normal termination of data transactions. It must always be asserted on the bus cycle coincident with the data that it is qualifying. It may be withheld by the slave for any number of clocks until valid data is ready to be supplied or accepted. core_drtry indicates invalid read data in the previous bus clock cycle. core_drtry extends the current data beat and does not terminate it. If it is asserted after the last (or only) data beat, the core negates core_dbb_out but still considers the data beat active and waits for another assertion of core_ta . core_drtry is ignored on write operations. core_tea indicates a nonrecoverable bus error event. Upon receiving a final (or only) termination condition, the core always negates core_dbb_out for one cycle.

If core_drtry is asserted by the memory system to extend the last (or only) data beat past the negation of core_dbb_out , the memory system should three-state the data bus on the clock after the final assertion of core_ta , even though it will negate core_drtry on that clock. This is to prevent a potential momentary data bus conflict if a write access begins on the following cycle.

The core_tea signal is used to signal a nonrecoverable error during the data transaction. It may be asserted on any cycle during core_dbb_out , or on the cycle after a qualified core_ta during a read operation, except when no- core_drtry mode is selected (where no- core_drtry mode cancels checking the cycle after core_ta). The assertion of core_tea terminates the data tenure immediately, even if in the middle of a burst; however, it does not prevent incorrect data that has just been acknowledged with core_ta from being written into the G2 core cache or GPRs. The assertion of core_tea initiates either a machine check exception or a checkstop condition based on the setting of the MSR.

An assertion of $\overline{\text{core_artry_in}}$ causes the data tenure to be terminated immediately if $\overline{\text{core_artry_in}}$ is for the address tenure associated with the data tenure in operation. If $\overline{\text{core_artry_out}}$ is connected for the G2 core, the earliest allowable assertion of $\overline{\text{core_ta}}$ to the core is directly dependent on the earliest possible assertion of $\overline{\text{core_artry_in}}$ to the G2 core; see Section 9.3.3, “Address Transfer Termination.”

If the G2 core clock is configured for 1:1 or 1.5:1 (processor clock to bus clock ratio) mode and the core is performing a burst read into its data cache, at least one wait state must be provided between the assertion of $\overline{\text{core_ts}}$ and the first assertion of $\overline{\text{core_ta}}$ for that transaction. If no- $\overline{\text{core_drtry}}$ mode is also selected, at least two wait states must be provided. The wait states are required due to possible resource contention in the data cache caused by a block replacement (or cast-out) required in connection with the new linefill. These wait states may be provided by withholding the assertion of $\overline{\text{core_ta}}$ to the G2 core for that data tenure, or by withholding $\overline{\text{core_dbg}}$ to the core, thereby delaying the start of the data tenure. This restriction applies only to burst reads into the data cache when configured in 1:1 or 1.5:1 clock modes. (It does not apply to instruction fetches, write operations, noncacheable read operations, or non-1:1 or non-1.5:1 clock modes.)

9.4.4.1 Normal Single-Beat Termination

Normal termination of a single-beat data read operation occurs when $\overline{\text{core_ta}}$ is asserted by a responding slave. The $\overline{\text{core_tea}}$ and $\overline{\text{core_drtry}}$ signals must remain negated during the transfer (see Figure 9-8).

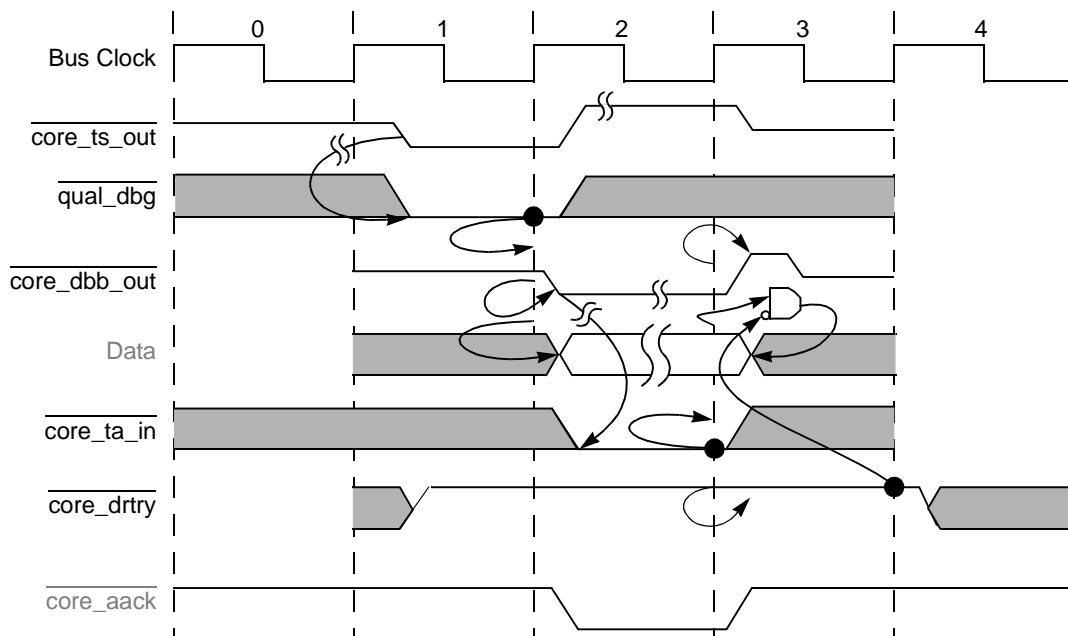


Figure 9-8. Normal Single-Beat Read Termination

The `core_drtry` signal is not sampled during data writes, as shown in Figure 9-9.

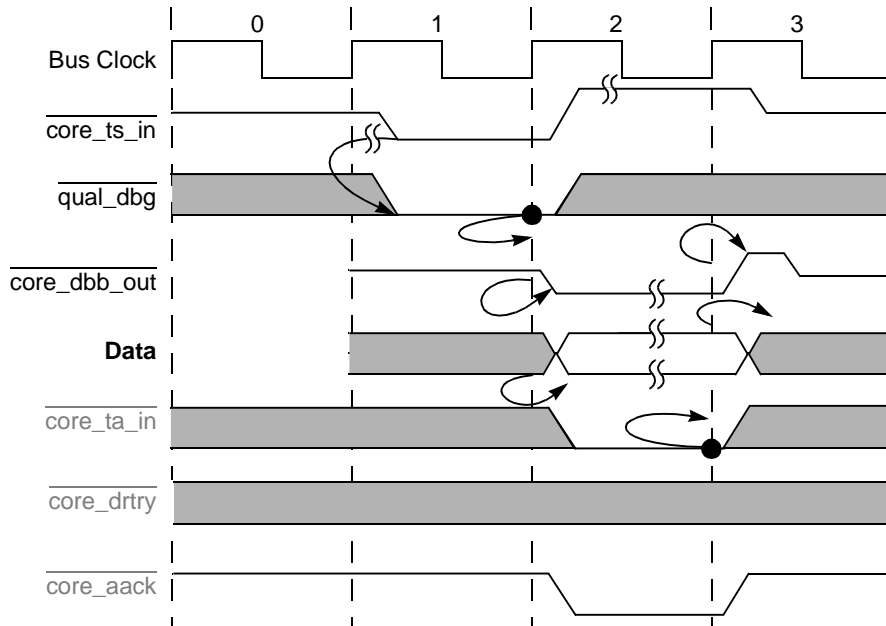


Figure 9-9. Normal Single-Beat Write Termination

9.4.4.2 Normal Burst Termination

Normal termination of a burst transfer occurs when `core_ta` is asserted for four bus clock cycles, as shown in Figure 9-10. The bus clock cycles in which `core_ta` is asserted need not be consecutive, thus allowing pacing of the data transfer beats. For read bursts to terminate successfully, `core_tea` and `core_drtry` must remain negated during the transfer. For write bursts, `core_tea` must remain negated for a successful transfer. `core_drtry` is ignored during data writes.

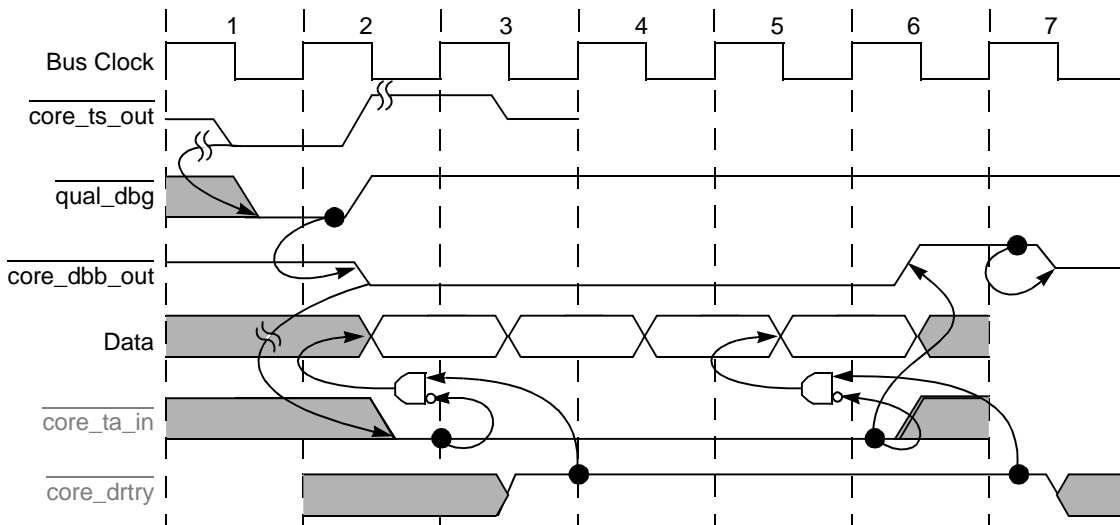


Figure 9-10. Normal Burst Transaction

For read bursts, $\overline{\text{core_drtry}}$ may be asserted one bus clock cycle after $\overline{\text{core_ta}}$ is asserted to signal that the data presented with $\overline{\text{core_ta}}$ is invalid and that the processor must wait for the negation of $\overline{\text{core_drtry}}$ before forwarding data to the processor (see Figure 9-11). Thus, a data beat can be terminated by a predicted branch with $\overline{\text{core_ta}}$ and then one bus clock cycle later confirmed with the negation of $\overline{\text{core_drtry}}$. The $\overline{\text{core_drtry}}$ signal is valid only for read transactions. $\overline{\text{core_ta}}$ must be asserted on the bus clock cycle before the first bus clock cycle of the assertion of $\overline{\text{core_drtry}}$; otherwise the results are undefined.

The $\overline{\text{core_drtry}}$ signal extends data bus mastership such that other processors cannot use the data bus until $\overline{\text{core_drtry}}$ is negated. Therefore, in the example shown in Figure 9-11, $\overline{\text{core_dbb_out}}$ cannot be asserted until bus clock cycle 5. This is true for both read and write operations even though $\overline{\text{core_drtry}}$ does not extend bus mastership for write operations.

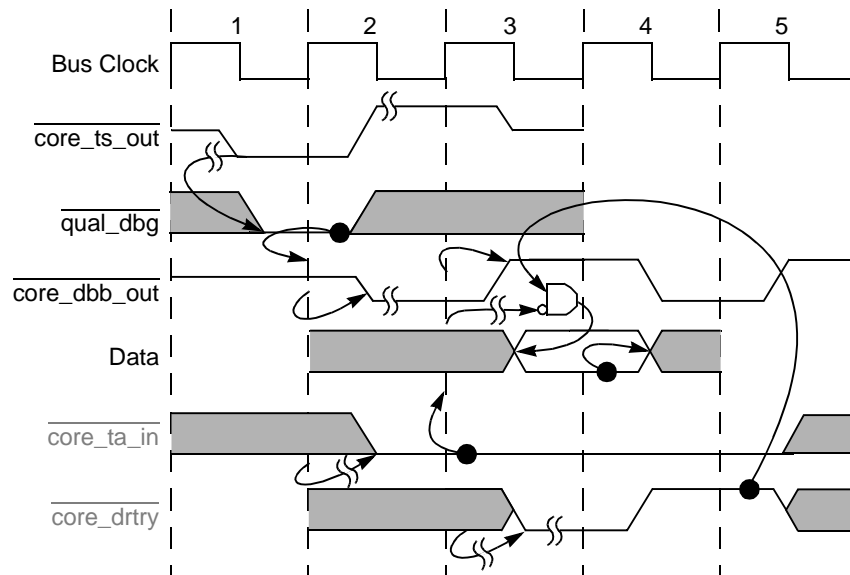


Figure 9-11. Termination with DRTRY

Figure 9-12 shows the effect of using $\overline{\text{core_drtry}}$ during a burst read. It also shows the effect of using $\overline{\text{core_ta}}$ to pace the data transfer rate. Notice that in bus clock cycle 3 in Figure 9-12, $\overline{\text{core_ta}}$ is negated for the second data beat. The G2 core data pipeline does not proceed until bus clock cycle 4, when $\overline{\text{core_ta}}$ is reasserted.

Note that $\overline{\text{core_drtry}}$ is useful for systems that implement predicted forwarding of data such as those with direct-mapped, second-level caches where hit/miss is determined on the following bus clock cycle, or for parity- or ECC-checked memory systems.

Note that $\overline{\text{core_drtry}}$ may not be implemented on other processors of this family.

9.4.4.3 Data Transfer Termination Due to a Bus Error

The $\overline{\text{core_tea}}$ signal indicates that a bus error occurred. It may be asserted while $\overline{\text{core_dbb_out}}$ (and/or $\overline{\text{core_drtry}}$ for read operations) is asserted. Asserting $\overline{\text{core_tea}}$ to the

core terminates the transaction; that is, further assertions of $\overline{\text{core_ta}}$ and $\overline{\text{core_drtry}}$ are ignored and $\overline{\text{core_dbb_out}}$ is negated; see Figure 9-12.

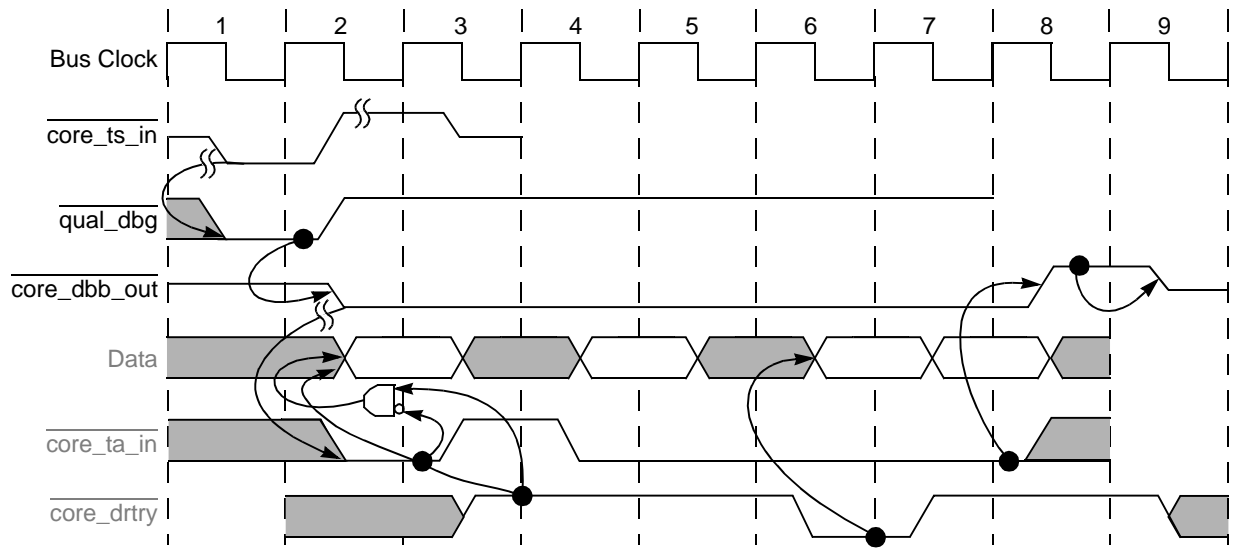


Figure 9-12. Read Burst with $\overline{\text{core_ta}}$ Wait States and $\overline{\text{core_drtry}}$

Assertion of the $\overline{\text{core_tea}}$ signal causes a machine check exception (and possibly a checkstop condition within the core). For more information, see Section 5.5.2, “Machine Check Exception (0x00200).” Note also that the G2 core does not implement a synchronous error capability for memory accesses. This means that the exception instruction pointer does not point to the memory operation that caused the assertion of $\overline{\text{core_tea}}$, but to the instruction about to be executed (perhaps several instructions later). However, assertion of $\overline{\text{core_tea}}$ does not invalidate data entering the GPR or the cache. Additionally, the corresponding address of the access that caused $\overline{\text{core_tea}}$ to be asserted is not latched by the G2 core. To recover, the exception handler must determine and remedy the cause of the $\overline{\text{core_tea}}$, or the G2 core must be reset; therefore, this function should only be used to flag fatal system conditions to the processor (such as parity or uncorrectable ECC errors).

After the G2 core has committed to run a transaction, that transaction must eventually complete. Address retry causes the transaction to be restarted; $\overline{\text{core_ta}}$ wait states and $\overline{\text{core_drtry}}$ assertion for reads delay termination of individual data beats. Eventually, however, the system must either terminate the transaction or assert the $\overline{\text{core_tea}}$ signal (and vector the core into a machine check exception.) For this reason, care must be taken to check for the end of physical memory and the location of certain system facilities to avoid memory accesses that result in the generation of machine check exceptions.

Note that $\overline{\text{core_tea}}$ generates a machine check exception depending on MSR[ME]. Clearing the machine check exception enable control bits leads to a true checkstop condition (instruction execution halted and processor clock stopped).

9.4.5 Memory Coherency—MEI Protocol

The G2 core provides dedicated hardware to provide memory coherency by snooping bus transactions. The address retry capability enforces the three-state, MEI cache-coherency protocol (see Figure 9-13).

The global ($\overline{\text{core_gbl_out}}$) output signal indicates whether the current transaction must be snooped by other snooping devices on the bus. Address bus masters assert core_gbl_out to indicate that the current transaction is a global access (that is, an access to memory shared by more than one device). If core_gbl_in is not asserted for the transaction, that transaction is not snooped. When other devices detect the core_gbl_in input asserted, they must respond by snooping the broadcast address.

Normally, $\overline{\text{core_gbl_out}}$ reflects the M-bit value specified for the memory reference in the corresponding translation descriptor. Note that care must be taken to minimize the number of pages marked as global, because the retry protocol discussed in the previous Section 9.4.4, “Data Transfer Termination” is used to enforce coherency and can require significant bus bandwidth.

When the G2 core is not the address bus master, $\overline{\text{core_gbl_out}}$ is an input. The core snoops a transaction if $\overline{\text{core_ts}}$ and $\overline{\text{core_gbl_out}}$ are asserted together in the same bus clock cycle (this is a qualified snooping condition). No snoop update to the core cache occurs if the snooped transaction is not marked global. This includes invalidation cycles.

When the G2 core detects a qualified snoop condition, the address associated with the core_ts is compared against the data cache tags. Snooping completes if no hit is detected. However, if the address hits in the cache, the core reacts according to the MEI protocol shown in Figure 9-13, assuming the WIM bits are set to write-back, caching-allowed, and coherency-enforced modes (WIM = 001).

The G2 core on-chip data cache is implemented as a four-way set-associative cache. To facilitate external monitoring of the internal cache tags, the cache set entry ($\text{core_cse}[0:1]$) signals indicate which cache set is being replaced on read operations. Note that these signals are valid only for core burst operations; for all other bus operations, $\text{core_cse}[0:1]$ should be ignored.

9.5 Timing Examples

This section shows timing diagrams for various scenarios. Figure 9-14 illustrates the fastest single-beat reads possible for the G2 core. This figure shows both minimal latency and maximum single-beat throughput. By delaying the data bus tenure, the latency increases, but, because of split-transaction pipelining, the overall throughput is not affected unless the data bus latency causes the third address tenure to be delayed.

Note that all bidirectional signals are three-stated between bus tenures.

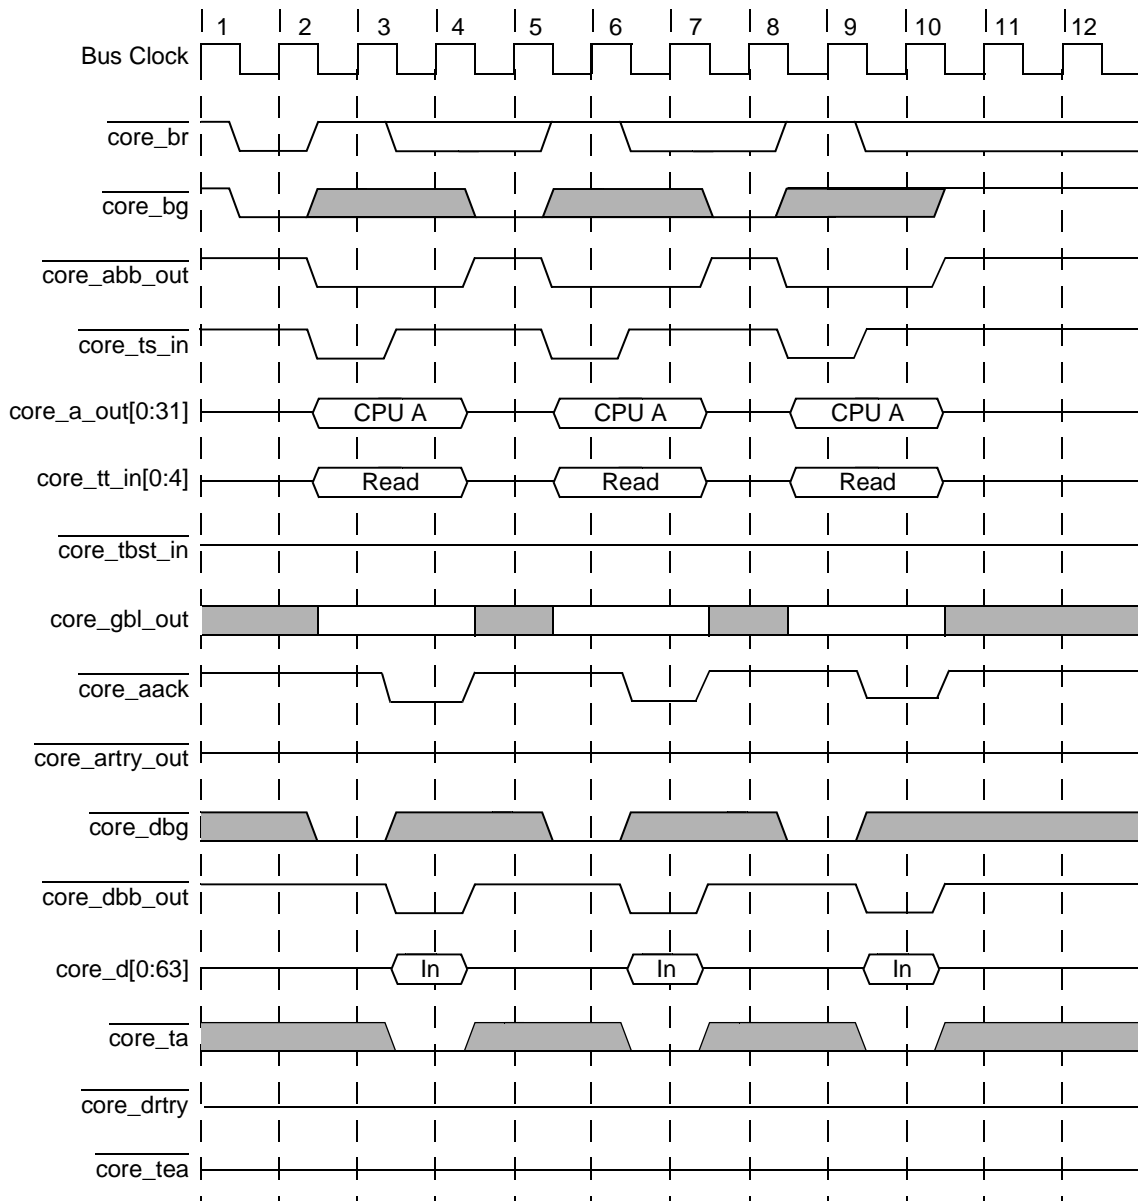


Figure 9-14. Fastest Single-Beat Reads

Figure 9-15 illustrates the fastest single-beat writes supported by the core. All bidirectional signals are three-stated between bus tenures.

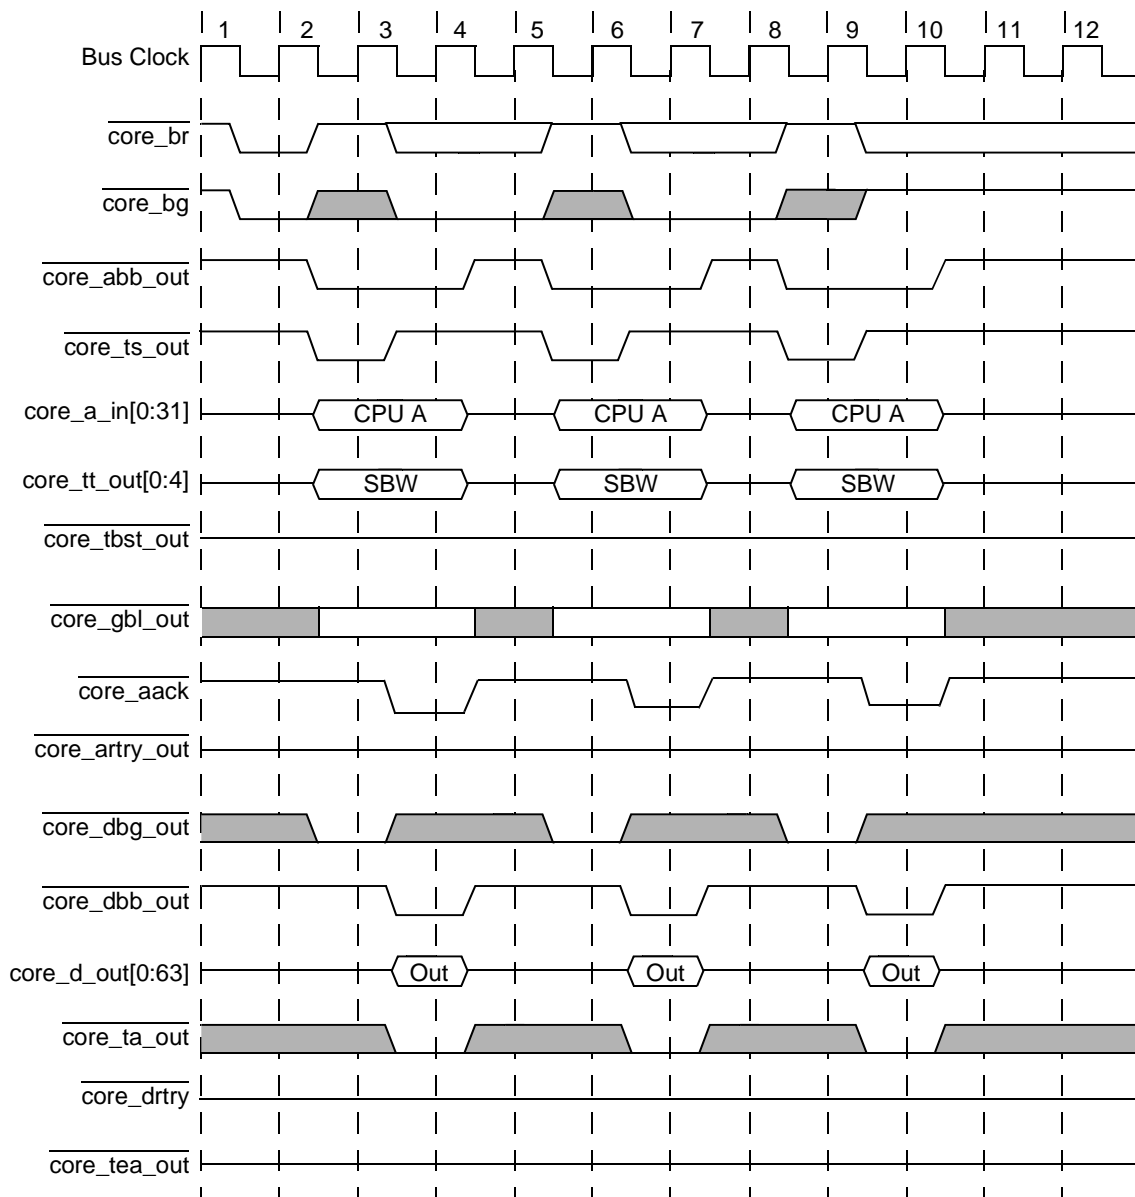


Figure 9-15. Fastest Single-Beat Writes

Figure 9-16 shows three ways to delay single-beat reads showing data-delay controls:

- The $\overline{\text{core_ta}}$ signal can remain negated to insert wait states in clock cycles 3 and 4.
- For the second access, $\overline{\text{core_dbg}}$ could have been asserted in clock cycle 6.
- In the third access, $\overline{\text{core_drtry}}$ is asserted in clock cycle 11 to flush the previous data.

Note that all bidirectional signals are three-stated between bus tenures. The pipelining shown in Figure 9-16 can occur if the second access is not another load (for example, an instruction fetch).

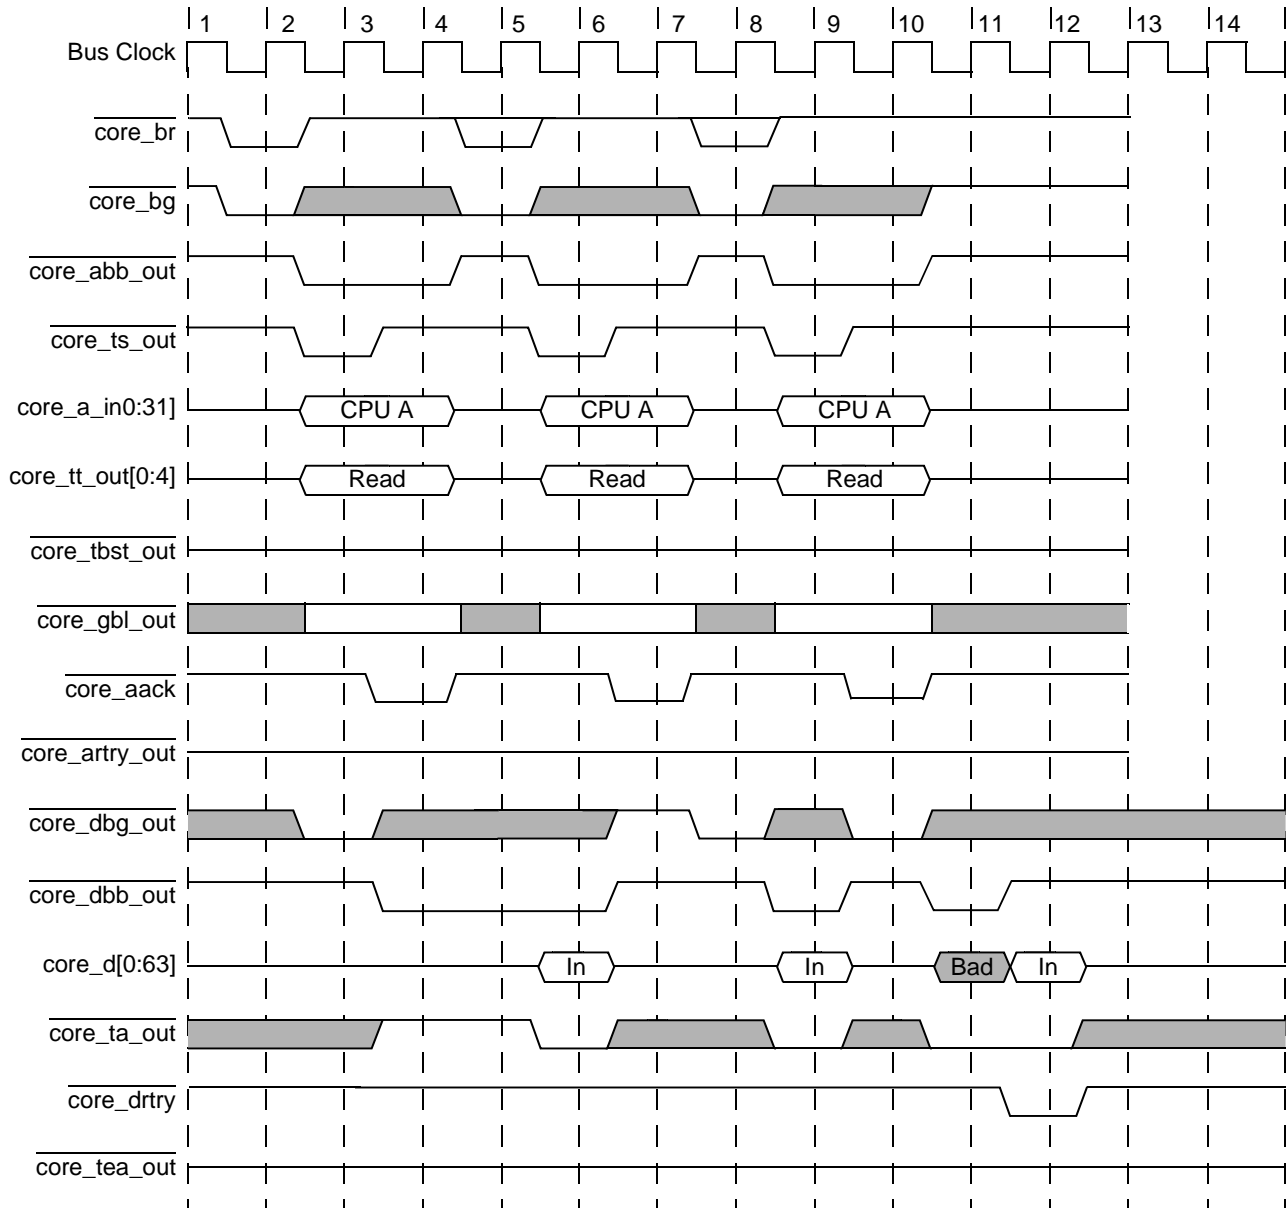


Figure 9-16. Single-Beat Reads Showing Data-Delay Controls

Figure 9-17 shows data-delay controls in a single-beat write operation. Note that all bidirectional signals are three-stated between bus tenures. Data transfers are delayed in the following ways:

- The $\overline{\text{core_ta}}$ signal is held negated to insert wait states in clocks 3 and 4.
- In clock 6, $\overline{\text{core_dbg}}$ is held negated, delaying the start of the data tenure.

The last access is not delayed ($\overline{\text{core_drtry}}$ is valid only for read operations).

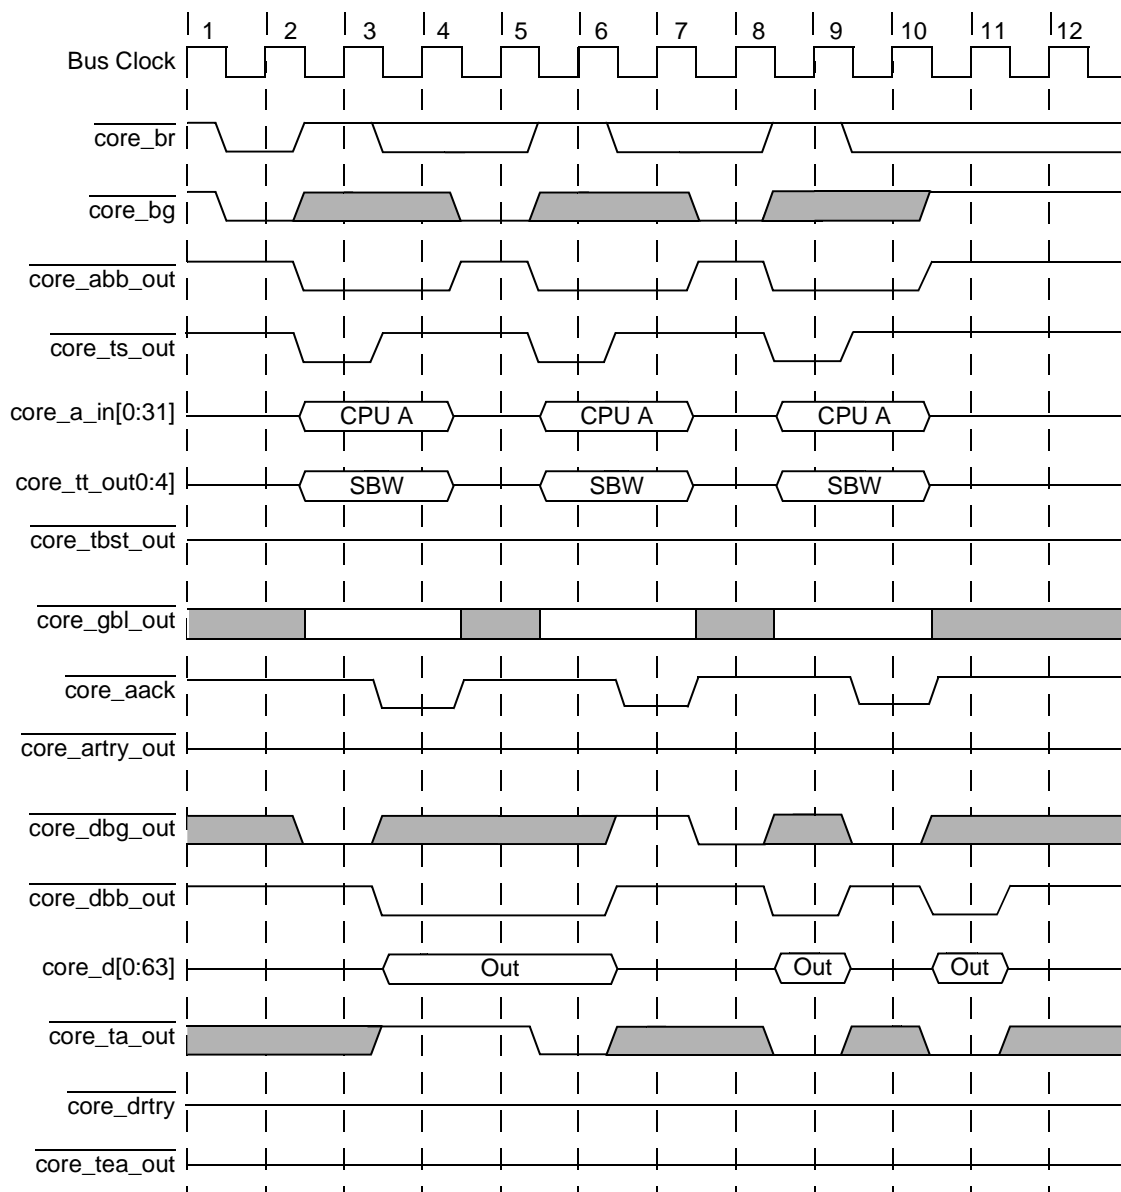


Figure 9-17. Single-Beat Writes Showing Data-Delay Controls

Figure 9-18 shows the use of data-delay controls with burst transfers. Note that all bidirectional signals are three-stated between bus tenures. Note the following:

- The first data beat of bursted read data (clock 0) is the critical quad word in 64-bit mode.
- The write burst shows the use of $\overline{\text{core_ta}}$ signal negation to delay the third data beat.
- The final read burst shows the use of $\overline{\text{core_drtry}}$ on the third data beat.
- The address for the third transfer is delayed until the first transfer completes.

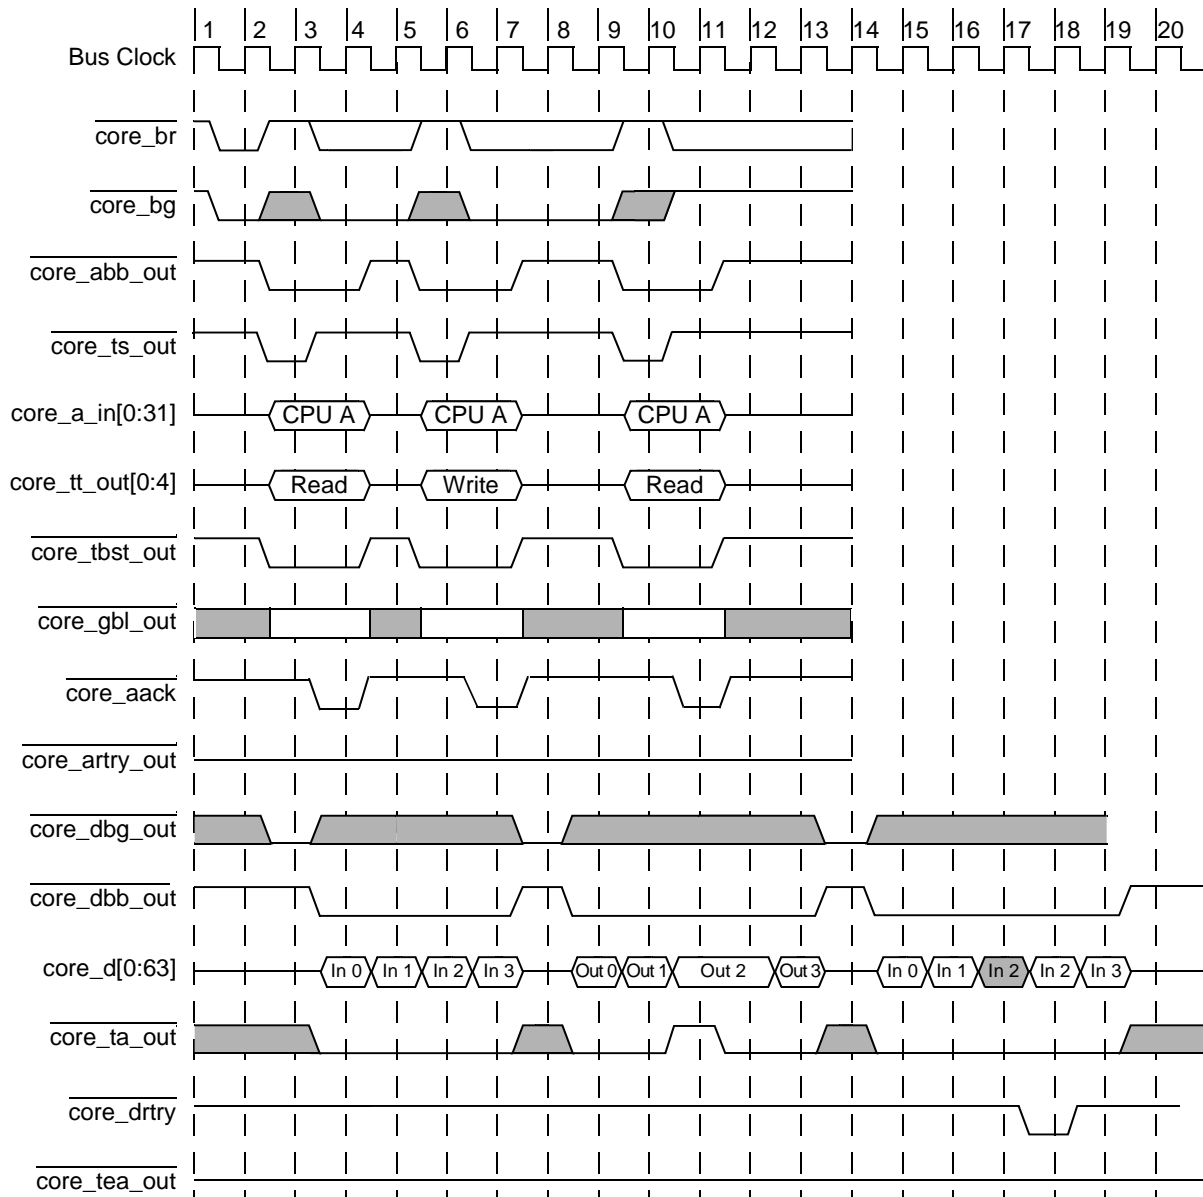


Figure 9-18. Burst Transfers with Data-Delay Controls

Figure 9-19 shows the use of the $\overline{\text{core_tea}}$ signal. Note that all bidirectional signals are three-stated between bus tenures. Note the following:

- The first data beat of the read burst (in clock 0) is the critical quad word.
- The $\overline{\text{core_tea}}$ signal truncates the burst write transfer on the third data beat.
- The G2 core eventually causes an exception to be taken on the $\overline{\text{core_tea}}$ event.

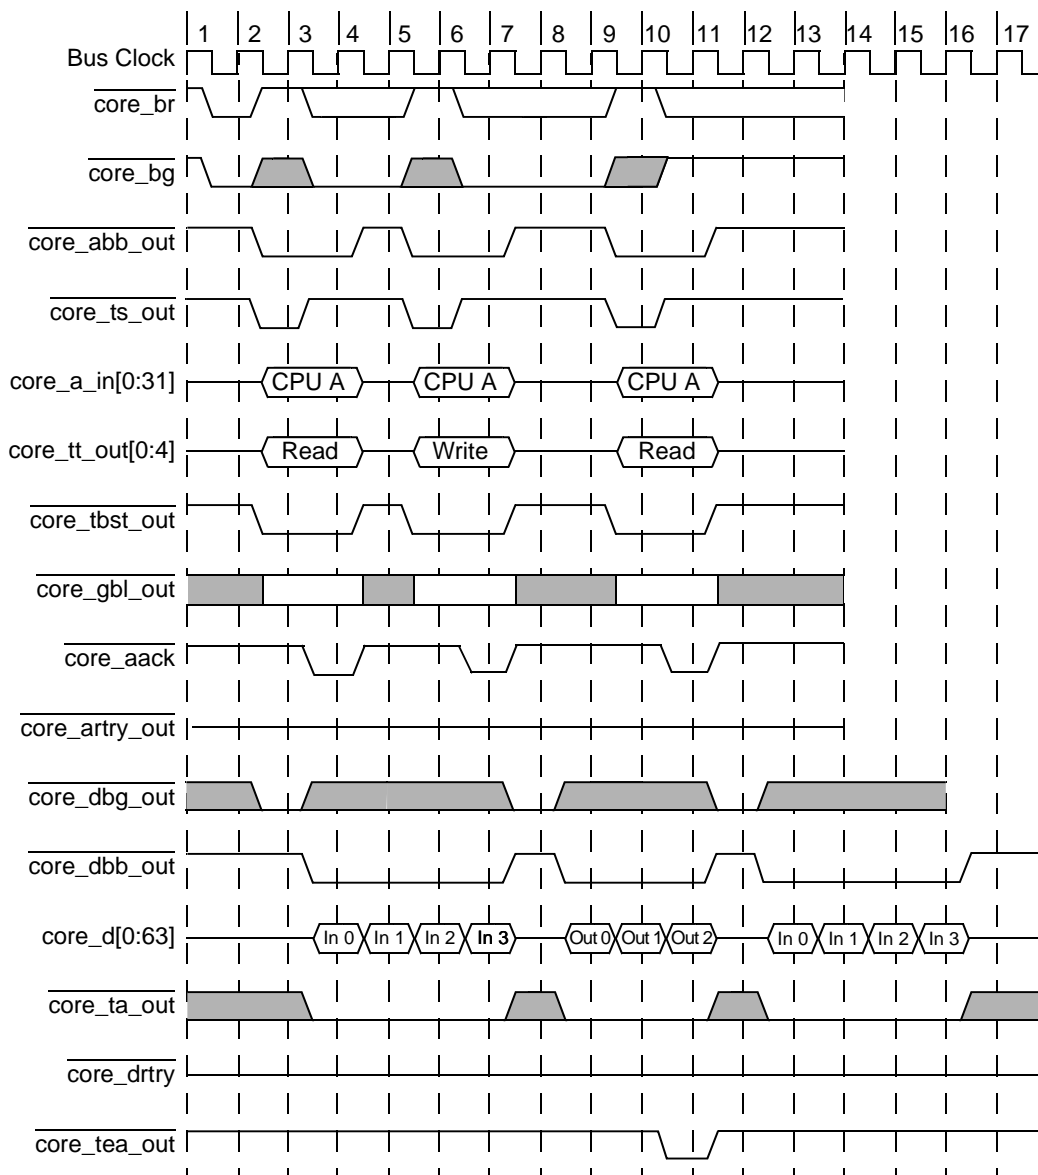


Figure 9-19. Use of Transfer Error Acknowledge (TEA)

9.6 Optional Bus Configurations

The G2 core supports the following three optional bus configurations that are selected by the assertion or negation of `core_drtry`, `core_tlbisync`, and `core_qack` during the negation of `core_hreset`.

- 32-bit data bus mode (see Section 9.6.1, “32-Bit Data Bus Mode,” for details)
- No-`core_drtry` mode (see Section 9.6.2, “No-`core_drtry` Mode,” for details)
- Reduced-pinout mode (see Section 9.6.3, “Reduced-Pinout Mode,” for details)

The operation and selection of the optional bus configurations are described in the following sections.

9.6.1 32-Bit Data Bus Mode

The G2 core supports an optional 32-bit data bus mode, which differs from the 64-bit data bus mode only in the byte lanes involved in the transfer and the number of data beats performed. When the G2 core in 32-bit data bus mode, only byte lanes 0 through 3 are used corresponding to `core_dh[0:31]` (both input and output signals) and `core_dp[0:3]` (both input and output signals). Byte lanes 4 through 7 corresponding to `core_dl[0:31]` (both input and output signals) and `core_dp[4:7]` (both input and output signals) are never used in this mode. The unused data bus signals are not sampled by the core during read operations, and they are driven low during write operations.

A data tenure in the 32-bit data bus mode takes one, two, or eight beats depending on the transfer size (see Table 9-2 for details) and the cache mode for the address. Data transactions of one or two data beats are performed for caching-inhibited load/store or write-through store operations. These transactions do not assert the `core_tbst_out` signal even though a two-beat burst may be performed (having the same `core_tbst_out` and `core_tsiz[0:2]` encodings as the 64-bit data bus mode). Single-beat data transactions are performed for bus operations of 4 bytes or less, and double-beat data transactions are performed for 8-byte operations only. The core only generates an 8-byte operation for a double-word-aligned load or store double operation to or from the floating-point GPRs. All cache-inhibited instruction fetches are performed as word (single-beat) operations.

Data transactions of eight data beats are performed for burst operations that load into or store from the core internal caches. These transactions transfer 32 bytes in the same way as in 64-bit data bus mode, asserting the `core_tbst_out` signal, and signaling a transfer size of 2 (`core_tsiz[0:2] = 0b010`).

The same bus protocols apply for arbitration, transfer, and termination of the address and data tenures in the 32-bit data bus mode as they apply to the 64-bit data bus mode. Late `core_artry_in` cancellation of the data tenure applies on the bus clock after the first data beat is acknowledged (after the first `core_ta`) for word or smaller transactions, or on the bus clock after the second data beat is acknowledged (after the second `core_ta`) for double-word

or burst operations (or coincident with respective $\overline{\text{core_ta}}$ if no $\overline{\text{core_drtry}}$ mode is selected).

An example of an 8-beat data transfer while the core is in the 32-bit data bus mode is shown in Figure 9-20.

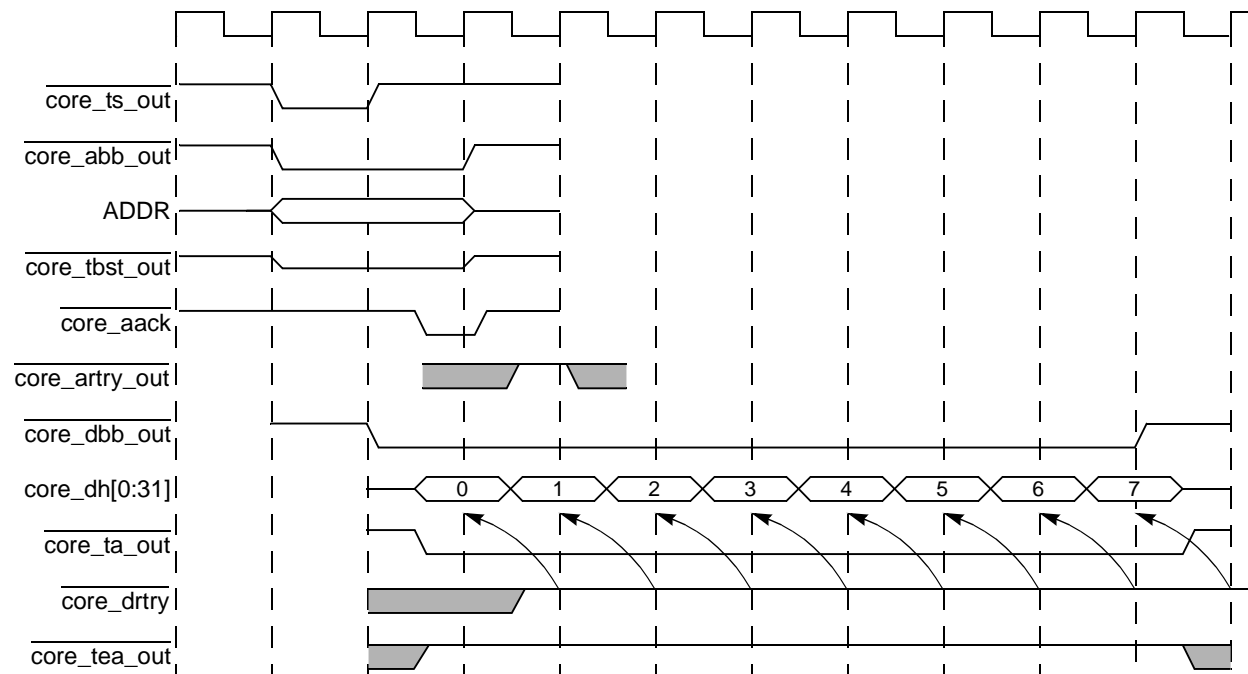


Figure 9-20. 32-Bit Data Bus Transfer (8-Beat Burst)

An example of a two-beat data transfer (with `core_drtry` asserted during each data tenure) is shown in Figure 9-21.

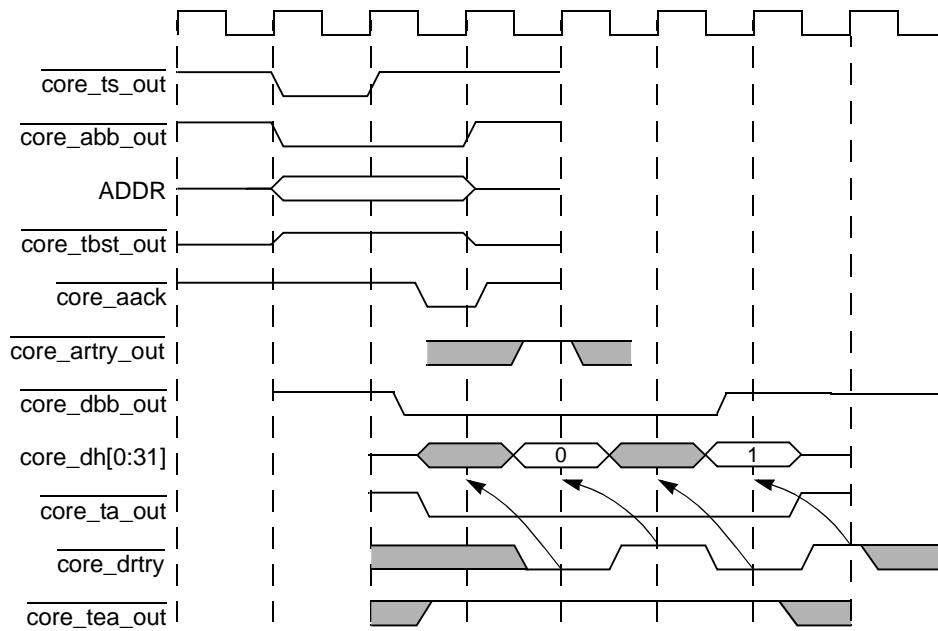


Figure 9-21. 32-Bit Data Bus Transfer (Two-Beat Burst with DRTRY)

The G2 core selects 64- or 32-bit data bus mode at startup by sampling the state of the `core_tlbisync` signal at the negation of `core_hreset`. If the `core_tlbisync` signal is negated at the negation of `core_hreset`, 64-bit data mode is entered by the core. If `core_tlbisync` is asserted at the negation of `core_hreset`, 32-bit data mode is entered.

9.6.2 No-core_drtry Mode

The G2 core supports an optional mode to disable the use of the data retry function provided through `core_drtry`. The no-`core_drtry` mode allows the forwarding of data during load operations to the processor core one bus cycle sooner than in the normal bus protocol.

The bus protocol specifies that, during load operations, the memory system can normally cancel data that was read by the master on the bus cycle after `core_ta` was asserted. This late cancellation protocol requires the core to hold any loaded data at the bus interface for one additional bus clock to verify that the data is valid before forwarding it to the processor core. For systems that do not implement the `core_drtry` function, the core provides an optional no-`core_drtry` mode that eliminates this one-cycle stall during all load operations and allows for the forwarding of data to the internal CPU immediately when `core_ta` is recognized.

When the G2 core is in no-`core_drtry` mode, data can no longer be canceled the cycle after it is acknowledged by an assertion of `core_ta`. Data is immediately forwarded to the

processor core, and any attempt at late cancellation by the system may cause improper operation by the core.

When the G2 core is following normal bus protocol, data may be canceled the bus cycle after $\overline{\text{core_ta}}$ by either of two means—late cancellation by $\overline{\text{core_drtry}}$, or late cancellation by core_artry_in . When no- $\overline{\text{core_drtry}}$ mode is selected, both cancellation cases must be disallowed in the system design for the bus protocol.

When no- $\overline{\text{core_drtry}}$ mode is selected, the system must ensure that $\overline{\text{core_drtry}}$ is not asserted to the core which may cause improper operation of the bus interface. The system must also ensure that an assertion of core_artry_in by a snooping device must occur no later than the first assertion of core_ta to the core but not on the cycle after the first assertion of core_ta .

Other than the inability to cancel data that was read by the master on the bus cycle after core_ta was asserted, the bus protocol for the core is identical to that for the basic transfer bus protocols described in this chapter, as well as for 32-bit data bus mode.

The G2 core selects the desired $\overline{\text{core_drtry}}$ mode at startup by sampling the state of the $\overline{\text{core_drtry}}$ signal itself at the negation of core_hreset . If $\overline{\text{core_drtry}}$ is negated at the negation of core_hreset , normal operation is selected. If $\overline{\text{core_drtry}}$ is asserted at the negation of core_hreset , no- $\overline{\text{core_drtry}}$ mode is selected.

9.6.3 Reduced-Pinout Mode

The G2 core provides an optional reduced-pinout mode, which idles the switching of numerous signals for reduced power consumption. Both input and output signals of the $\text{core_dl}[0:31]$, $\text{core_dp}[0:7]$, $\text{core_ap}[0:3]$, core_ape , core_dpe , and core_rsrv signals are disabled when the reduced-pinout mode is selected. Note that the 32-bit data bus mode is implicitly selected when the reduced-pinout mode is enabled.

In reduced-pinout mode, the bidirectional and output signals disabled are always driven low during the periods when they would normally have been driven by the core. The open-drain outputs (core_ape and core_dpe) are always three-stated. The bidirectional inputs are always turned-off at the input receivers of the core and are not sampled.

The G2 core selects either full-pinout or reduced-pinout mode at startup by sampling the state of the $\overline{\text{core_qack}}$ signal at the negation of core_hreset . If $\overline{\text{core_qack}}$ is asserted at the negation of core_hreset , full-pinout mode is selected by the core. If $\overline{\text{core_qack}}$ is negated at the negation of core_hreset , reduced-pinout mode is selected.

9.7 Interrupt, Checkstop, and Reset Signals

This section describes external interrupts, checkstop operations, and hard and soft reset inputs.

9.7.1 External Interrupts

Asserting the external interrupt input signals ($\overline{\text{core_int}}$, $\overline{\text{core_smi}}$, and $\overline{\text{core_mcp}}$) of the core eventually forces the processor to take an external interrupt exception, or a system management interrupt exception if the MSR[EE] is set, or the machine check interrupt if MSR[ME] and HID0[EMCP] are set.

9.7.2 Checkstops

Asserting the G2 core has two checkstop input signals— $\overline{\text{core_ckstp_in}}$ (non-maskable) and $\overline{\text{core_mcp}}$ (enabled when MSR[ME] is cleared and HID0[EMCP] is set), and a checkstop output ($\overline{\text{core_ckstp_out}}$). If $\overline{\text{core_ckstp_in}}$ or $\overline{\text{core_mcp}}$ is asserted, the core halts operations by gating off all internal clocks. The core asserts $\overline{\text{core_ckstp_out}}$ if $\overline{\text{core_ckstp_in}}$ is asserted.

If $\overline{\text{core_ckstp_out}}$ is asserted by the core, it has entered the checkstop state and processing has halted internally. The $\overline{\text{core_ckstp_out}}$ signal can be asserted for various reasons including receiving a $\overline{\text{core_tea}}$ signal and detection of external parity errors. For more information about checkstop state, see Section 5.5.2.2, “Checkstop State (MSR[ME] = 0).”

9.7.3 Reset Inputs

The G2 core has two reset inputs, described as follows:

- $\overline{\text{core_hreset}}$ (hard reset)— $\overline{\text{core_hreset}}$ is used for power-on reset sequences, or for situations in which the core must go through the entire cold-start sequence of internal hardware initializations.
- $\overline{\text{core_sreset}}$ (soft reset)—The soft reset input provides warm reset capability. This input can be used to avoid forcing the core to complete the cold start sequence.

When either reset input is negated, the processor attempts to fetch code from the system reset exception vector. The vector is located at offset 0x00100 from the exception prefix (all zeros or ones, depending on the setting of the exception prefix bit in the machine state register (MSR[IP])). The IP bit is set for $\overline{\text{core_hreset}}$.

9.7.4 Core Quiesce Control Signals

The core quiesce control signals ($\overline{\text{core_qreq}}$ and $\overline{\text{core_qack}}$) allow the processor to enter a low power state and bring bus activity to a quiescent state in an orderly fashion.

The system quiesce state is entered by configuring the processor to assert the $\overline{\text{core_qreq}}$ output. This signal allows the system to terminate or pause any bus activities that are normally snooped. When the system is ready to enter the system quiesce state, it asserts $\overline{\text{core_qack}}$. At this time, the core may enter a quiescent (low-power) state during which it stops snooping bus activity.

9.8 Processor State Signals

This section describes the G2 core support for atomic update and memory through the use of the **lwarx/stwcx**. instruction pair and includes a description of the core `core_tlbisync` input.

9.8.1 Support for the lwarx/stwcx. Instruction Pair

The Load Word and Reserve Indexed (**lwarx**) and the Store Word Conditional Indexed (**stwcx**.) instructions provide a means for atomic memory updating. Memory can be updated atomically by setting a reservation on the load and checking that the reservation is still valid before the store is performed. In the core, the reservations are made on behalf of aligned, 32-byte sections of the memory address space.

The reservation (`core_rsrv`) output signal is driven synchronously with the bus clock and reflects the status of the reservation coherency bit in the reservation address buffer (see Section 3.9, “Instruction and Data Cache Operation” for more information). See Section 8.3.11.3, “Reservation `core_rsrv`—Output,” for information about timing.

9.8.2 `core_tlbisync` Input

The `core_tlbisync` input allows for the hardware synchronization of changes to MMU tables when the core and another DMA master share the same MMU translation tables in system memory. It is asserted by a DMA master when it is using shared addresses that could be changed in the MMU tables by the core during the DMA master’s tenure.

Asserting the `core_tlbisync` input to the G2 core prevents it from completing any instructions past a **tlbsync** instruction. Generally, during the execution of an **eciwx** or **ecowx** instruction by the core, the selected DMA device should assert the `core_tlbisync` signal and keep it asserted during its DMA tenure if it is using a shared translation address. Subsequent instructions should include a **sync** and **tlbsync** instruction before any MMU table changes are performed. This prevents the core from making table changes disruptive to the other master during the DMA period.

9.9 IEEE 1149.1-Compliant Interface

The G2 core boundary-scan interface is a fully-compliant implementation of the IEEE 1149.1 standard. This section describes the core IEEE 1149.1 (JTAG) interface.

9.9.1 IEEE 1149.1 Interface Description

The G2 core has five dedicated JTAG signals (described in Table 9-11). The `core_tdi` and `core_tdo` scan ports are used to scan instructions, as well as data, into the various scan registers for JTAG operations. The scan operation is controlled by the test access port

(core_tap_en) controller, which in turn is controlled by the core_tms input sequence. The scan data is latched in at the rising edge of core_tck.

Table 9-11. IEEE Interface Pin Descriptions

Signal Name	Input/Output	Weak Pullup Provided	IEEE 1149.1 Function
core_tdi	Input	Yes	Serial scan input signal
core_tdo	Output	No	Serial scan output signal
core_tms	Input	Yes	TAP controller mode signal
core_tck	Input	Yes	Scan clock
core_trst	Input	Yes	TAP controller reset

Test reset (core_trst) is a JTAG optional signal used to reset the TAP controller asynchronously. The core_trst signal assures that the JTAG logic does not interfere with the normal operation of the chip, and can be asserted coincident with the assertion of the core_hreset signal.

The G2_LE core implements the JTAG/COP in the same manner as does the G2 core implementation with the exception of the introduction of the 33-bit Run_N counter register in which the most-significant 32 bits form a 32-bit counter. The function of the least-significant bit remains unchanged. The Run_N counter is used by the COP to control the number of processor cycles that the processor runs before halting.

9.10 Using core_dbwo (Data Bus Write Only)

The G2 core supports split-transaction pipelined transactions. It supports a limited out-of-order capability for its own pipelined transactions through the core_dbwo signal. When recognized on the clock of a qualified core_dbg, the assertion of core_dbwo directs the core to perform the next pending data write tenure (if any), even if a pending read tenure would have normally been performed because of address pipelining. The core_dbwo signal does not reorder write tenures with respect to other write tenures from the same core. It only allows that a write tenure be performed ahead of a pending read tenure from the same core.

In general, an address tenure on the bus is followed strictly in order by its associated data tenure. Transactions pipelined by the core complete strictly in order. However, the core can run bus transactions out of order only when the external system allows the core to perform a cache-line-snoop-push-out operation (or other write transaction, if pending in the core write queues) between the address and data tenures of a read operation through the use of core_dbwo. This effectively envelops the write operation within the read operation. Figure 9-22 shows how core_dbwo is used to perform an enveloped write transaction.

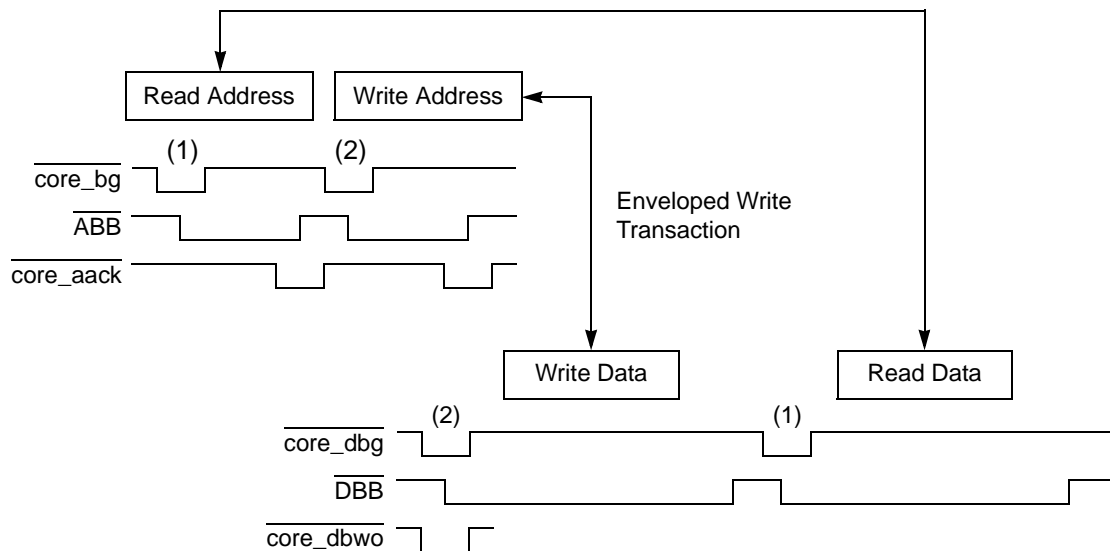


Figure 9-22. core_dbwo Transaction

Note that although the G2 core can pipeline any write transaction behind the read transaction, special care should be used when using the enveloped write feature. It is envisioned that most core implementations will not need this capability; for these applications, core_dbwo should remain negated. In cores where this capability is needed, core_dbwo should be asserted under the following scenario:

1. The G2 core initiates a read transaction (either single-beat or burst) by completing the read address tenure with no address retry.
2. Then, the G2 core initiates a write transaction by completing the write address tenure, with no address retry.
3. At this point, if core_dbwo is asserted with a qualified data bus grant to the G2 core, the G2 core asserts core_dbb_out and drives the write data onto the data bus, out of order with respect to the address pipeline. The write transaction concludes with the core negating core_dbb_out.
4. The next qualified data bus grant signals the G2 core to complete the outstanding read transaction by latching the data on the bus. This assertion of core_dbg should not be accompanied by an asserted core_dbwo.

Any number of bus transactions by other bus masters can be attempted between any of these steps.

Note the following regarding core_dbwo:

- core_dbwo can be asserted if no data bus read is pending, but it has no effect on write ordering.
- The ordering and presence of data bus writes is determined by the writes in the write queues at the time core_bg is asserted for the write address (not core_dbg). If a particular write is desired (for example, a cache-line-snoop-push-out operation),

then $\overline{\text{core_bg}}$ must be asserted after that particular write is in the queue and it must be the highest priority write in the queue at that time. A cache-line-snoop-push-out operation may be the highest priority write, but more than one may be queued.

- Because more than one write may be in the write queue when $\overline{\text{core_dbg}}$ is asserted for the write address, more than one data bus write may be enveloped by a pending data bus read.

The arbiter must monitor bus operations and coordinate the various masters and slaves with respect to the use of the data bus when core_dbwo is used. Individual core_dbg signals associated with each bus device should allow the arbiter to synchronize both pipelined and split-transaction bus organizations. Individual core_dbg and core_dbwo signals provide a primitive form of source-level tagging for the granting of the data bus.

Note that use of $\overline{\text{core_dbwo}}$ allows some operation-level tagging with respect to the G2 core and the use of the data bus.



Jsing core_dbwo (Data Bus Write Only)

Chapter 10

Power Management

The G2 core is the first processor core specifically designed for low-power operation. It provides both automatic and program-controllable power reduction modes for progressive reduction of power consumption. This chapter describes the hardware support provided by the G2 core for power management.

10.1 Overview

The G2 core has explicit power management features that are described in this chapter. Note that the design of the G2 core is fully static, allowing the internal processor core state to be preserved when no internal clock is present.

The device drivers must be modified for power management as operating systems service I/O requests by system calls to the device drivers. When a device driver is called to reduce the power of a device, it needs to be able to check the power state of the device, save the device configuration parameters, and put the device into a power-saving mode. Furthermore, every time the device driver is called, it needs to check the power status of the device and restore the device to the full-on state, if the device is in a power-saving mode.

10.2 Dynamic Power Management

Dynamic power management (DPM) automatically powers up and down the individual execution units of the G2 core, based on the contents of the instruction stream. For example, if no floating-point instructions are being executed, the floating-point unit is automatically powered down. Power is not actually removed from the execution unit; instead, each execution unit has an independent clock input, which is automatically controlled on a clock-by-clock basis. Because CMOS circuits consume negligible power when they are not switching, stopping the clock to an execution unit effectively eliminates its power consumption. The operation of DPM is completely transparent to software or any external hardware. Dynamic power management is enabled by setting `HID0[DPM]` on power-up following a hard reset sequence (`core_hreset`).

10.3 Programmable Power Modes

Hardware can enable a power management state through external asynchronous interrupts. The hardware interrupt causes the transfer of program flow to interrupt handler code. The appropriate mode is then set by the software. The G2 core provides a separate interrupt and interrupt vector for power management—the system management interrupt (`core_smi`). The G2 core also contains a decremter timer that allows it to enter the nap or doze mode for a predetermined period and then return to full power operation through the decremter interrupt exception.

The G2 core provides four power modes selectable by setting the appropriate control bits in the MSR and HID0. The four power modes are described briefly as follows:

- **Full-power**—This is the default power state of the G2 core. The G2 core is fully powered and the internal functional units are operating at the full processor clock speed. If the dynamic power management mode is enabled, functional units that are idle will automatically enter a low-power state without affecting performance, software execution, or external hardware.
- **Doze**—All the functional units of the G2 core are disabled except for the time base/decremter registers and the bus snooping logic. When the processor is in doze mode, an external asynchronous interrupt, system management interrupt, decremter exception, hard or soft reset, or machine check input (`core_mcp`) brings the G2 core into the full-power state. The core in doze mode maintains the a phase-locked loop (PLL) in a fully powered state and locked to the system external clock input (`core_sysclk`) so a transition to the full-power state takes only a few processor clock cycles.
- **Nap**—The nap mode further reduces power consumption by disabling bus snooping, leaving only the time base register and the PLL in a powered state. The core returns to the full-power state upon receipt of an external asynchronous interrupt, system management interrupt, decremter exception, hard or soft reset, or machine check input (`core_mcp`) signal. A return to full-power state from a nap state takes only a few processor clock cycles.
- **Sleep**—Sleep mode reduces power consumption to a minimum by disabling all internal functional units; then external system logic may disable the PLL and `core_sysclk`. Returning the core to the full-power state requires the enabling of the PLL and `core_sysclk`, followed by the assertion of an external asynchronous interrupt, system management interrupt, hard or soft reset, or `core_mcp` signal after the time required to relock the PLL.

Note that the G2 core cannot switch from one power management mode to another without first returning to full-on mode. The nap and sleep modes disable bus snooping; therefore, a hardware handshake using `core_qreq` and `core_qack` is provided to ensure coherency before the core enters these power management modes. Table 10-1 summarizes the four power states for the core.

Table 10-1. G2 core Programmable Power Modes

PM Mode	Functioning Units	Activation Method	Full-Power Wake-Up Method
Full power	All units active	—	—
Full power (with DPM)	Requested logic by demand	By instruction dispatch	—
Doze	<ul style="list-style-type: none"> • Bus snooping • Data cache as needed • Decrementer timer 	Controlled by SW	External asynchronous exceptions Decrementer interrupt Reset
Nap	Decrementer timer	Controlled by hardware and software	External asynchronous exceptions Decrementer interrupt Reset
Sleep	None	Controlled by hardware and software	External asynchronous exceptions Reset

10.3.1 Power Management Modes

The following sections describe the characteristics of the G2 core power management modes, the requirements for entering and exiting the various modes, and the system capabilities provided by the G2 core while the power management modes are active.

10.3.1.1 Full-Power Mode with DPM Disabled

Full-power mode with DPM disabled is selected when the DPM enable bit in HID0[DPM] is cleared. The following characteristics apply:

- Default state following power-up and $\overline{\text{core_hreset}}$
- All functional units are operating at full processor speed at all times

10.3.1.2 Full-Power Mode with DPM Enabled

Full-power mode with DPM enabled (HID0[DPM] = 1) provides on-chip power management without affecting the functionality or performance of the G2 core as follows:

- Required functional units are operating at full processor speed
- Functional units are clocked only when needed
- No software or hardware intervention required after mode is set
- Software/hardware and performance transparent

10.3.1.3 Doze Mode

Doze mode disables most functional units but maintains cache coherency by enabling the bus interface unit and snooping. A snoop hit causes the G2 core to enable the data cache, copy the data back to memory, disable the cache, and fully return to the doze mode.

Doze mode is characterized by the following features:

- Most functional units disabled
- Bus snooping and time base/decrementer still enabled
- PLL running and locked to internal core_sysclk

To enter the doze mode, the following conditions must occur:

- Set doze bit (HID0[8] = 1), MSR[POW] is set
- G2 core enters doze mode after several processor clocks

To return to full-power mode, the following conditions must occur:

- Assert internal $\overline{\text{core_int}}$, $\overline{\text{core_smi}}$, or $\overline{\text{core_mcp}}$ signals or decremter interrupts
- Hard reset or soft reset
- Transition to full-power state occurs only after a few processor cycles

10.3.1.4 Nap Mode

The nap mode disables the G2 core except for the processor PLL and time base/decrementer. The time base can be used to restore the core to a full-on state after a specified period.

Because bus snooping is disabled for nap and sleep mode, a hardware handshake using the quiesce request ($\overline{\text{core_qreq}}$) and quiesce acknowledge ($\overline{\text{core_qack}}$) signals are required to maintain data coherency. The G2 core asserts the $\overline{\text{core_qreq}}$ signal to indicate that it is ready to disable bus snooping, including all bus activity. Once the processor has entered a quiescent state, it no longer snoops bus activity.

When the system logic has ensured that snooping is no longer necessary, it allows the processor to enter the nap (or sleep) mode and causes the assertion of the G2 core $\overline{\text{core_qack}}$ input signal for the duration of the nap mode period.

Nap mode is characterized by the following features:

- Time base/decrementer still enabled
- Most functional units disabled (including bus snooping)
- PLL running and locked to internal core_sysclk

To enter the nap mode, the following conditions must occur:

- Set nap bit (HID0[9] = 1), MSR[POW] bit is set
- G2 core asserts $\overline{\text{core_qreq}}$
- System asserts $\overline{\text{core_qack}}$
- The processor core enters nap mode after several processor clocks

To return to full-power mode, one of the following conditions must occur:

- Assert $\overline{\text{core_int}}$, $\overline{\text{core_smi}}$, or $\overline{\text{core_mcp}}$ internal signals
- Decrementer exception interrupt
- Hard reset or soft reset

Transition to full-power takes only a few processor cycles. $\overline{\text{core_qack}}$ can remain asserted; however, $\overline{\text{core_qreq}}$ negates before any bus transaction begins.

10.3.1.5 Sleep Mode

Sleep mode consumes the least amount of power of the four modes, since all functional units are disabled. To conserve the maximum amount of power, the PLL and internal core_sysclk signals can be disabled. Due to the fully static design of the G2 core, the internal processor state is preserved when no internal clock is present. Because the time base and decrementer are disabled while the G2 core is in sleep mode, the time base contents must be updated from an external time base following sleep mode, if accurate time-of-day maintenance is required.

Before entering sleep mode, the G2 core asserts $\overline{\text{core_qreq}}$ to indicate that it is ready to disable bus snooping. When the system has ensured that snooping is no longer necessary, the system logic allows the G2 core to enter sleep mode by asserting $\overline{\text{core_qack}}$ for the duration of the sleep mode period.

Sleep mode is characterized by the following features:

- All functional units disabled (including bus snooping and time base)
- All nonessential input receivers disabled
- Internal clock regenerators disabled
- PLL and core_sysclk can be disabled

To enter sleep mode, the following conditions must occur:

- Set sleep bit ($\text{HID0}[10] = 1$), $\text{MSR}[\text{POW}]$ is set
- G2 core asserts $\overline{\text{core_qreq}}$
- System logic asserts $\overline{\text{core_qack}}$
- G2 core enters sleep mode after several processor clocks

To return to full-power mode, the following conditions must occur:

- Assert $\overline{\text{core_int}}$, $\overline{\text{core_smi}}$, or $\overline{\text{core_mcp}}$ internal signals
- Hard reset or soft reset

To return to full-power mode after PLL and core_sysclk are disabled in sleep mode, the following conditions must occur:

- Enable core_sysclk

- Reconfigure PLL into desired processor clock mode
- System logic waits for PLL startup and relock time (100 μ sec)
- System logic asserts one of the sleep recovery signals (for example, $\overline{\text{core_int}}$ or $\overline{\text{core_smi}}$)

10.3.2 Power Management Software Considerations

Because the G2 core is a dual issue processor core with out-of-order execution capability, care must be taken in how the power management modes are entered. Furthermore, nap and sleep modes require all outstanding bus operations to be completed before the power management mode is entered. Section 10.4, “Example Code Sequence for Entering Processor Sleep Mode,” provides an example software sequence for putting the G2 core into sleep mode.

Normally, during system configuration time, one of the power management modes would be selected by setting the appropriate H1D0 mode bit. Later on, the power management mode is invoked by setting MSR[POW]. To ensure a clean transition into and out of the power management mode, set MSR[EE] (external interrupt enable) and execute the following code sequence:

```

sync
mtmsr[POW = 1]
isync
loop: b loop

```

10.4 Example Code Sequence for Entering Processor Sleep Mode

The following is a sample code sequence for entering G2 core sleep mode.

```

*****
# set up G2 core H1D0 power management bits
#*****
#*****processor HID and external interrupt initialization*****
#
# set up HID registers for the various processors of this family
# hid setup taken from minix's mpxPowerPC.s

    mfspr    r31, pvr          # pvr reg
    srwi    r31, r31, 16
resetTest603:
    cmpi    0, 0, r31, 3
    bne     cr0, endHIDSetup

    addi    r0, r0, 0
    oris    r0, r0, 0x8000    # enable machine check pin EMCP
    oris    r0, r0, 0x0010    # enable dynamic power mgmt DPM
    oris    r0, r0, 0x0020    # enable SLEEP power mode
    ori     r0, r0, 0x8000    # enable the Icache ICE

```



Example Code Sequence for Entering Processor Sleep Mode

```
    ori    r0, r0, 0x4000    # enable the Dcache DCE
    ori    r0, r0, 0x0800    # invalidate Icache ICFI
    ori    r0, r0, 0x0400    # invalidate Dcache DCFI
    mtspr  hid0, r0
    isync

#*****
# then when the processor is in a loop, force an SMI interrupt
#*****

.orig 0x00001400                # System Management Interrupt

# force big-endian mode
    stw    r0,0x05f8,r0        # need nop every second inst.
    stw    r0,0x05fc,r0
    mfmsr  r0
    ori    r0,r0,r0
    ori    r0,r0,0x0001        # force big-endian LE bit
    ori    r0,r0,r0
    xori   r0,r0,0x0001        # force big-endian LE bit
    ori    r0,r0,r0
    mtmsr  r0
    ori    r0,r0,r0
    isync
    ori    r0,r0,r0

# save off additional registers to be corrupted
    stw    r20,0x05f4,r0
    mfspr  r21, srr0           # put srr0 in r21
    stw    r21,0x05f0,r0      # put r21 in 0x05f0
    mfspr  r22, srr1           # put srr1 in r22
    stw    r22,0x05ec,r0      # put r22 in 0x05ec
    stw    r23,0x05e8,r0
    mfcr   r23
    stw    r23,0x05e4,r0
    xor    r0,r0,r0

#*****
# set msr pow bit to go into sleep mode

    sync
    mfmsr  r5                  # get MSR
    addis  r3, r0, 0x0004      # turn on POW bit
    ori    r3, r3, 0x0000      # turn on ME bit 19
    or     r5, r3, r5
    mtmsr  r5
    isync

    addis  r20, r0, 0x0000
    ori    r20, r20, 0x0002

stay_here:
    addic. r20, r20, -1        # subtract 1 from r20 and set cc
    bgt    cr0, stay_here     # loop if positive

# restore corrupted registers
    lwz    r23,0x05e4,r0
```



Example Code Sequence for Entering Processor Sleep Mode

```
        mtcrf        0xff,r23
        lwz         r23,0x05e8,r0
        lwz         r22,0x05ec,r0
        mtspr       srr1, r22
        lwz         r21,0x05f0,r0
        mtspr       srr0, r21
        lwz         r20,0x05f4,r0
        lwz         r0,0x05fc,r0

        sync
        rfi

#*****
# to get out of sleep mode, do a Soft Reset
#*****

.orig 0x00000100                # Reset handler in low memory

# force big-endian mode
        stw         r0,0x05f8,r0        # need nop every second inst.
        stw         r0,0x05fc,r0
        mfmsr       r0
        ori         r0,r0,r0
        ori         r0,r0,0x0001        # force big-endian LE bit
        ori         r0,r0,r0
        xori        r0,r0,0x0001        # force big-endian LE bit
        ori         r0,r0,r0
        mtmsr       r0
        ori         r0,r0,r0
        isync
        ori         r0,r0,r0

# save off additional registers to be corrupted
        stw         r20,0x05f4,r0
        stw         r21,0x05f0,r0
        stw         r22,0x05ec,r0
        stw         r23,0x05e8,r0
        mfcr        r23
        stw         r23,0x05e4,r0
        xor         r0,r0,r0

# restore corrupted registers
        lwz         r23,0x05e4,r0
        mtcrf       0xff,r23
        lwz         r23,0x05e8,r0
        lwz         r22,0x05ec,r0
        lwz         r21,0x05f0,r0
        lwz         r20,0x05f4,r0
        lwz         r0,0x05fc,r0

        sync
        rfi
#*****
```

Chapter 11

Debug Features

This chapter describes the debug features of the PowerPC architecture with respect to the G2_LE core. Both the G2 and G2_LE include the trace facility debug feature. However, the G2_LE core has improved debug capability by enhancing the JTAG/COP interface. The enhanced debug features are described as follows:

- Addition of three breakpoint registers
- Inclusions of watchpoint/breakpoint indication signals— $\overline{\text{core_iabr}}$, $\overline{\text{core_iabr2}}$, $\overline{\text{core_dabr}}$, and $\overline{\text{core_dabr2}}$
- Addition of COP_SVR instruction
- New force-single-step operation instruction

11.1 Breakpoint Facilities

The G2_LE core provides enhanced debug facilities—instruction address breakpoint, data address breakpoint, and program single stepping to enable software debug events. The existing IABR and single-step functions are facilitated by the new debug features. The debug facilities consist of a set of debug control registers (DBCR, IBCR), a set of instruction address breakpoint registers (IABR, IABR2), and a set of data address breakpoint registers (DABR, DABR2). The basic operation of the DABRs are similar to that of the MPC750 processor. For information on the MPC750, see the *MPC750 RISC Microprocessor Family User's Manual*. These registers are used together to enable various breakpoint functions.

These registers are accessible to only supervisor-level programs by the **mf spr** and **mt spr** instructions. The SPR address for the registers can be found in Table 3-33 of Chapter 3, “Instruction Set Model.”

11.1.1 Instruction Address Breakpoint Registers (IABR, IABR2)

IABR and IABR2 can be used to cause a breakpoint exception if a specified instruction address is encountered. The IABR and IABR2 control the instruction address breakpoint exception. IABR[CEA] and IABR2[CEA] hold the effective address to which each

instruction's address is compared. The exception is enabled by setting IABR[BE]. The exception is taken when there is an instruction address breakpoint match on the next instruction to complete. The instruction tagged with the match cannot complete before the instruction address breakpoint exception (0x01300) is taken. The address of the instruction that matches the breakpoint condition is stored in SRR0. The tagged instruction retires after returning from the exception (**rfi** or **rfti**). The results are then committed to the destination registers and address.

If the IABR or IABR2 values are set to any exception vector range, an unrecoverable state occurs. The IABR or IABR2 values should never be set to match within the instruction address breakpoint exception handler. Failure to prohibit a breakpoint within any handler may result in an indeterminate or unrecoverable processor state. See Section 2.1.2.14, "Instruction Address Breakpoint Registers (IABR and IABR2)," for bit descriptions.

11.1.2 Instructional Address Control Register (IBCR)

IBCR is a supervisor-level SPR. It controls the compare and match type conditions for IABR and IABR2. Note that IABR and IABR2 must be enabled before the effects of IBCR are realized. See Section 2.1.2.14.1, "Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only," for bit descriptions.

11.1.3 Data Address Breakpoint Registers (DABR, DABR2)

The DABR and DABR2 registers are used to cause a breakpoint exception if the specified address is encountered. DABR[CEA] and DABR2[CEA] hold an effective address to which each address of data access is compared. The breakpoint translation matches when a data address breakpoint matches ($MSR[DR] = DABR[BT]$). The data address write and data address read exceptions are enabled by setting DABR[WBE,RBE] and DABR2[WBE,RBE]. The data tagged with the match does not complete before the breakpoint exception is taken.

The DSI exception (0x00300) occurs when there is a data address breakpoint match. The DSI exception is taken before the load or store instruction is executed. When the exception is taken, DAR is set to the data address that causes the breakpoint and DSISR[DABR] is set to indicate a data address breakpoint. The address of the instruction associated with the breakpoint condition is stored in SRR0. The instruction retires after returning from the DSI exception, and all registers and memory accesses are committed to memory.

An unrecoverable state occurs whenever DABR or DABR2 values are set to an exception vector. These values must not be set to match within the DSI exception handler or the G2_LE core may enter an indeterminate or unrecoverable processor core state.

11.1.4 Data Address Control Register (DBCR)

DBCR is a supervisor-level SPR on the G2_LE core that controls the compare type and match type conditions for DABR and DABR2. Note that DABR or DABR2 or both breakpoint registers must be enabled before the effects of DBCR are realized. See Section 2.1.2.15.1, “Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only,” for bit descriptions.

11.1.5 Other Debug Resources

In addition to the four breakpoint registers and the two breakpoint control registers, other internal register values control and monitor the effects of breakpoint conditions. Table 11-1 shows these registers and their bits.

Table 11-1. Other Debug and Support Register Bits

Register	Bits	Name	Description
MSR	17	PR	Privilege level. Breakpoint registers can only be accessed when this bit is cleared (supervisor mode).
	21	SE	Single-step trace enable. 0 The processor executes instructions normally. 1 The processor generates a trace exception upon the successful completion of the next instruction.
	22	BE	Branch trace enable 0 The processor executes branch instructions normally. 1 The processor generates a trace exception upon the successful completion of a branch instruction.
HID0	0–31	—	See Table 2-5 for details.
DAR	0–31	—	Data address register. DAR is loaded with the effective address of a data breakpoint condition that matches.
DSISR	9	DABR	Set if DABR exception occurs.

11.1.6 Software Debug Features

Software programming model interface controls debug features including instruction and data breakpoints. When an instruction or data address breakpoint register is enabled and the conditions are met, an instruction address breakpoint exception (0x01300) or DSI exception (0x00300) occurs.

The cause of a DSI exception can be determined by the setting of DSISR[DABR]. A data address breakpoint exception occurs when the data in DABR[BT] or DABR2[BT] matches the next data access (load or store instruction) to complete in the completion unit (see Section 2.1.2.15.1, “Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only,” for more details). The DAR contains the address of the matching data address breakpoint determined by DABR, DABR2, and DBCR.

IABR[BE] and IABR2[BE] enable and control instruction address breakpoint, and IBCR controls match conditions (see Section 2.1.2.14.1, “Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only,” for more details). DABR[29–31] and DABR2[29–31] enable and control data address breakpoint and DBCR controls match conditions.

When MSR[SE] (single-step trace enable) is set, the processor core generates a trace exception (0x00D00) upon the successful completion of the next instruction. When the MSR[BE] (branch trace enable) is set, the processor core generates a trace exception (0x00D00) upon the successful completion of a branch instruction.

11.2 Expanded Debugging Facilities in Breakpoint Registers

Breakpoint, single-step, and branch trace enable, address and combinational matching are additional debugging facilities provided by the breakpoint registers (DABR, DABR2, IABR, and IABR2).

11.2.1 Breakpoint Enabled

When an instruction address breakpoint is set, and a condition is matched, an instruction address breakpoint exception (0x01300) occurs along with executing the matched instruction. The instruction retires after it has returned from the exception. When a data and a condition are matched, a DSI exception (0x00300) occurs along with executing the matched instruction. The instruction retires after it has returned from the exception and has updated all memory or registers.

11.2.2 Single-Step Enabled

Single-stepping can be a very useful tool in software debugging. This debug feature executes one instruction before it takes a trace exception. In trace exception, the result is being examined after that one instruction has executed.

When MSR[SE] (single-step trace enable) is set, the processor generates a trace exception (0x00D00) upon the successful completion of the next instruction. A trace exception is not taken for an **isync**, **sync**, **rfi**, **rfdi**, or **mtmsr** instructions. If softstop or hardstop is enabled, and MSR[SE] bit is set, the machine will stop before the present instruction is retired and not take a trace exception.

MSR can be set by using **mtmsr** or by setting the SRR0 bit corresponding to MSR[SE] before returning from an interrupt. If the SRR0 is set after returning from the interrupt, single-step is enabled by executing one instruction along with taking the trace exception.

A typical software debugging procedure is to set a instruction address breakpoint at the instruction address to be single stepped. When the IABR exception is taken, the exception

routine should disable the instruction address breakpoint and set SRR0 to set the MSR[SE] on the **rfi**. The trace exception is taken after the IABR exception is executed. In any exception, the value of MSR is saved in SRR0. MSR[SE] is no longer set along with single-step is disabled. Finally, the trace exception examines the result through a routine and sets SRR0 to disable MSR[SE] on the **rfi** or to execute the next instruction.

Single-stepping skips **isync**, **sync**, **rfi**, **rfdi**, and branch instructions because they do not enter the instruction pipeline. The branch trace may be used for **rfi**, **rfdi**, and branch instructions.

Also, single-step debugging condition over a **mtmsr** may give unwanted results. Once MSR is updated, single stepping may be disabled and the G2 core continues executing instructions without this debugging conditions. Thus, it is recommended to disable and enable MSR[SE] by using SRR0 within an interrupt. Therefore, **rfdi** is responsible for setting or configuring MSR[SE].

11.2.3 Branch Trace Enabled

When MSR[BE] (branch trace enable) is set, the processor generates a trace exception (0x00D00) upon the successful completion of a branch instruction. If softstop or hardstop is enabled, and MSR[SE] is set, the machine stops before the present instruction is retired, and does not take a trace exception.

11.2.4 Address Matching

On G2 and G2_LE a match occurs when an address equals to an effective address in a breakpoint register. The G2_LE can match addresses on greater than or equal or less than as an additional matching condition for IBCR and DBCR.

11.2.5 Combinational Matching

An address match can be signaled after an OR function of the two compared addresses match or the AND of the two addresses match, depending on the setting of IBCR and DBCR associated with the enabled breakpoint registers. This feature along with matching on greater than and less than allows a breakpoint to be set inside or outside a range of two addresses. The instruction address breakpoints and data address breakpoints always work independently of each other. For more details, see Section 2.1.2.14.1, “Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only” and Section 2.1.2.15.1, “Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only.”

11.3 Watchpoint Signaling

There is a mechanism to enable address matching but it also disables the signaling of an exception on a softstop. This allows observing address matching on the watchpoint signals

only. Watchpoint signals allow external observability of breakpoint matches and address matching is the output to the watchpoint signals (`core_iabr`, `core_iabr2`, `core_dabr`, and `core_dabr2`). These four watchpoint signals are asserted for at least one bus clock cycle. When DBCR and IBCR are configured for combinational signal type OR, the watchpoint signals—`core_iabr`, `core_iabr2` and `core_dabr`, `core_dabr2`—reflect their respective breakpoints. When DBCR and IBCR are configured for combinational signal type AND, only the `core_iabr2` and `core_dabr2` watchpoint signals are asserted when the AND condition is met. IBCR[DNS] and DBCR[DNS] inhibit the signal transition on the core signal pins.

11.4 Exception Vectors and Priority

Table 11-2 lists exception vectors which are associated with debug and breakpoint events. Breakpoint events do not change other exception vectors and conditions

Table 11-2. Related Debug Exceptions and Conditions

Exception Type	Vector Offset	Causing Condition
Data access	00300	A data address breakpoint exception occurs when a match condition exists for the effective address of the data access in either DABR or DABR2 for the next read or write data access, and WBE and RBE, DABR enable bits are set for read or write, respectively. A data breakpoint event is determined by setting DSISR[DABR], which causes a data access exception. The DAR contains the address of the breakpoint match condition.
Trace	00D00	The trace exception is taken when MSR[SE] = 1 or when the currently completing instruction is a branch instruction and MSR[BE] = 1.
Instruction address breakpoint	01300	An instruction address breakpoint exception occurs when a match condition exists for the effective address of the instruction access in either IABR or IABR2 for the next instruction to complete in the completion unit, and WBE, IABR enable bit is set.

11.5 Instruction Address Breakpoint Examples

The address matching for the instruction address breakpoint register has the following four possible conditions for the specific register signals:

- Instruction's effective address = IABR_ADDR (value in IABR[CEA])
- Instruction's effective address = IABR_ADDR OR instruction's effective address = IABR2_ADDR (value in IABR2[CEA])
- IABR_ADDR < instruction's effective address < IABR2_ADDR
- Instruction's effective address < IABR_ADDR OR instruction's effective address > IABR2_ADDR

Table 11-3 describes the instruction address breakpoint register for a single address matching conditions.

Table 11-3. Single Address Matching (G2 Core Emulation)

Register Field Name	Condition	Register Field Name	Condition
IABR[CEA]	IABR_ADDR	IABR2[CEA]	—
IABR[BE]	1	IABR2[BE]	0
IBCR[CNT]	0	—	—
IBCR[SIG_TYPE]	OR	—	—
IBCR[CMP1]	=	IBCR[CMP2]	—

This matches when the instruction’s effective address = IABR_ADDR.

Table 11-4 describes the instruction address breakpoint registers when an address can match one or the other possible addresses (an OR condition).

Table 11-4. Two Addresses OR Matching

Register Field Name	Condition	Register Field Name	Condition
IABR[CEA]	IABR_ADDR	IABR2[CEA]	IABR2_ADDR
IABR[BE]	1	IABR2[BE]	1
IBCR[CNT]	0	—	—
IBCR[SIG_TYPE]	OR	—	—
IBCR[CMP1]	=	IBCR[CMP2]	=

This matches when the instruction’s effective address = IABR_ADDR OR the instructions effective address = IABR2_ADDR.

Table 11-5 describes the instruction address breakpoint register for an address matching inside an address range condition.

Table 11-5. Address Matching for Inside Address Range

Register Field Name	Condition	Register Field Name	Condition
IABR[CEA]	IABR_ADDR	IABR2[CEA]	IABR2_ADDR
IABR[BE]	1	IABR2[BE]	1
IBCR[CNT]	0	—	—
IBCR[SIG_TYPE]	AND	—	—
IBCR[CMP1]	>	IBCR[CMP2]	<

This matches when $IABR_ADDR \leq \text{instruction's effective address} < IABR2_ADDR$.

Table 11-6 describes the instruction address breakpoint register for an address matching outside an address range condition.

Table 11-6. Address Matching for Outside Address Range

Signal	Condition	Signal	Condition
IABR[CEA]	IABR_ADDR	IABR2[CEA]	IABR2_ADDR
IABR[BE]	1	IABR2[BE]	1
IBCR[CNT]	0	—	—
IBCR[SIG_TYPE]	OR	—	—
IBCR[CMP1]	<	IBCR[CMP2]	>

This matches when the instruction’s effective address < IABR_ADDR OR the instruction’s effective address ≤ IABR2_ADDR.

The breakpoint match can be observed externally through watchpoint signals. When DCBR[SIG_TYPE] or ICBR[SIG_TYPE] is cleared for an OR signal type, the watchpoint signals—core_iabr, core_iabr2, core_dabr, and core_dabr2—reflect their respective breakpoints. When DCBR[SIG_TYPE] or ICBR[SIG_TYPE] is set for an AND signal type, the watchpoint signals core_iabr2 and core_dabr2 are asserted when the AND condition is met. The watchpoint signal has to asserted for at least one bus clock cycle. For more details, see Section 2.1.2.14.1, “Instruction Address Breakpoint Control Registers (IBCR)—G2_LE Only,” and Section 2.1.2.15.1, “Data Address Breakpoint Control Registers (DBCR)—G2_LE-Only.”

11.6 Synchronization Requirements

An **isync** instruction must follow the setting of the **mtspr** of the breakpoint related registers, MSR, HID0, IABR, IABR2, DABR, DABR2, IBCR, and DBCR to ensure that the breakpoint condition is set. IBCR and DBCR should be set before enabling the breakpoint. The breakpoint should be cleared before changing bits in the IBCR and DCBR. For more details, see Section 5.5.16, “Instruction Address Breakpoint Exception (0x01300).”

An unrecoverable state occurs at anytime if one of the register values of IABR, IABR2, DABR, and DABR2 are set to an exception vector. The IABR or IABR2 values must not be set to match within the instruction address breakpoint exception handler. The DABR or DABR2 values must not be set to the DSI exception handler. Failure to prohibit a breakpoint within the instruction address breakpoint exception or DSI handler may result an unrecoverable and indeterminate processor core state.

If an IABR match and DABR match occur on the same instruction, the instruction address breakpoint exception is taken before the DSI exception.

If an IABR match occurs on a branch instruction, the instruction address breakpoint exception is set to the effective address of the branch instruction.

Appendix A

PowerPC Instruction Set Listings

This appendix lists the G2 core microprocessor’s instruction set as well as the additional PowerPC instructions not implemented in the G2. Instructions are sorted by mnemonic, opcode, function, and form. Also included is a quick reference table that contains general information, such as the architecture level, privilege level, and form, and indicates if the instruction is 64-bit and optional.

Note that split fields representing the concatenation of sequences from left to right, are shown in lowercase. For more information refer to Chapter 8, “Instruction Set,” in the *Programming Environments Manual*.

The following key applies to the tables in this appendix.

Key: Reserved Bits Instruction Not Implemented in the G2 Core

A.1 Instructions Sorted by Mnemonic

Table A-1 lists the instructions implemented in the PowerPC architecture in alphabetical order by mnemonic.

Table A-1. Complete Instruction List Sorted by Mnemonic

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
addx	31				D					A					B			OE											Rc
addcx	31				D					A					B			OE											Rc
addex	31				D					A					B			OE											Rc
addi	14				D					A																			
addic	12				D					A																			
addic.	13				D					A																			
addis	15				D					A																			
addmex	31				D					A					0	0	0	0	0	0									Rc
addzex	31				D					A					0	0	0	0	0	0									Rc
andx	31				S					A					B														Rc
andcx	31				S					A					B														Rc
andi.	28				S					A																			
andis.	29				S					A																			



Table A-1. Complete Instruction List Sorted by Mnemonic (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
bx	18	LI														AA	LK											
bcx	16	BO				BI				BD								AA	LK									
bcctrx	19	BO				BI				0 0 0 0 0				528				LK										
bclrx	19	BO				BI				0 0 0 0 0				16				LK										
cmp	31	crfD	0	L	A				B				0				0											
cmpi	11	crfD	0	L	A				SIMM																			
cmpl	31	crfD	0	L	A				B				32				0											
cmpli	10	crfD	0	L	A				UIMM																			
cntlzdx¹	31	S				A				0 0 0 0 0				58				Rc										
cntlzwx	31	S				A				0 0 0 0 0				26				Rc										
crand	19	crbD				crbA				crbB				257				0										
crandc	19	crbD				crbA				crbB				129				0										
creqv	19	crbD				crbA				crbB				289				0										
crnand	19	crbD				crbA				crbB				225				0										
crnor	19	crbD				crbA				crbB				33				0										
cror	19	crbD				crbA				crbB				449				0										
crorc	19	crbD				crbA				crbB				417				0										
crxor	19	crbD				crbA				crbB				193				0										
dcbf	31	0 0 0 0 0				A				B				86				0										
dcbi²	31	0 0 0 0 0				A				B				470				0										
dcbst	31	0 0 0 0 0				A				B				54				0										
dcbt	31	0 0 0 0 0				A				B				278				0										
dcbtst	31	0 0 0 0 0				A				B				246				0										
dcbz	31	0 0 0 0 0				A				B				1014				0										
divdx¹	31	D				A				B				OE	489				Rc									
divdux¹	31	D				A				B				OE	457				Rc									
divwx	31	D				A				B				OE	491				Rc									
divwux	31	D				A				B				OE	459				Rc									
eciwx	31	D				A				B				310				0										
ecowx	31	S				A				B				438				0										
eieio	31	0 0 0 0 0				0 0 0 0 0				0 0 0 0 0				854				0										
eqvx	31	S				A				B				284				Rc										
extsbx	31	S				A				0 0 0 0 0				954				Rc										
extshx	31	S				A				0 0 0 0 0				922				Rc										
extswx¹	31	S				A				0 0 0 0 0				986				Rc										
fabsx	63	D				0 0 0 0 0				B				264				Rc										
faddx	63	D				A				B				0 0 0 0 0				21	Rc									

Table A-1. Complete Instruction List Sorted by Mnemonic (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
faddsx	59	D			A			B			00000			21			Rc											
fcfid_x¹	63	D			00000			B			846			Rc														
fcmpo	63	crfD	00		A			B			32			0														
fcmpu	63	crfD	00		A			B			0			0														
fctid_x¹	63	D			00000			B			814			Rc														
fctidz_x¹	63	D			00000			B			815			Rc														
fctiw_x	63	D			00000			B			14			Rc														
fctiwz_x	63	D			00000			B			15			Rc														
fdiv_x	63	D			A			B			00000			18	Rc													
fdivs_x	59	D			A			B			00000			18	Rc													
fmadd_x	63	D			A			B			C			29	Rc													
fmaddsx	59	D			A			B			C			29	Rc													
fmr_x	63	D			00000			B			72			Rc														
fmsub_x	63	D			A			B			C			28	Rc													
fmsubs_x	59	D			A			B			C			28	Rc													
fmul_x	63	D			A			00000			C			25	Rc													
fmuls_x	59	D			A			00000			C			25	Rc													
fnabs_x	63	D			00000			B			136			Rc														
fneg_x	63	D			00000			B			40			Rc														
fnmadd_x	63	D			A			B			C			31	Rc													
fnmaddsx	59	D			A			B			C			31	Rc													
fnmsub_x	63	D			A			B			C			30	Rc													
fnmsubs_x	59	D			A			B			C			30	Rc													
fres_x³	59	D			00000			B			00000			24	Rc													
frsp_x	63	D			00000			B			12			Rc														
frsqrt_x³	63	D			00000			B			00000			26	Rc													
fsel_x³	63	D			A			B			C			23	Rc													
fsqrt_x³	63	D			00000			B			00000			22	Rc													
fsqrts_x³	59	D			00000			B			00000			22	Rc													
fsub_x	63	D			A			B			00000			20	Rc													
fsubs_x	59	D			A			B			00000			20	Rc													
icbi	31	00000			A			B			982			0														
isync	19	00000			00000			00000			150			0														
lbz	34	D			A			d																				
lbzu	35	D			A			d																				
lbzux	31	D			A			B			119			0														
lbzx	31	D			A			B			87			0														



Table A-1. Complete Instruction List Sorted by Mnemonic (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
ld ¹	58	D			A			ds										0										
ldarx ¹	31	D			A			B			84										0							
ldu ¹	58	D			A			ds										1										
ldux ¹	31	D			A			B			53										0							
ldx ¹	31	D			A			B			21										0							
lfd	50	D			A			d																				
lfd	51	D			A			d																				
lfd	31	D			A			B			631										0							
lfdx	31	D			A			B			599										0							
lfs	48	D			A			d																				
lfsu	49	D			A			d																				
lfsux	31	D			A			B			567										0							
lfsx	31	D			A			B			535										0							
lha	42	D			A			d																				
lhau	43	D			A			d																				
lhau	31	D			A			B			375										0							
lhax	31	D			A			B			343										0							
lhbrx	31	D			A			B			790										0							
lhz	40	D			A			d																				
lhzu	41	D			A			d																				
lhzux	31	D			A			B			311										0							
lhzx	31	D			A			B			279										0							
lmw ⁴	46	D			A			d																				
lswi ⁴	31	D			A			NB			597										0							
lswx ⁴	31	D			A			B			533										0							
lwa ¹	58	D			A			ds										2										
lwarx	31	D			A			B			20										0							
lwaux ¹	31	D			A			B			373										0							
lwax ¹	31	D			A			B			341										0							
lwbrx	31	D			A			B			534										0							
lwz	32	D			A			d																				
lwz	33	D			A			d																				
lwzux	31	D			A			B			55										0							
lwzx	31	D			A			B			23										0							
mcrf	19	crfD		00		crfS		00		00000						0		0										
mcrfs	63	crfD		00		crfS		00		00000						64		0										
mcrxr	31	crfD		00		00000				00000				512		0												

Table A-1. Complete Instruction List Sorted by Mnemonic (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
mfcrl	31		D					0	0	0	0	0	0	0	0	0	0					19						0
mffsx	63		D					0	0	0	0	0	0	0	0	0	0					583						Rc
mfmsr ²	31		D					0	0	0	0	0	0	0	0	0	0					83						0
mfspr ⁵	31		D					spr														339						0
mfsr ²	31		D				0		SR				0	0	0	0	0					595						0
mfsrin ²	31		D					0	0	0	0	0	0		B							659						0
mftb	31		D					tbr														371						0
mtrcf	31		S				0		CRM					0								144						0
mtfsb0x	63		crbD					0	0	0	0	0	0	0	0	0	0					70						Rc
mtfsb1x	63		crbD					0	0	0	0	0	0	0	0	0	0					38						Rc
mtfsfx	63	0						FM					0		B							711						Rc
mtfsfix	63		crfD	00				0	0	0	0	0	0	0	0	0	0	IMM		0		134						Rc
mtmsr ²	31		S					0	0	0	0	0	0	0	0	0	0					146						0
mtspr ⁵	31		S					spr														467						0
mtsr ²	31		S				0		SR				0	0	0	0	0					210						0
mtsrin ²	31		S					0	0	0	0	0	0		B							242						0
mulhd ¹	31		D					A					B		0							73						Rc
mulhdu ¹	31		D					A					B		0							9						Rc
mulhw ¹	31		D					A					B		0							75						Rc
mulhwu ¹	31		D					A					B		0							11						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE							233						Rc
mulld ¹	31		D					A					B		OE													



Table A-1. Complete Instruction List Sorted by Mnemonic (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
rlwimix	20		S						A					SH					MB					ME				Rc
rlwinmx	21		S						A					SH					MB					ME				Rc
rlwnmx	23		S						A					B					MB					ME				Rc
sc	17		00000					00000											0000000000000000							1	0	
slbia^{1,2,3}	31		00000					00000						00000										498				0
slbie^{1,2,3}	31		00000					00000						B										434				0
sldx¹	31		S						A					B										27				Rc
slwx	31		S						A					B										24				Rc
sradx¹	31		S						A					B										794				Rc
sradix¹	31		S						A					sh										413		sh		Rc
srawx	31		S						A					B										792				Rc
srawix	31		S						A					SH										824				Rc
srdx¹	31		S						A					B										539				Rc
srwx	31		S						A					B										536				Rc
stb	38		S						A															d				
stbu	39		S						A															d				
stbux	31		S						A					B										247				0
stbx	31		S						A					B										215				0
std¹	62		S						A															ds				0
stdcx.¹	31		S						A					B										214				1
stdu¹	62		S						A															ds				1
stdux¹	31		S						A					B										181				0
stdx¹	31		S						A					B										149				0
stfd	54		S						A															d				
stfdu	55		S						A															d				
stfdux	31		S						A					B										759				0
stfdx	31		S						A					B										727				0
stfiwx³	31		S						A					B										983				0
stfs	52		S						A															d				
stfsu	53		S						A															d				
stfsux	31		S						A					B										695				0
stfsx	31		S						A					B										663				0
sth	44		S						A															d				
sthrx	31		S						A					B										918				0
sthu	45		S						A															d				
sthux	31		S						A					B										439				0
sthx	31		S						A					B										407				0

A.2 Instructions Sorted by Opcode

Table A-2 lists the instructions defined in the PowerPC architecture in numeric order by opcode.

Table A-2. Complete Instruction List Sorted by Opcode

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
tdi ¹	000010	TO			A			SIMM																				
twi	000011	TO			A			SIMM																				
mulli	000111	D			A			SIMM																				
subfic	001000	D			A			SIMM																				
cmpli	001010	crfD	0	L	A			UIMM																				
cmpi	001011	crfD	0	L	A			SIMM																				
addic	001100	D			A			SIMM																				
addic.	001101	D			A			SIMM																				
addi	001110	D			A			SIMM																				
addis	001111	D			A			SIMM																				
bcx	010000	BO			BI			BD										AA	LK									
sc	010001	00000			00000			0000000000000000										1	0									
bx	010010	LI										AA	LK															
mcrf	010011	crfD	00	crfS	00	00000			0000000000							0												
bclrx	010011	BO			BI			00000			0000010000				LK													
crnor	010011	crbD			crbA			crbB			0000100001				0													
rfi	010011	00000			00000			00000			0000110010				0													
crandc	010011	crbD			crbA			crbB			0010000001				0													
isync	010011	00000			00000			00000			0010010110				0													
crxor	010011	crbD			crbA			crbB			0011000001				0													
crnand	010011	crbD			crbA			crbB			0011100001				0													
crand	010011	crbD			crbA			crbB			0100000001				0													
creqv	010011	crbD			crbA			crbB			0100100001				0													
crorc	010011	crbD			crbA			crbB			0110100001				0													
cror	010011	crbD			crbA			crbB			0111000001				0													
bcctrx	010011	BO			BI			00000			1000010000				LK													
rlwimix	010100	S			A			SH			MB			ME			Rc											
rlwinmx	010101	S			A			SH			MB			ME			Rc											
rlwnmx	010111	S			A			B			MB			ME			Rc											
ori	011000	S			A			UIMM																				
oris	011001	S			A			UIMM																				
xori	011010	S			A			UIMM																				
xoris	011011	S			A			UIMM																				

Table A-2. Complete Instruction List Sorted by Opcode (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
andi.	0 1 1 1 1 0 0	S								A										UIMM									
andis.	0 1 1 1 1 0 1	S								A										UIMM									
rldicl^{x1}	0 1 1 1 1 1 0	S								A		sh		mb						0 0 0		sh	Rc						
rldicr^{x1}	0 1 1 1 1 1 0	S								A		sh		me							0 0 1		sh	Rc					
rldic^{x1}	0 1 1 1 1 1 0	S								A		sh		mb							0 1 0		sh	Rc					
rldimix¹	0 1 1 1 1 1 0	S								A		sh		mb							0 1 1		sh	Rc					
rldcl^{x1}	0 1 1 1 1 1 0	S								A			B		mb						0 1 0 0 0			Rc					
rldcr^{x1}	0 1 1 1 1 1 0	S								A			B		me						0 1 0 0 1			Rc					
cmp	0 1 1 1 1 1 1	crfD	0	L						A			B		0 0 0 0 0 0 0 0 0 0							0							
tw	0 1 1 1 1 1 1	TO								A			B		0 0 0 0 0 0 0 1 0 0							0							
subfcx	0 1 1 1 1 1 1	D								A			B	OE	0 0 0 0 0 0 1 0 0 0							Rc							
mulhd^{x1}	0 1 1 1 1 1 1	D								A			B	0	0 0 0 0 0 0 1 0 0 1							Rc							
addcx	0 1 1 1 1 1 1	D								A			B	OE	0 0 0 0 0 0 1 0 1 0							Rc							
mulhw^x	0 1 1 1 1 1 1	D								A			B	0	0 0 0 0 0 0 1 0 1 1							Rc							
mfc^r	0 1 1 1 1 1 1	D				0 0 0 0 0					0 0 0 0 0			0 0 0 0 0 1 0 0 1 1							0								
lwar^x	0 1 1 1 1 1 1	D								A			B		0 0 0 0 0 1 0 1 0 0							0							
ld^{x1}	0 1 1 1 1 1 1	D								A			B		0 0 0 0 0 1 0 1 0 1							0							
lwz^x	0 1 1 1 1 1 1	D								A			B		0 0 0 0 0 1 0 1 1 1							0							
slw^x	0 1 1 1 1 1 1	S								A			B		0 0 0 0 0 1 1 0 0 0							Rc							
cntlz^{wx}	0 1 1 1 1 1 1	S								A		0 0 0 0 0		0 0 0 0 0 1 1 0 1 0							Rc								
sld^{x1}	0 1 1 1 1 1 1	S								A			B		0 0 0 0 0 1 1 0 1 1							Rc							
and^x	0 1 1 1 1 1 1	S								A			B		0 0 0 0 0 1 1 1 0 0							Rc							
cmpl	0 1 1 1 1 1 1	crfD	0	L						A			B		0 0 0 0 1 0 0 0 0 0							0							
subfx	0 1 1 1 1 1 1	D								A			B	OE	0 0 0 0 1 0 1 0 0 0							Rc							
ld^{ux1}	0 1 1 1 1 1 1	D								A			B		0 0 0 0 1 1 0 1 0 1							0							
dcbst	0 1 1 1 1 1 1	0 0 0 0 0								A			B		0 0 0 0 1 1 0 1 1 0							0							
lwz^{ux}	0 1 1 1 1 1 1	D								A			B		0 0 0 0 1 1 0 1 1 1							0							
cntlz^{dx1}	0 1 1 1 1 1 1	S								A		0 0 0 0 0		0 0 0 0 1 1 1 0 1 0							Rc								
andc^x	0 1 1 1 1 1 1	S								A			B		0 0 0 0 1 1 1 1 0 0							Rc							
td¹	0 1 1 1 1 1 1	TO								A			B		0 0 0 1 0 0 0 1 0 0							0							
mulhd^{x1}	0 1 1 1 1 1 1	D								A			B	0	0 0 0 1 0 0 1 0 0 1							Rc							
mulhw^x	0 1 1 1 1 1 1	D								A			B	0	0 0 0 1 0 0 1 0 1 1							Rc							
mfms^r	0 1 1 1 1 1 1	D				0 0 0 0 0					0 0 0 0 0			0 0 0 1 0 1 0 0 1 1							0								
ldar^{x1}	0 1 1 1 1 1 1	D								A			B		0 0 0 1 0 1 0 1 0 0							0							
dcbf	0 1 1 1 1 1 1	0 0 0 0 0								A			B		0 0 0 1 0 1 0 1 1 0							0							
lbz^x	0 1 1 1 1 1 1	D								A			B		0 0 0 1 0 1 0 1 1 1							0							
neg^x	0 1 1 1 1 1 1	D								A		0 0 0 0 0		OE	0 0 0 1 1 0 1 0 0 0							Rc							



Table A-2. Complete Instruction List Sorted by Opcode (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
lbusx	0111111	D								A					B														0
norx	0111111	S								A					B														Rc
subfex	0111111	D								A					B			OE											Rc
addex	0111111	D								A					B			OE											Rc
mtcrf	0111111	S				0				CRM							0												0
mtmsr	0111111	S							00000				00000																0
stdx ¹	0111111	S								A					B														0
stwcx.	0111111	S								A					B														1
stwx	0111111	S								A					B														0
stdux ¹	0111111	S								A					B														0
stwux	0111111	S								A					B														0
subfzex	0111111	D								A			00000					OE											Rc
addzex	0111111	D								A			00000					OE											Rc
mtsr	0111111	S				0				SR			00000																0
stdcx. ¹	0111111	S								A					B														1
stbx	0111111	S								A					B														0
subfmex	0111111	D								A			00000					OE											Rc
mulld ¹	0111111	D								A					B			OE											Rc
addmex	0111111	D								A			00000					OE											Rc
mullwx	0111111	D								A					B			OE											Rc
mtsrin	0111111	S							00000						B														0
dcbst	0111111		00000							A					B														0
stbux	0111111	S								A					B														0
addx	0111111	D								A					B			OE											Rc
dcbt	0111111		00000							A					B														0
lhzx	0111111	D								A					B														0
eqvx	0111111	S								A					B														Rc
tlbie ^{2,3}	0111111		00000					00000							B														0
eciwx	0111111	D								A					B														0
lhzux	0111111	D								A					B														0
xorx	0111111	S								A					B														Rc
mfspr ⁴	0111111	D											spr																0
lwax ¹	0111111	D								A					B														0
lhax	0111111	D								A					B														0
tlbia ^{2,3}	0111111		00000					00000					00000																0
mftb	0111111	D											tbr																0
lwaux ¹	0111111	D								A					B														0

Table A-2. Complete Instruction List Sorted by Opcode (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
lhax	011111	D								A					B						0101110111							0
sthx	011111	S								A					B						0110010111							0
orc_x	011111	S								A					B						0110011100							Rc
sradi_x¹	011111	S								A				sh							1100111011			sh				Rc
slbie^{1,2,3}	011111	00000								00000					B						0110110010							0
ecowx	011111	S								A					B						0110110110							0
sthux	011111	S								A					B						0110110111							0
orx	011111	S								A					B						0110111100							Rc
divdu_x¹	011111	D								A				B	OE						0111001001							Rc
divwu_x	011111	D								A				B	OE						0111001011							Rc
mtspr⁴	011111	S											spr								0111010011							0
dcbi	011111	00000								A				B							0111010110							0
nand_x	011111	S								A				B							0111011100							Rc
divd_x¹	011111	D								A				B	OE						0111101001							Rc
divw_x	011111	D								A				B	OE						0111101011							Rc
slbia^{1,2,3}	011111	00000								00000				00000							0111110010							0
mcrxr	011111	crfD	00							00000				00000							1000000000							0
lsw_x⁵	011111	D								A				B							1000010101							0
lwbrx	011111	D								A				B							1000010110							0
lfsx	011111	D								A				B							1000010111							0
srw_x	011111	S								A				B							1000011000							Rc
srds_x¹	011111	S								A				B							1000011011							Rc
tlbsync^{2,3}	011111	00000								00000				00000							1000110110							0
lfsux	011111	D								A				B							1000110111							0
mfsr	011111	D	0							SR				00000							1001010011							0
lswi_x⁵	011111	D								A				NB							1001010101							0
sync	011111	00000								00000				00000							1001010110							0
lfdx	011111	D								A				B							1001010111							0
lfdux	011111	D								A				B							1001110111							0
mfsrin²	011111	D								00000				B							1010010011							0
stsw_x⁵	011111	S								A				B							1010010101							0
stwbrx	011111	S								A				B							1010010110							0
stfsx	011111	S								A				B							1010010111							0
stfsux	011111	S								A				B							1010110111							0
stswi_x⁵	011111	S								A				NB							1011010101							0
stfdx	011111	S								A				B							1011010111							0
stfdux	011111	S								A				B							1011110111							0



Table A-2. Complete Instruction List Sorted by Opcode (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
lbrx	011111	D								A																			0
srawx	011111	S								A																			Rc
sradx ¹	011111	S								A																			Rc
srawix	011111	S								A					SH														Rc
eieio	011111	00000								00000					00000														0
sthbrx	011111	S								A					B														0
extshx	011111	S								A					00000														Rc
extsbx	011111	S								A					00000														Rc
tlbld ^{2,6}	011111	00000								00000					B														0
icbi	011111	00000								A					B														0
stfiwx ³	011111	S								A					B														0
extsw ¹	011111	S								A					00000														Rc
tlblid ^{2,6}	011111	00000								00000					B														0
dcbz	011111	00000								A					B														0
lwz	100000	D								A																			d
lwzu	100001	D								A																			d
lbz	100010	D								A																			d
lbzu	100011	D								A																			d
stw	100100	S								A																			d
stwu	100101	S								A																			d
stb	100110	S								A																			d
stbu	100111	S								A																			d
lhz	101000	D								A																			d
lhzu	101001	D								A																			d
lha	101010	D								A																			d
lhau	101011	D								A																			d
sth	101100	S								A																			d
sthu	101101	S								A																			d
lmw ⁵	101110	D								A																			d
stmw ⁵	101111	S								A																			d
lfs	110000	D								A																			d
lfsu	110001	D								A																			d
lfd	110010	D								A																			d
lfdv	110011	D								A																			d
stfs	110100	S								A																			d
stfsu	110101	S								A																			d
stfd	110110	S								A																			d



Table A-2. Complete Instruction List Sorted by Opcode (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31						
stfdu	1 1 0 1 1 1	S		A		d																												
ld¹	1 1 1 0 1 0	D		A		ds										0 0																		
ldu¹	1 1 1 0 1 0	D		A		ds										0 1																		
lwa¹	1 1 1 0 1 0	D		A		ds										1 0																		
fdivsx	1 1 1 0 1 1	D		A		B		0 0 0 0 0			1 0 0 1 0			Rc																				
fsubsx	1 1 1 0 1 1	D		A		B		0 0 0 0 0			1 0 1 0 0			Rc																				
faddsx	1 1 1 0 1 1	D		A		B		0 0 0 0 0			1 0 1 0 1			Rc																				
fsqrts³	1 1 1 0 1 1	D		0 0 0 0 0			B		0 0 0 0 0			1 0 1 1 0			Rc																			
fresx³	1 1 1 0 1 1	D		0 0 0 0 0			B		0 0 0 0 0			1 1 0 0 0			Rc																			
fmulsx	1 1 1 0 1 1	D		A		0 0 0 0 0			C			1 1 0 0 1			Rc																			
fmsubsx	1 1 1 0 1 1	D		A		B		C			1 1 1 0 0			Rc																				
fmaddsx	1 1 1 0 1 1	D		A		B		C			1 1 1 0 1			Rc																				
fnmsubsx	1 1 1 0 1 1	D		A		B		C			1 1 1 1 0			Rc																				
fnmaddsx	1 1 1 0 1 1	D		A		B		C			1 1 1 1 1			Rc																				
std¹	1 1 1 1 1 0	S		A		ds										0 0																		
stdu¹	1 1 1 1 1 0	S		A		ds										0 1																		
fcmpu	1 1 1 1 1 1	crfD		0 0		A		B		0 0 0 0 0 0 0 0 0 0							0																	
frsp^x	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 0 0 0 1 1 0 0							Rc																		
fctiw^x	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 0 0 0 1 1 1 0																									
fctiwz^x	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 0 0 0 1 1 1 1							Rc																		
fdiv^x	1 1 1 1 1 1	D		A		B		0 0 0 0 0			1 0 0 1 0			Rc																				
fsub^x	1 1 1 1 1 1	D		A		B		0 0 0 0 0			1 0 1 0 0			Rc																				
fadd^x	1 1 1 1 1 1	D		A		B		0 0 0 0 0			1 0 1 0 1			Rc																				
fsqrt^{x3}	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 0 0			1 0 1 1 0			Rc																			
fsel^{x3}	1 1 1 1 1 1	D		A		B		C			1 0 1 1 1			Rc																				
fmul^x	1 1 1 1 1 1	D		A		0 0 0 0 0			C			1 1 0 0 1			Rc																			
frsqrte^{x3}	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 0 0			1 1 0 1 0			Rc																			
fmsub^x	1 1 1 1 1 1	D		A		B		C			1 1 1 0 0			Rc																				
fmadd^x	1 1 1 1 1 1	D		A		B		C			1 1 1 0 1			Rc																				
fnmsub^x	1 1 1 1 1 1	D		A		B		C			1 1 1 1 0			Rc																				
fnmadd^x	1 1 1 1 1 1	D		A		B		C			1 1 1 1 1			Rc																				
fcmpo	1 1 1 1 1 1	crfD		0 0		A		B		0 0 0 0 1 0 0 0 0 0							0																	
mtfsb1^x	1 1 1 1 1 1	crbD			0 0 0 0 0			0 0 0 0 0			0 0 0 0 1 0 0 1 1 0							Rc																
fneg^x	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 0 1 0 1 0 0 0							Rc																		
mcrfs	1 1 1 1 1 1	crfD		0 0		crfS		0 0		0 0 0 0 0			0 0 0 1 0 0 0 0 0 0							0														
mtfsb0^x	1 1 1 1 1 1	crbD			0 0 0 0 0			0 0 0 0 0			0 0 0 1 0 0 0 1 1 0							Rc																
fmr^x	1 1 1 1 1 1	D		0 0 0 0 0			B		0 0 0 1 0 0 1 0 0 0							Rc																		



Table A-2. Complete Instruction List Sorted by Opcode (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
mtfsfix	1	1	1	1	1	1	1	1	crfD	0	0	0	0	0	0	0	0	IMM	0	0	0	1	0	0	0	0	1	1	0	Rc
fnabsxv	1	1	1	1	1	1	1	1	D		0	0	0	0	0	0	B		0	0	1	0	0	0	1	0	0	0	0	Rc
fabsx	1	1	1	1	1	1	1	1	D		0	0	0	0	0	0	B		0	1	0	0	0	0	1	0	0	0	0	Rc
mffsx	1	1	1	1	1	1	1	1	D		0	0	0	0	0	0	00000		1	0	0	1	0	0	0	1	1	1	1	Rc
mtfsfx	1	1	1	1	1	1	1	1	0	FM				0	B		1	0	1	1	0	0	0	1	1	1	1	1	1	Rc
fctidx¹	1	1	1	1	1	1	1	1	D		0	0	0	0	0	0	B		1	1	0	0	1	0	1	1	1	1	0	Rc
fctidzx¹	1	1	1	1	1	1	1	1	D		0	0	0	0	0	0	B		1	1	0	0	1	0	1	1	1	1	1	Rc
fcfidx¹	1	1	1	1	1	1	1	1	D		0	0	0	0	0	0	B		1	1	0	1	0	0	1	1	1	1	0	Rc

- ¹ 64-bit instruction
- ² Supervisor-level instruction
- ³ Optional in the PowerPC architecture
- ⁴ Supervisor- and user-level instruction
- ⁵ Load and store string or multiple instruction
- ⁶ G2-implementation-specific instruction

A.3 Instructions Grouped by Functional Categories

Table A-3 through Table A-30 list the PowerPC instructions grouped by function.

Table A-3. Integer Arithmetic Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
addx	31				D					A					B			OE										266	Rc
addcx	31				D					A					B			OE										10	Rc
addex	31				D					A					B			OE										138	Rc
addi	14				D					A									SIMM										
addic	12				D					A									SIMM										
addic.	13				D					A									SIMM										
addis	15				D					A									SIMM										
addmex	31				D					A			0	0	0	0	0	0	OE									234	Rc
addzex	31				D					A			0	0	0	0	0	0	OE									202	Rc
divdx¹	31				D					A					B			OE										489	Rc
divdux¹	31				D					A					B			OE										457	Rc
divwx	31				D					A					B			OE										491	Rc
divwux	31				D					A					B			OE										459	Rc
mulhd¹	31				D					A					B			0										73	Rc
mulhdux¹	31				D					A					B			0										9	Rc
mulhwx	31				D					A					B			0										75	Rc
mulhwux	31				D					A					B			0										11	Rc
mulld¹	31				D					A					B			OE										233	Rc
mulli	07				D					A									SIMM										
mullwx	31				D					A					B			OE										235	Rc
negx	31				D					A			0	0	0	0	0	0	OE									104	Rc
subfx	31				D					A					B			OE										40	Rc
subfcx	31				D					A					B			OE										8	Rc
subficx	08				D					A									SIMM										
subfex	31				D					A					B			OE										136	Rc
subfmex	31				D					A			0	0	0	0	0	0	OE									232	Rc
subfzex	31				D					A			0	0	0	0	0	0	OE									200	Rc

¹ 64-bit instruction



Table A-4. Integer Compare Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
cmp	31	crfD	0	L		A		B	0 0 0 0 0 0 0 0 0 0										0									
cmpi	11	crfD	0	L		A	SIMM																					
cmpl	31	crfD	0	L		A		B	32										0									
cmpli	10	crfD	0	L		A	UIMM																					

Table A-5. Integer Logical Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
andx	31	S				A		B	28										Rc									
andcx	31	S				A		B	60										Rc									
andi	28	S				A	UIMM																					
andis	29	S				A	UIMM																					
cntlzdx ¹	31	S				A	0 0 0 0 0					58					Rc											
cntlzwx	31	S				A	0 0 0 0 0					26					Rc											
eqvx	31	S				A		B	284										Rc									
extsbx	31	S				A	0 0 0 0 0					954					Rc											
extshx	31	S				A	0 0 0 0 0					922					Rc											
extswx ¹	31	S				A	0 0 0 0 0					986					Rc											
nandx	31	S				A		B	476										Rc									
norx	31	S				A		B	124										Rc									
orx	31	S				A		B	444										Rc									
orcx	31	S				A		B	412										Rc									
ori	24	S				A	UIMM																					
oris	25	S				A	UIMM																					
xorx	31	S				A		B	316										Rc									
xori	26	S				A	UIMM																					
xoris	27	S				A	UIMM																					

¹ 64-bit instruction

Table A-6. Integer Rotate Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
rldclx ¹	30		S				A						B							mb				8				Rc
rldcrx ¹	30		S				A						B							me				9				Rc
rldicx ¹	30		S				A						sh							mb			2		sh			Rc
rldicl ¹	30		S				A						sh							mb			0		sh			Rc
rldicr ¹	30		S				A						sh							me			1		sh			Rc
rldimix ¹	30		S				A						sh							mb			3		sh			Rc
rlwimix	22		S				A						SH							MB								Rc
rlwinmx	20		S				A						SH							MB								Rc
rlwnmx	21		S				A						SH							MB								Rc

¹ 64-bit instruction

Table A-7. Integer Shift Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
sldx ¹	31		S				A						B											27				Rc
slwx	31		S				A						B											24				Rc
sradx ¹	31		S				A						B											794				Rc
sradx ¹	31		S				A						sh											413		sh		Rc
srawx	31		S				A						B											792				Rc
srawix	31		S				A						SH											824				Rc
srdx ¹	31		S				A						B											539				Rc
srwx	31		S				A						B											536				Rc

¹ 64-bit instruction



Table A-8. Floating-Point Arithmetic Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
faddx	63		D				A			B			0 0 0 0 0										21					Rc
faddsx	59		D				A			B			0 0 0 0 0										21					Rc
fdivx	63		D				A			B			0 0 0 0 0										18					Rc
fdivsx	59		D				A			B			0 0 0 0 0										18					Rc
fmulx	63		D				A			0 0 0 0 0									C				25					Rc
fmulsx	59		D				A			0 0 0 0 0									C				25					Rc
fresx¹	59		D				0 0 0 0 0			B			0 0 0 0 0										24					Rc
frsqrtox¹	63		D				0 0 0 0 0			B			0 0 0 0 0										26					Rc
fsubx	63		D				A			B			0 0 0 0 0										20					Rc
fsubsx	59		D				A			B			0 0 0 0 0										20					Rc
fselx¹	63		D				A			B									C				23					Rc
fsqrtx¹	63		D				0 0 0 0 0			B			0 0 0 0 0										22					Rc
fsqrtsx¹	59		D				0 0 0 0 0			B			0 0 0 0 0										22					Rc

¹ Optional in the PowerPC architecture

Table A-9. Floating-Point Multiply-Add Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
fmaddx	63		D				A			B										C				29					Rc
fmaddsx	59		D				A			B										C				29					Rc
fmsubx	63		D				A			B										C				28					Rc
fmsubsx	59		D				A			B										C				28					Rc
fnmaddx	63		D				A			B										C				31					Rc
fnmaddsx	59		D				A			B										C				31					Rc
fnmsubx	63		D				A			B										C				30					Rc
fnmsubsx	59		D				A			B										C				30					Rc

Table A-10. Floating-Point Rounding and Conversion Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
fcfidx¹	63		D				0 0 0 0 0			B													846						Rc
fctidx¹	63		D				0 0 0 0 0			B													814						Rc
fctidzx¹	63		D				0 0 0 0 0			B													815						Rc
fctiw¹	63		D				0 0 0 0 0			B													14						Rc
fctiwzx	63		D				0 0 0 0 0			B													15						Rc
frsp¹	63		D				0 0 0 0 0			B													12						Rc

¹ 64-bit instruction

Table A-11. Floating-Point Compare Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
fcmpo	63	crfD			00		A			B			32			0												
fcmpu	63	crfD			00		A			B			0			0												

Table A-12. Floating-Point Status and Control Register Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
mcrfs	63	crfD			00		crfS			00			00000			64			0									
mffsx	63	D			00000			00000			583			Rc														
mtfsb0x	63	crbD			00000			00000			70			Rc														
mtfsb1x	63	crbD			00000			00000			38			Rc														
mtfsfx	31	0	FM			0	B			711			Rc															
mtfsfix	63	crfD			00		00000			IMM		0	134			Rc												

Table A-13. Integer Load Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
lbz	34	D			A			d																				
lbzu	35	D			A			d																				
lbzux	31	D			A			B			119			0														
lbzx	31	D			A			B			87			0														
ld¹	58	D			A			ds			0																	
ldu¹	58	D			A			ds			1																	
ldux¹	31	D			A			B			53			0														
ldx¹	31	D			A			B			21			0														
lha	42	D			A			d																				
lhau	43	D			A			d																				
lhaux	31	D			A			B			375			0														
lhax	31	D			A			B			343			0														
lhz	40	D			A			d																				
lhzu	41	D			A			d																				
lhzux	31	D			A			B			311			0														
lhzx	31	D			A			B			279			0														
lwa¹	58	D			A			ds			2																	
lwaux¹	31	D			A			B			373			0														
lwax¹	31	D			A			B			341			0														
lwz	32	D			A			d																				
lwzu	33	D			A			d																				

Table A-13. Integer Load Instructions (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
lwzux	31			D						A					B														0
lwzx	31			D						A					B														0

¹ 64-bit instruction

Table A-14. Integer Store Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
stb	38			S						A																			
stbu	39			S						A																			
stbux	31			S						A					B														0
stbx	31			S						A					B														0
std¹	62			S						A																			0
stdu¹	62			S						A																			1
stdux¹	31			S						A					B														0
stdx¹	31			S						A					B														0
sth	44			S						A																			
sthu	45			S						A																			
sthux	31			S						A					B														0
sthx	31			S						A					B														0
stw	36			S						A																			
stwu	37			S						A																			
stwux	31			S						A					B														0
stwx	31			S						A					B														0

¹ 64-bit instruction

Table A-15. Integer Load and Store with Byte-Reverse Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
lhbrx	31			D						A					B														0
lwbrx	31			D						A					B														0
sthbrx	31			S						A					B														0
stwbrx	31			S						A					B														0

Table A-16. Integer Load and Store Multiple Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
lmw¹	46			D						A																			
stmw¹	47			S						A																			

¹ Load and store string or multiple instruction

Table A-17. Integer Load and Store String Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
lswi ¹	31				D					A					NB													0
lswx ¹	31				D					A					B													0
stswi ¹	31				S					A					NB													0
stswx ¹	31				S					A					B													0

¹ Load and store string or multiple instruction

Table A-18. Memory Synchronization Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
eieio	31				0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
isync	19				0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
ldarx ¹	31				D					A					B													0	
lwarx	31				D					A					B													0	
stdcx ¹	31				S					A					B													1	
stwcx.	31				S					A					B													1	
sync	31				0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

¹ 64-bit instruction

Table A-19. Floating-Point Load Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
lfd	50				D					A																			
lfdx	51				D					A																			
lfdx	31				D					A					B													0	
lfdx	31				D					A					B													0	
lfs	48				D					A																			
lfsu	49				D					A																			
lfsux	31				D					A					B													0	
lfsx	31				D					A					B													0	

Table A-20. Floating-Point Store Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
stfd	54				S					A																		
stfdx	55				S					A																		
stfdx	31				S					A					B													0
stfdx	31				S					A					B													0
stfiwx ¹	31				S					A					B													0

Table A-20. Floating-Point Store Instructions (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
stfs	52	S			A			d																				
stfsu	53	S			A			d																				
stfsux	31	S			A			B			695												0					
stfsx	31	S			A			B			663												0					

¹ Optional in the PowerPC architecture

Table A-21. Floating-Point Move Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
fabsx	63	D			00000			B			264												Rc					
fmr_x	63	D			00000			B			72												Rc					
fnabs_x	63	D			00000			B			136												Rc					
fneg_x	63	D			00000			B			40												Rc					

Table A-22. Branch Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
bx	18	LI																								AA	LK	
bc_x	16	BO			BI			BD												AA	LK							
bcctr_x	19	BO			BI			00000			528												LK					
bclr_x	19	BO			BI			00000			16												LK					

Table A-23. Condition Register Logical Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
crand	19	crbD			crbA			crbB			257												0					
crandc	19	crbD			crbA			crbB			129												0					
creqv	19	crbD			crbA			crbB			289												0					
crnand	19	crbD			crbA			crbB			225												0					
crnor	19	crbD			crbA			crbB			33												0					
cror	19	crbD			crbA			crbB			449												0					
crorc	19	crbD			crbA			crbB			417												0					
crxor	19	crbD			crbA			crbB			193												0					
mcrf	19	crfD	00	crfS	00	00000			0000000000												0							

Table A-24. System Linkage Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
rfi ¹	19	00000			00000			00000			50						0											
sc	17	00000			00000			0000000000000000																	1	0		

¹ Supervisor-level instruction

Table A-25. Trap Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
td ¹	31	TO			A			B			68						0											
tdi ¹	03	TO			A			SIMM																				
tw	31	TO			A			B			4						0											
twi	03	TO			A			SIMM																				

¹ 64-bit instruction

Table A-26. Processor Control Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
mcrxr	31	crfS		00		00000			00000			512						0										
mfcrr	31	D			00000			00000			19						0											
mfmsr ¹	31	D			00000			00000			83						0											
mfspr ²	31	D			spr						339						0											
mftb	31	D			tpr						371						0											
mtcrf	31	S		0	CRM			0	144						0													
mtmsr ¹	31	S			00000			00000			146						0											
mtspr ²	31	D			spr						467						0											

¹ Supervisor-level instruction

² Supervisor- and user-level instruction

Table A-27. Cache Management Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
dcbf	31	00000			A			B			86						0											
dcbi ¹	31	00000			A			B			470						0											
dcbst	31	00000			A			B			54						0											
dcbt	31	00000			A			B			278						0											
dcbtst	31	00000			A			B			246						0											
dcbz	31	00000			A			B			1014						0											
icbi	31	00000			A			B			982						0											

¹ Supervisor-level instruction



Table A-28. Segment Register Manipulation Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
mfsr ¹	31			D			0		SR				0	0	0	0												595	0
mfsrin ¹	31			D					0	0	0	0																659	0
mtsr ¹	31			S			0		SR				0	0	0	0												210	0
mtsrin ¹	31			S					0	0	0	0																242	0

¹ Supervisor-level instruction

Table A-29. Lookaside Buffer Management Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
slbia ^{1, 2, 3}	31				0	0	0	0					0	0	0	0													498	0
slbie ^{1, 2, 3}	31				0	0	0	0																					434	0
tlbia ^{1, 3}	31				0	0	0	0					0	0	0	0													370	0
tlbie ^{1, 3}	31				0	0	0	0																					306	0
tlbld ^{1, 4}	31				0	0	0	0																					978	0
tlbli ^{1, 4}	31				0	0	0	0																					1010	0
tlbsync ^{1, 3}	31				0	0	0	0					0	0	0	0													566	0

¹ Supervisor-level instruction

² 64-bit instruction

³ Optional in the PowerPC architecture

⁴ G2-implementation specific instruction

Table A-30. External Control Instructions

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
eciwx	31			D					A																				310	0
ecowx	31			S					A																				438	0

A.4 Instructions Sorted by Form

Table A-31 through Table A-45 list the PowerPC instructions grouped by form.

Table A-31. I-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
	OPCD					LI															AA	LK							
Specific Instruction																													
bcx	18	LI															AA	LK											

Table A-32. B-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
	OPCD					BO			BI			BD							AA	LK									
Specific Instruction																													
bcx	16	BO			BI			BD							AA	LK													

Table A-33. SC-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
	OPCD					00000			00000			000000000000000000												1	0				
Specific Instruction																													
sc	17	00000			00000			000000000000000000												1	0								

Table A-34. D-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
	OPCD					D			A			d																	
	OPCD					D			A			SIMM																	
	OPCD					S			A			d																	
	OPCD					S			A			UIMM																	
	OPCD					crfD		0	L	A			SIMM																
	OPCD					crfD		0	L	A			UIMM																
	OPCD					TO			A			SIMM																	
Specific Instruction																													
addi	14	D			A			SIMM																					
addic	12	D			A			SIMM																					

Table A-34. D-Form (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
tdi ²	02	TO			A			SIMM																				
twi	03	TO			A			SIMM																				
xori	26	S			A			UIMM																				
xoris	27	S			A			UIMM																				

¹ Load and store string or multiple instruction

² 64-bit instruction

Table A-35. DS-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD	D			A			ds										XO										
	OPCD	S			A			ds										XO										

Specific Instructions

ld ¹	58	D			A			ds										0
ldu ¹	58	D			A			ds										1
lwa ¹	58	D			A			ds										2
std ¹	62	S			A			ds										0
stdu ¹	62	S			A			ds										1

¹ 64-bit instruction

Table A-36. X-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD	D			A			B			XO										0							
	OPCD	D			A			NB			XO										0							
	OPCD	D			00000			B			XO										0							
	OPCD	D			00000			00000			XO										0							
	OPCD	D			0	SR		00000			XO										0							
	OPCD	S			A			B			XO										Rc							
	OPCD	S			A			B			XO										1							
	OPCD	S			A			B			XO										0							
	OPCD	S			A			NB			XO										0							
	OPCD	S			A			00000			XO										Rc							
	OPCD	S			00000			B			XO										0							
	OPCD	S			00000			00000			XO										0							
	OPCD	S			0	SR		00000			XO										0							
	OPCD	S			A			SH			XO										Rc							

Table A-36. X-Form (continued)

Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

OPCD	crfD	0	L	A	B	XO	0
OPCD	crfD	00		A	B	XO	0
OPCD	crfD	00	crfS	00	00000	XO	0
OPCD	crfD	00	00000	00000	00000	XO	0
OPCD	crfD	00	00000	IMM	0	XO	Rc
OPCD	TO			A	B	XO	0
OPCD	D			00000	B	XO	Rc
OPCD	D			00000	00000	XO	Rc
OPCD	crbD			00000	00000	XO	Rc
OPCD	00000			A	B	XO	0
OPCD	00000			00000	B	XO	0
OPCD	00000			00000	00000	XO	0

Specific Instructions

andx	31	S		A	B	28	Rc
andcx	31	S		A	B	60	Rc
cmp	31	crfD	0	L	A	B	0
cmpl	31	crfD	0	L	A	B	32
cntlzdx¹	31	S		A	00000	58	Rc
cntlzwx	31	S		A	00000	26	Rc
dcbf	31	00000		A	B	86	0
dcbi²	31	00000		A	B	470	0
dcbst	31	00000		A	B	54	0
dcbt	31	00000		A	B	278	0
dcbtst	31	00000		A	B	246	0
dcbz	31	00000		A	B	1014	0
eciwx	31	D		A	B	310	0
ecowx	31	S		A	B	438	0
eieio	31	00000		00000	00000	854	0
eqvx	31	S		A	B	284	Rc
extsbx	31	S		A	00000	954	Rc
extshx	31	S		A	00000	922	Rc
extswx¹	31	S		A	00000	986	Rc
fabsx	63	D		00000	B	264	Rc
fcfidx¹	63	D		00000	B	846	Rc
fcmpo	63	crfD	00	A	B	32	0

Table A-36. X-Form (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
fcmpu	63	crfD		00					A					B								0						0
fctid^{x1}	63		D						00000					B								814						Rc
fctidz^{x1}	63		D						00000					B								815						Rc
fctiw^x	63		D						00000					B								14						Rc
fctiwz^x	63		D						00000					B								15						Rc
fmr^x	63		D						00000					B								72						Rc
fnabs^x	63		D						00000					B								136						Rc
fneg^x	63		D						00000					B								40						Rc
frsp^x	63		D						00000					B								12						Rc
icbi	31			00000					A					B								982						0
lbzux	31		D						A					B								119						0
lbzx	31		D						A					B								87						0
ldar^{x1}	31		D						A					B								84						0
ldux^{x1}	31		D						A					B								53						0
ldx^{x1}	31		D						A					B								21						0
ldux	31		D						A					B								631						0
ldx	31		D						A					B								599						0
ldsux	31		D						A					B								567						0
ldsx	31		D						A					B								535						0
lhaux	31		D						A					B								375						0
lhax	31		D						A					B								343						0
lhbrx	31		D						A					B								790						0
lhzux	31		D						A					B								311						0
lhzx	31		D						A					B								279						0
lswi³	31		D						A					NB								597						0
lswx³	31		D						A					B								533						0
lwarx	31		D						A					B								20						0
lwaux^{x1}	31		D						A					B								373						0
lwax^{x1}	31		D						A					B								341						0
lwbrx	31		D						A					B								534						0
lwzux	31		D						A					B								55						0
lwzx	31		D						A					B								23						0
mcrfs	63	crfD		00				crfS		00				00000								64						0
mcrxr	31	crfD		00				00000		00000				00000								512						0
mfc^r	31		D					00000		00000				00000								19						0
mffs^x	63		D					00000		00000				00000								583						Rc



Table A-36. X-Form (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
mfmsr ²	31				D				0	0	0	0	0	0	0	0	0	0											0
mfsr ²	31				D			0		SR																			0
mfsrin ²	31				D				0	0	0	0	0	0	0	0	0	0											0
mtfsb0x	63				crbD				0	0	0	0	0	0	0	0	0	0											Rc
mtfsb1x	63				crfD				0	0	0	0	0	0	0	0	0	0											Rc
mtfsfix	63				crbD		0	0							IMM		0												Rc
mtmsr ²	31				S				0	0	0	0	0	0	0	0	0	0											0
mtsr ²	31				S			0		SR																			0
mtsrin ²	31				S				0	0	0	0	0	0	0	0	0	0											0
nandx	31				S					A					B														Rc
norx	31				S					A					B														Rc
orx	31				S					A					B														Rc
orcx	31				S					A					B														Rc
slbia ^{1,2,4}	31					0	0	0	0	0	0	0	0	0	0	0	0	0											0
slbie ^{1,2,4}	31					0	0	0	0	0	0	0	0	0	B														0
sldx ¹	31				S					A					B														Rc
slwx	31				S					A					B														Rc
sradx ¹	31				S					A					B														Rc
srawx	31				S					A					B														Rc
srawix	31				S					A					SH														Rc
srdx ¹	31				S					A					B														Rc
srwx	31				S					A					B														Rc
stbux	31				S					A					B														0
stbx	31				S					A					B														0
stdcx ¹	31				S					A					B														1
stdux ¹	31				S					A					B														0
stdx ¹	31				S					A					B														0
stfdux	31				S					A					B														0
stfdx	31				S					A					B														0
stfiwx ⁴	31				S					A					B														0
stfsux	31				S					A					B														0
stfsx	31				S					A					B														0
sthbrx	31				S					A					B														0
sthux	31				S					A					B														0
sthx	31				S					A					B														0
stswi ³	31				S					A					NB														0

Table A-36. X-Form (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
stswx ³	31				S					A					B														0
stwbrx	31				S					A					B														0
stwcx.	31				S					A					B														1
stwux	31				S					A					B														0
stwx	31				S					A					B														0
sync	31				0	0	0	0	0	0			0	0	0	0	0	0											0
td ¹	31				TO					A					B														0
tlbia ^{2,4}	31				0	0	0	0	0	0			0	0	0	0	0	0											0
tlbie ^{2,4}	31				0	0	0	0	0	0					B														0
tlbld ^{2,5}	31				0	0	0	0	0	0					B														0
tlbli ^{2,5}	31				0	0	0	0	0	0					B														0
tlbsync ^{2,4}	31				0	0	0	0	0	0			0	0	0	0	0	0											0
tw	31				TO					A					B														0
xorx	31				S					A					B														Rc

- ¹ 64-bit instruction
- ² Supervisor- and user-level instruction
- ³ Load and store string or multiple instruction
- ⁴ Optional in the PowerPC architecture
- ⁵ G2-implementation specific instruction

Table A-37. XL-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
	OPCD				BO					BI			0	0	0	0													LK
	OPCD				crbD					crbA			crbB																0
	OPCD				crfD	0	0			crfS	0	0	0	0	0	0	0												0
	OPCD				0	0	0	0		0	0	0	0	0	0	0	0												0

Specific Instructions

bcctx	19				BO					BI			0	0	0	0													LK
bclrx	19				BO					BI			0	0	0	0													LK
crand	19				crbD					crbA			crbB																0
crandc	19				crbD					crbA			crbB																0
creqv	19				crbD					crbA			crbB																0
crnand	19				crbD					crbA			crbB																0
crnor	19				crbD					crbA			crbB																0
cror	19				crbD					crbA			crbB																0
crorc	19				crbD					crbA			crbB																0

Table A-37. XL-Form (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
crxor	19	crbD			crbA			crbB			193			0														
isync	19	0 0 0 0			0 0 0 0			0 0 0 0			150			0														
mcrf	19	crfD		0 0		crfS		0 0		0 0 0 0			0			0												
rfi ¹	19	0 0 0 0			0 0 0 0			0 0 0 0			50			0														

¹ Supervisor-level instruction

Table A-38. XFX-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD	D			spr			XO			0																	
	OPCD	D			0	CRM			0	XO			0															
	OPCD	S			spr			XO			0																	
	OPCD	D			tbr			XO			0																	

Specific Instructions

mf spr ¹	31	D			spr			339			0				
mftb	31	D			tbr			371			0				
mtcrf	31	S			0	CRM			0	144			0		
mt spr ¹	31	D			spr			467			0				

¹ Supervisor- and user-level instruction

Table A-39. XFL-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD	0	FM			0	B			XO			Rc															

Specific Instructions

mtfsfx	63	0	FM			0	B			711			Rc		
---------------	----	---	----	--	--	---	---	--	--	-----	--	--	----	--	--

Table A-40. XS-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD	S			A			sh			XO			sh Rc														

Specific Instructions

sradix ¹	31	S			A			sh			413			sh Rc		
----------------------------	----	---	--	--	---	--	--	----	--	--	-----	--	--	-------	--	--

¹ 64-bit instruction

Table A-41. XO-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD	D					A						B	OE					XO									Rc
	OPCD	D					A						B	0					XO									Rc
	OPCD	D					A		0	0	0	0	0	0	OE				XO									Rc

Specific Instructions

addx	31	D					A						B	OE					266									Rc
addcx	31	D					A						B	OE					10									Rc
addex	31	D					A						B	OE					138									Rc
addmex	31	D					A		0	0	0	0	0	OE					234									Rc
addzex	31	D					A		0	0	0	0	0	OE					202									Rc
divdx¹	31	D					A						B	OE					489									Rc
divdux¹	31	D					A						B	OE					457									Rc
divwx	31	D					A						B	OE					491									Rc
divwux	31	D					A						B	OE					459									Rc
mulhd¹	31	D					A						B	0					73									Rc
mulhdux¹	31	D					A						B	0					9									Rc
mulhwx	31	D					A						B	0					75									Rc
mulhwux	31	D					A						B	0					11									Rc
mulldx¹	31	D					A						B	OE					233									Rc
mullwx	31	D					A						B	OE					235									Rc
negx	31	D					A		0	0	0	0	0	OE					104									Rc
subfx	31	D					A						B	OE					40									Rc
subfcx	31	D					A						B	OE					8									Rc
subfex	31	D					A						B	OE					136									Rc
subfmex	31	D					A		0	0	0	0	0	OE					232									Rc
subfzex	31	D					A		0	0	0	0	0	OE					200									Rc

¹ 64-bit instruction

Table A-42. A-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
	OPCD	D					A						B	0	0	0	0	0	0	XO									Rc
	OPCD	D					A						B					C		XO									Rc
	OPCD	D					A		0	0	0	0	0					C		XO									Rc
	OPCD	D					0	0	0	0	0	0	B	0	0	0	0	0		XO									Rc



Table A-42. A-Form (continued)

Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Specific Instructions

faddx	63	D	A	B	0 0 0 0 0	21	Rc
faddsx	59	D	A	B	0 0 0 0 0	21	Rc
fdivx	63	D	A	B	0 0 0 0 0	18	Rc
fdivsx	59	D	A	B	0 0 0 0 0	18	Rc
fmaddx	63	D	A	B	C	29	Rc
fmaddsx	59	D	A	B	C	29	Rc
fmsubx	63	D	A	B	C	28	Rc
fmsubsx	59	D	A	B	C	28	Rc
fmulx	63	D	A	0 0 0 0 0	C	25	Rc
fmulsx	59	D	A	0 0 0 0 0	C	25	Rc
fnmaddx	63	D	A	B	C	31	Rc
fnmaddsx	59	D	A	B	C	31	Rc
fnmsubx	63	D	A	B	C	30	Rc
fnmsubsx	59	D	A	B	C	30	Rc
fresx¹	59	D	0 0 0 0 0	B	0 0 0 0 0	24	Rc
frsqrte¹	63	D	0 0 0 0 0	B	0 0 0 0 0	26	Rc
fselx¹	63	D	A	B	C	23	Rc
fsqrtx¹	63	D	0 0 0 0 0	B	0 0 0 0 0	22	Rc
fsqrtsx¹	59	D	0 0 0 0 0	B	0 0 0 0 0	22	Rc
fsubx	63	D	A	B	0 0 0 0 0	20	Rc
fsubsx	59	D	A	B	0 0 0 0 0	20	Rc

¹ Optional in the PowerPC architecture

Table A-43. M-Form

Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

OPCD	S	A	SH	MB	ME	Rc
OPCD	S	A	B	MB	ME	Rc

Specific Instructions

rlwimix	20	S	A	SH	MB	ME	Rc
rlwinmx	21	S	A	SH	MB	ME	Rc
rlwnmx	23	S	A	B	MB	ME	Rc

Table A-44. MD-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD				S					A				sh						mb			XO	sh	Rc			
	OPCD				S					A				sh						me			XO	sh	Rc			

Specific Instructions

ridicx ¹	30				S					A				sh					mb			2	sh	Rc				
rldiclx ¹	30				S					A				sh					mb			0	sh	Rc				
rldicrx ¹	30				S					A				sh					me			1	sh	Rc				
rldimix ¹	30				S					A				sh					mb			3	sh	Rc				

¹ 64-bit instruction

Table A-45. MDS-Form

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	OPCD				S					A				B						mb			XO		Rc			
	OPCD				S					A				B						me			XO		Rc			

Specific Instructions

rldclx ¹	30				S					A				B					mb			8		Rc				
rldcrx ¹	30				S					A				B					me			9		Rc				

¹ 64-bit instruction

A.5 Instruction Set Legend

Table A-46 provides general information on the PowerPC instruction set (such as the architectural level, privilege level, and form).

Table A-46. PowerPC Instruction Set Legend

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
addx	√						XO
addcx	√						XO
addex	√						XO
addi	√						D
addic	√						D
addic.	√						D
addis	√						D
addmex	√						XO
addzex	√						XO
andx	√						X
andcx	√						X
andi.	√						D
andis.	√						D
bx	√						I
bcx	√						B
bcctrx	√						XL
bclrx	√						XL
cmp	√						X
cmpi	√						D
cmpl	√						X
cmpli	√						D
cntlzdx ¹	√				√		X
cntlzwx	√						X
crand	√						XL
crandc	√						XL
creqv	√						XL
crnand	√						XL
crnor	√						XL
cror	√						XL
crorc	√						XL
crxor	√						XL

Table A-46. PowerPC Instruction Set Legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
dcbf		√					X
dcbi ²			√	√			X
dcbst		√					X
dcbt		√					X
dcbtst		√					X
dcbz		√					X
divd _x ¹	√				√		XO
divdu _x ¹	√				√		XO
divw _x	√						XO
divw _u _x	√						XO
eciw _x		√				√	X
ecow _x		√				√	X
eieio		√					X
eqv _x	√						X
extsb _x	√						X
extsh _x	√						X
extsw _x ¹	√				√		X
fabs _x	√						X
fadd _x	√						A
fadd _s _x	√						A
fcfid _x ¹	√				√		X
fcmpo	√						X
fcmpu	√						X
fctid _x ¹	√				√		X
fctid _z _x ¹	√				√		X
fctiw _x	√						X
fctiw _z _x	√						X
fdiv _x	√						A
fdiv _s _x	√						A
fmadd _x	√						A
fmadd _s _x	√						A
fmr _x	√						X
fmsub _x	√						A
fmsub _s _x	√						A
fmul _x	√						A

Table A-46. PowerPC Instruction Set Legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
fmulsx	√						A
fnabsx	√						X
fnegx	√						X
fnmaddx	√						A
fnmaddsx	√						A
fnmsubx	√						A
fnmsubsx	√						A
fresx ³	√					√	A
frsp ^x	√						X
frsqrte ^{x3}	√					√	A
fsel ^{x3}	√					√	A
fsqrt ^{x3}	√					√	A
fsqrts ^{x3}	√					√	A
fsub ^x	√						A
fsubsx	√						A
icbi		√					X
isync		√					XL
lbz	√						D
lbzu	√						D
lbzux	√						X
lbzx	√						X
ld ¹	√				√		DS
ldarx ¹	√				√		X
ldu ¹	√				√		DS
ldux ¹	√				√		X
ldx ¹	√				√		X
lfd	√						D
lfdu	√						D
lfdux	√						X
lfdx	√						X
lfs	√						D
lfsu	√						D
lfsux	√						X
lfsx	√						X
lha	√						D

Table A-46. PowerPC Instruction Set Legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
lhau	√						D
lhaux	√						X
lhax	√						X
lhbrx	√						X
lhz	√						D
lhzu	√						D
lhzux	√						X
lhzx	√						X
lmw ⁴	√						D
lswi ⁴	√						X
lswx ⁴	√						X
lwa¹	√				√		DS
lwarx	√						X
lwaux¹	√				√		X
lwax¹	√				√		X
lwbrx	√						X
lwz	√						D
lwzu	√						D
lwzux	√						X
lwzx	√						X
mcrf	√						XL
mcrfs	√						X
mcrxr	√						X
mfcrr	√						X
mffsx	√						X
mfmsr ²			√	√			X
mfspr ⁵	√		√	√			AFX
mfsr ²			√	√			X
mfsrin ²			√	√			X
mtfb		√					AFX
mtcrf	√						AFX
mtfsb0x	√						X
mtfsb1x	√						X
mtfsfx	√						AXL
mtfsfix	√						X

Table A-46. PowerPC Instruction Set Legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
mtmsr ²			√	√			X
mtspr ⁵	√		√	√			XFX
mtsr ²			√	√			X
mtsrin ²			√	√			X
mulhd _x ¹	√				√		XO
mulhd _{ux} ¹	√				√		XO
mulhw _x	√						XO
mulhw _{ux}	√						XO
mulld _x ¹	√				√		XO
mulli	√						D
mullw _x	√						XO
nand _x	√						X
neg _x	√						XO
nor _x	√						X
or _x	√						X
orc _x	√						X
ori	√						D
oris	√						D
rfi ²			√	√			XL
rdcl _x ¹	√				√		MDS
rdcr _x ¹	√				√		MDS
rdic _x ¹	√				√		MD
rdic _{lx} ¹	√				√		MD
rdic _{rx} ¹	√				√		MD
rdimix ¹	√				√		MD
rlwimix	√						M
rlwinm _x	√						M
rlwnm _x	√						M
sc	√		√				SC
slbia ^{1, 2, 3}			√	√	√	√	X
slbie ^{1, 2, 3}			√	√	√	√	X
sld _x ¹	√				√		X
slw _x	√						X
srad _x ¹	√				√		X
sradix ¹	√				√		XS

Table A-46. PowerPC Instruction Set Legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
srawx	√						X
srawix	√						X
srdx ¹	√				√		X
srwx	√						X
stb	√						D
stbu	√						D
stbux	√						X
stbx	√						X
std ¹	√				√		DS
stdcx. ¹	√				√		X
stdu ¹	√				√		DS
stdux ¹	√				√		X
stdx ¹	√				√		X
stfd	√						D
stfdu	√						D
stfdux	√						X
stfdx	√						X
stfiwx ³	√					√	X
stfs	√						D
stfsu	√						D
stfsux	√						X
stfsx	√						X
sth	√						D
sthbrx	√						X
sthu	√						D
sthux	√						X
sthx	√						X
stmw ⁴	√						D
stswi ⁴	√						X
stswx ⁴	√						X
stw	√						D
stwbrx	√						X
stwcx.	√						X
stwu	√						D
stwux	√						X

Table A-46. PowerPC Instruction Set Legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	64-Bit	Form
stwx	√						X
subfx	√						XO
subfcx	√						XO
subfex	√						XO
subfic	√						D
subfmex	√						XO
subfzex	√						XO
sync	√						X
td ¹	√				√		X
tdi ¹	√				√		D
tlbia ^{2, 3}			√	√		√	X
tlbie ^{2, 3}			√	√		√	X
tlbl ^{2, 6}				√			X
tlbli ^{2, 6}				√			X
tlbsync ^{2, 3}			√	√			X
tw	√						X
twi	√						D
xorx	√						X
xori	√						D
xoris	√						D

¹ 64-bit instruction

² Supervisor-level instruction

³ Optional in the PowerPC architecture

⁴ Load and store string or multiple instruction

⁵ Supervisor- and user-level instruction

⁶ G2-implementation specific instruction



Appendix B

Revision History

This appendix provides a list of the major differences between the *G2 PowerPC Core Reference Manual*, Revision 0 and the *G2 PowerPC Core Reference Manual*, Revision 1.

B.1 Revision Changes From Revision 0 to Revision 1

Major changes to the *G2 PowerPC Core Reference Manual* from Revision 0 to Revision 1 are as follows:

Section, Page	Changes
Book	Added trademark information for PowerPC. Added tab pages, Glossary, Appendix B, and Index.
xxv	Under the heading ‘Organization,’ in the first bullet, replace the statement in parenthesis in the second sentence with the following: (including instruction and data cache way-locking for the G2 core)
1.3.3.3, 1-26	Cache way-locking is a feature of both the G2 core and the G2_LE. Remove ‘G2_LE-Only’ from the heading and replace the paragraph with the following: The G2 core implements instruction and data cache way-locking, which guarantees that certain memory accesses will hit in the cache. This provides deterministic access times for those accesses. See Chapter 4, “Instruction and Data Cache Operation,” for more information.
1.4, 1-40	In Table 1-6, “Differences Between G2 and G2_LE Cores,” replace the rows on cache locking with the following:

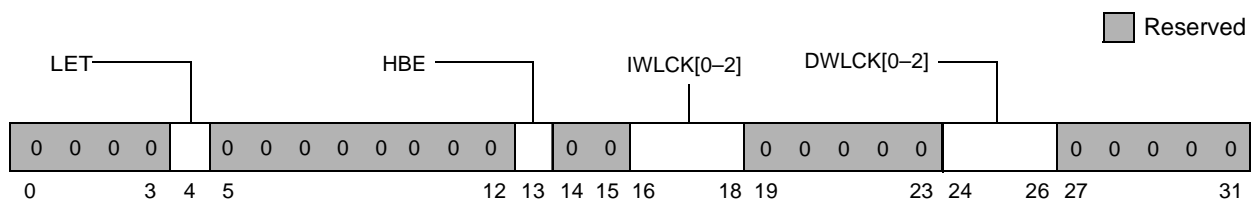
Supports instruction cache way-locking in addition to entire instruction cache locking	HID2 register controls instruction cache way-locking. The instruction cache way-locking is useful for locking blocks of instructions into the instruction cache for time-critical applications that require deterministic behavior.
Supports data cache way-locking in addition to entire data cache locking	HID2 register controls data cache way-locking. It is useful for locking blocks of data into the data cache for time-critical applications where deterministic behavior is required.

2.1.2.1, 2-11 In Table 2-5, “HID0 Bit Functions,” replace the description of bit 1 with the following:

1	—	Reserved
---	---	----------

2.1.2.3, 2-14 Replace the first sentence of the first paragraph with the following:
The G2 core implements an additional hardware implementation-dependent HID2 register, shown in Figure 2-4, which enables cache way-locking; the G2_LE core also enables true little-endian mode and the new additional BAT registers.

2.1.2.3, 2-15 Replace Figure 2-4, “Hardware Implementation-Dependent Register 2 (HID2)” with the following:



2.1.2.3, 2-15 In Table 2-8, “HID2 Bit Descriptions,” replace the description of bit 15 with the following:

15	—	Reserved
----	---	----------

Chapter 4, 4-1 Replace the last sentence of the second paragraph with the following:
It also describes the cache way-locking features provided in the G2 core.

4.2.3.3, 4-5 Replace the second paragraph with the following:
Note that the G2 core also provides instruction cache way-locking in addition to entire instruction cache locking as described in Section 4.12, “Cache Locking.”

4.3.3.3, 4-7 Replace the second paragraph with the following:
Note that the G2 core also provides instruction cache way-locking in addition to entire data cache locking as described in Section 4.12, “Cache Locking.”

4.5.2, 4-10 In Figure 4-3, “Double-Word Address Ordering—Critical-Double-Word-First,” remove ‘G2_LE Core Cache Address’ from the first heading and replace it with the ‘G2 Core Cache Address.’

- 4.12, 4-32 Replace the first paragraph with the following:
This section describes the entire cache locking and cache way-locking features of the G2 core.
- 4.12.1, 4-32 The title of the second bullet should be: ‘Way-Locking.’
- 4.12.2, 4-33 The title of Table 4-11 should read, “HID2 Bits Used to Perform Cache Way-Locking.”
- 4.12.3.1, 4-34 Replace the first paragraph with the following:
This section describes the procedures for performing data cache locking on the G2 core.
- 4.12.3.1.3, 4-35 In Table 4-14, “MSR Bits for Disabling Exceptions,” replace the description of bit 24 with the following:
- | | | |
|----|----|---------------------------|
| 24 | CE | Critical interrupt enable |
|----|----|---------------------------|
- 4.12.3.1.7, 4-37 Replace the first paragraph with the following:
Data cache way-locking is controlled by HID2[DWLCK], bits 24–26. Table 4-15 shows the HID2[DWLCK[0–2]] settings for the G2 core embedded processor.
- 4.12.3.1.7, 4-37 The title of Table 4-15 should read, “G2 Core DWLCK[0–2] Encodings.”
Replace the paragraph after Table 4-15 with the following:
The following assembly code locks way 0 of the G2 core data cache:
- 4.12.3.2, 4-38 Replace the first paragraph with the following:
This section describes the procedures for performing instruction cache locking on the G2 core.
- 4.12.3.2.3, 4-40 In Table 4-17, “MSR Bits for Disabling Exceptions,” replace the description of bit 24 with the following:
- | | | |
|----|----|---------------------------|
| 24 | CE | Critical interrupt enable |
|----|----|---------------------------|
- 4.12.3.2.6, 4-42 Remove ‘(G2_LE Only)’ from the heading and replace the first paragraph with the following:
Instruction cache way-locking is controlled by the HID2[IWLCK], bits 16–18. Table 4-18 shows the HID2[IWLCK[0–2]] settings for the G2 core embedded processor.
- 4.12.3.2.6, 4-42 The title of Table 4-18 should read, “G2 Core IWLCK[0–2] Encodings.” Replace the paragraph after Table 4-18 with the following:

The following assembly code locks way 0 of the G2 core instruction cache:

4.12.3.2.7, 4-42

Replace the last paragraph with the following:

In the second method, the instruction cache block invalidate (**icbi**) instruction can be used to invalidate individual cache blocks. The **icbi** instruction invalidates blocks in an entirely locked instruction cache. The **icbi** instruction also may invalidate way-locked blocks within the instruction cache.

6.5.2.2.2, 6-43

Replace the third line, `'bdnzf 0, im1,'` of function `'im1,'` with the following:

```
bdnzf    eq, im1
```

Replace the fourteenth line, `'srw r1, r1, 8,'` of function `'im1,'` with the following:

```
srwi    r1, r1, 8
```

8.3.15.3, 8-55

Replace 'Timing Comments,' with the following:

Assertion/Negation—Must remain stable during operation; should only be changed during the assertion of `core_hreset` or during sleep mode.

Glossary of Terms and Abbreviations

The glossary contains an alphabetical list of terms, phrases, and abbreviations used in this book. Some of the terms and definitions included in the glossary are reprinted from IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*, copyright ©1985 by the Institute of Electrical and Electronics Engineers, Inc. with the permission of the IEEE.

A **Architecture.** A detailed specification of requirements for a processor or computer system. It does not specify details of how the processor or computer system must be implemented; instead it provides a template for a family of compatible *implementations*.

Asynchronous exception. *Exceptions* that are caused by events external to the processor's execution. In this document, the term *asynchronous exception* is used interchangeably with the word *interrupt*.

Atomic access. A bus access that attempts to be part of a read-write operation to the same address uninterrupted by any other access to that address (the term refers to the fact that the transactions are indivisible). The PowerPC architecture implements *atomic accesses* through the **lwarx/stwex** instruction pair.

B **BAT (block address translation) mechanism.** A software-controlled array that stores the available block address translations on-chip.

Beat. A single state on the G2 bus interface that may extend across multiple bus cycles. A G2 transaction can be composed of multiple address or data *beats*.

Biased exponent. An *exponent* whose range of values is shifted by a constant (bias). Typically a bias is provided to allow a range of positive values to express a range that includes both positive and negative values.

Big-endian. A byte-ordering method in memory where the address n of a word corresponds to the *most-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 0, 1, 2, 3, with 0 being the *most-significant byte*. See *Little-endian*.

Block. An area of memory that ranges from 128 Kbytes to 256 Mbytes whose size, translation, and protection attributes are controlled by the *BAT* mechanism.

Boundedly undefined. A characteristic of certain operation results that are not rigidly prescribed by the PowerPC architecture. Boundedly-undefined results for a given operation may vary among implementations and between execution attempts in the same implementation.

Although the architecture does not prescribe the exact behavior for when results are allowed to be *boundedly undefined*, the results of executing instructions in contexts where results are allowed to be *boundedly undefined* are constrained to ones that could have been achieved by executing an arbitrary sequence of defined instructions, in valid form, starting in the state the machine was in before attempting to execute the given instruction.

Branch folding. The replacement with target instructions of a branch instruction and any instructions along the not-taken path when a branch is either taken or predicted as taken.

Branch prediction. The process of guessing whether a branch will be taken. Such predictions can be correct or incorrect; the term ‘predicted’ as it is used here does not imply that the prediction is correct (successful). The PowerPC architecture defines a means for *static branch* prediction as part of the instruction encoding.

Branch resolution. The determination of whether a branch is taken or not taken. A branch is said to be resolved when the processor can determine which instruction path to take. If the branch is resolved as predicted, the instructions following the predicted branch that may have been speculatively executed can complete (see *Completion*). If the branch is not resolved as predicted, instructions on the mispredicted path, and any results of speculative execution, are purged from the pipeline and fetching continues from the nonpredicted path.

Burst. A multiple-beat data transfer whose total size is typically equal to a cache block.

Bus clock. Clock that causes the bus state transitions.

Bus master. The owner of the address or data bus; the device that initiates or requests the transaction.

C

Cache. High-speed memory containing recently accessed data or instructions (subset of main memory).

Cache block. A small region of contiguous memory that is copied from memory into a *cache*. The size of a *cache block* may vary among processors; the maximum block size is one *page*. In PowerPC processors, *cache coherency* is maintained on a cache-block basis. Note that the term *cache block* is often used interchangeably with ‘cache line.’

Cache coherency. An attribute wherein an accurate and common view of memory is provided to all devices that share the same memory system. Caches are coherent if a processor performing a read from its cache is supplied with data corresponding to the most recent value written to memory or to another processor’s cache.

Cache flush. An operation that removes from a cache any data from a specified address range. This operation ensures that any modified data within the specified address range is written back to main memory. This operation is generated typically by a Data Cache Block Flush (**dcbf**) instruction.

Caching-inhibited. A memory update policy in which the cache is bypassed and the load or store is performed to or from main memory.

Cast out. A *cache block* that must be written to memory when a cache miss causes a *cache block* to be replaced.

Changed bit. One of two *page history bits* found in each *page table entry* (PTE). The processor sets the changed bit if any store is performed into the *page*. See also *Page access history bits* and *Referenced bit*.

Clean. An operation that causes a *cache block* to be written to memory, if modified, and then left in a valid, unmodified state in the cache.

Clear. To cause a bit or bit field to register a value of zero. See also *Set*.

Completion. Completion occurs when an instruction has finished executing, written back any results, and is removed from the completion queue (CQ). When an instruction completes, it is guaranteed that this instruction and all previous instructions can cause no exceptions.

Context synchronization. An operation that ensures that all instructions in execution complete past the point where they can produce an *exception*, that all instructions in execution complete in the context in which they began execution, and that all subsequent instructions are *fetched* and executed in the new context. Context synchronization may result from executing specific instructions (such as **isync** or **rfi**) or when certain events occur (such as an *exception*).

Copy-back operation. A cache operation in which a cache line is copied back to memory to enforce cache coherency. Copy-back operations consist of snoop push-out operations and cache cast-out operations.

D

Denormalized number. A nonzero floating-point number whose exponent has a reserved value, usually the format's minimum, and whose explicit or implicit leading significand bit is zero.

Direct-mapped cache. A cache in which each main memory address can appear in only one location within the cache, operates more quickly when the memory request is a cache hit.

Direct-store segment access. An access to an I/O address space. The G2 defines separate memory-mapped and I/O address spaces, or segments, distinguished by the corresponding segment register T bit in the address translation logic of the G2. If the T bit is cleared, the memory reference is a normal memory-mapped access and can use the virtual memory management hardware of the G2. If the T bit is set, the memory reference is a direct-store access.

E

Effective address (EA). The 32-bit address specified for a load, store, or an instruction fetch. This address is then submitted to the MMU for translation to either a *physical memory* address or an I/O address.

Exception. A condition encountered by the processor that requires special, supervisor-level processing.

Exception handler. A software routine that executes when an exception is taken. Normally, the exception handler corrects the condition that caused the exception, or performs some other meaningful task (that may include aborting the program that caused the exception). The address for each exception handler is identified by an exception vector offset defined by the architecture and a prefix selected via the MSR.

Exclusive state. MEI state (E) in which only one caching device contains data that is also in system memory.

Execution synchronization. A mechanism by which all instructions in execution are architecturally complete before beginning execution (appearing to begin execution) of the next instruction. Similar to context synchronization but doesn't force the contents of the instruction buffers to be deleted and refetched.

Exponent. In the binary representation of a floating-point number, the exponent is the component that normally signifies the integer power to which the value two is raised in determining the value of the represented number. See also *Biased exponent*.

F

Fall-through (branch fall-through). A not-taken branch. On the G2 core, fall-through branch instructions are removed from the instruction stream at dispatch. That is, these instructions are allowed to fall through the instruction queue through the dispatch mechanism, without either being passed to an execution unit and or given a position in the CQ.

Feed-forwarding. A G2 feature that reduces the number of clock cycles that an execution unit must wait to use a register. When the source register of the current instruction is the same as the destination register of the previous instruction, the result of the previous instruction is routed to the current instruction at the same time that it is written to the register file. With feed-forwarding, the destination bus is gated to the waiting execution unit over the appropriate source bus, saving the cycles which would be used for the write and read.

Fetch. Retrieving instructions from either the cache or main memory and placing them into the instruction queue.

Finish. Finishing occurs in the last cycle of execution. In this cycle, the CQ entry is updated to indicate that the instruction has finished executing.

Floating-point register (FPR). Any of the 32 registers in the floating-point register file. These registers provide the source operands and destination results for floating-point instructions. Load instructions move data from memory to FPRs and store instructions move data from FPRs to memory. The FPRs are 64 bits wide and store floating-point values in double-precision format.

Floating-point unit. The functional unit in the G2 processor responsible for executing all floating-point instructions.

Flush. An operation that causes a cache block to be invalidated and the data, if modified, to be written to memory.

Folding. See *Branch folding*.

Fraction. In the binary representation of a floating-point number, the field of the *significand* that lies to the right of its implied binary point.

G **General-purpose register (GPR).** Any of the 32 registers in the general-purpose register file. These registers provide the source operands and destination results for all integer data manipulation instructions. Integer load instructions move data from memory to GPRs and store instructions move data from GPRs to memory.

Guarded. The guarded attribute pertains to out-of-order execution. When a page is designated as guarded, instructions and data cannot be accessed out-of-order.

H **Harvard architecture.** An architectural model featuring separate caches and other memory management resources for instructions and data.

Hashing. An algorithm used in the *page table* search process.

I **IEEE 754.** A standard written by the Institute of Electrical and Electronics Engineers that defines operations and representations of binary floating-point numbers.

Illegal instructions. A class of instructions that are not implemented for a particular PowerPC processor. These include instructions not defined by the PowerPC architecture. In addition, for 32-bit implementations, instructions that are defined only for 64-bit implementations are considered to be illegal instructions. For 64-bit implementations instructions that are defined only for 32-bit implementations are considered to be illegal instructions.

Implementation. A particular processor that conforms to the PowerPC architecture, but may differ from other architecture-compliant implementations for example in design, feature set, and implementation of *optional* features. The PowerPC architecture has many different implementations.

Imprecise exception. A type of *synchronous exception* that is allowed not to adhere to the precise exception model (see *Precise exception*). The PowerPC architecture allows only floating-point exceptions to be handled imprecisely.

Instruction queue. A holding place for instructions fetched from the current instruction stream.

Integer unit. The functional unit in the G2 responsible for executing all integer instructions.

In-order. An aspect of an operation that adheres to a sequential model. An operation is said to be performed in-order if, at the time that it is performed, it is known to be required by the sequential execution model. See *Out-of-order*.

Instruction latency. The total number of clock cycles necessary to execute an instruction and make ready the results of that instruction.

Interrupt. An external signal that causes the G2 to suspend current execution and take a predefined exception.

K

Key bits. A set of key bits referred to as Ks and Kp in each segment register and each BAT register. The key bits determine whether supervisor or user programs can access a *page* within that *segment* or *block*.

Kill. An operation that causes a *cache block* to be invalidated without writing any modified data to memory.

L

Latency. The number of clock cycles necessary to execute an instruction and make ready the results of that execution for a subsequent instruction.

L2 cache. See *Secondary cache*.

Least-significant bit (lsb). The bit of least value in an address, register, field, data element, or instruction encoding.

Least-significant byte (LSB). The byte of least value in an address, register, data element, or instruction encoding.

Little-endian. A byte-ordering method in memory where the address *n* of a word corresponds to the *least-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 3, 2, 1, 0, with 3 being the *most-significant byte*. See *Big-endian*.

M

Mantissa. The decimal part of logarithm.

MEI (modified/exclusive/invalid). *Cache coherency* protocol used to manage caches on different devices that share a memory system. Note that the PowerPC architecture does not specify the implementation of a MEI protocol to ensure cache coherency.

Memory access ordering. The specific order in which the processor performs load and store memory accesses and the order in which those accesses complete.

Memory-mapped accesses. Accesses whose addresses use the page or block address translation mechanisms provided by the MMU and that occur externally with the bus protocol defined for memory.

Memory coherency. An aspect of caching in which it is ensured that an accurate view of memory is provided to all devices that share system memory.

Memory consistency. Refers to agreement of levels of memory with respect to a single processor and system memory (for example, on-chip cache, secondary cache, and system memory).

Memory management unit (MMU). The functional unit that is capable of translating an *effective* (logical) *address* to a physical address, providing protection mechanisms, and defining caching methods.

Modified state. MEI state (M) in which one, and only one, caching device has the valid data for that address. The data at this address in external memory is not valid.

Most-significant bit (msb). The highest-order bit in an address, registers, data element, or instruction encoding.

Most-significant byte (MSB). The highest-order byte in an address, registers, data element, or instruction encoding.

N

NaN. An abbreviation for not a number; a symbolic entity encoded in floating-point format. There are two types of NaNs—signaling NaNs and quiet NaNs.

No-op. No-operation. A single-cycle operation that does not affect registers or generate bus activity.

Normalization. A process by which a floating-point value is manipulated such that it can be represented in the format for the appropriate precision (single- or double-precision). For a floating-point value to be representable in the single- or double-precision format, the leading implied bit must be a 1.

O

OEA (operating environment architecture). The level of the architecture that describes PowerPC memory management model, supervisor-level registers, synchronization requirements, and the exception model. It also defines the time-base feature from a supervisor-level perspective. Implementations that conform to the PowerPC OEA also conform to the PowerPC UISA and VEA.

Optional. A feature, such as an instruction, a register, or an exception, that is defined by the PowerPC architecture but not required to be implemented.

Out-of-order. An aspect of an operation that allows it to be performed ahead of one that may have preceded it in the sequential model, for example, speculative operations. An operation is said to be performed out-of-order if, at the time that it is performed, it is not known to be required by the sequential execution model. See *In-order*.

Out-of-order execution. A technique that allows instructions to be issued and completed in an order that differs from their sequence in the instruction stream.

Overflow. An condition that occurs during arithmetic operations when the result cannot be stored accurately in the destination register(s). For example, if two 32-bit numbers are multiplied, the result may not be representable in 32 bits. Since the 32-bit registers of the G2 cannot represent this sum, an overflow condition occurs.

P

Page. A region in memory. The OEA defines a page as a 4-Kbyte area of memory, aligned on a 4-Kbyte boundary.

Page access history bits. The *changed* and *referenced* bits in the PTE keep track of the access history within the page. The referenced bit is set by the MMU whenever the page is accessed for a read or write operation. The changed bit is set when the page is stored into. See *Changed bit* and *Referenced bit*.

Page fault. A page fault is a condition that occurs when the processor attempts to access a memory location that does not reside within a *page* not currently resident in *physical memory*. On PowerPC processors, a page fault exception condition occurs when a matching, valid *page table entry* ($PTE[V] = 1$) cannot be located.

Page table. A table in memory is comprised of *page table entries*, or PTEs. It is further organized into eight PTEs per PTEG (page table entry group). The number of PTEGs in the page table depends on the size of the page table (as specified in the SDR1 register).

Page table entry (PTE). Data structures containing information used to translate *effective address* to physical address on a 4-Kbyte page basis. A PTE consists of 8 bytes of information in a 32-bit processor and 16 bytes of information in a 64-bit processor.

Park. The act of allowing a bus master to maintain bus mastership without having to arbitrate.

Physical memory. The actual memory that can be accessed through the system's memory bus.

Pipelining. A technique that breaks operations, such as instruction processing or bus transactions, into smaller distinct stages or tenures (respectively) so that a subsequent operation can begin before the previous one has completed.

Precise exceptions. A category of exception for which the pipeline can be stopped so instructions that preceded the faulting instruction can complete and subsequent instructions can be flushed and redispached after exception handling has completed. See *Imprecise exceptions*.

Primary opcode. The most-significant 6 bits (bits 0–5) of the instruction encoding that identifies the type of instruction.

Program order. The order of instructions in an executing program. More specifically, this term is used to refer to the original order in which program instructions are fetched into the instruction queue from the cache.

Protection boundary. A boundary between *protection domains*.

Protection domain. A protection domain is a segment, a virtual page, a BAT area, or a range of unmapped effective addresses. It is defined only when the appropriate relocate bit in the MSR (IR or DR) is 1.

Q

Quiesce. To come to rest. The processor is said to quiesce when an exception is taken or a **sync** instruction is executed. The instruction stream is stopped at the decode stage and executing instructions are allowed to complete to create a controlled context for instructions that may be affected by out-of-order, parallel execution. See *Context synchronization*.

Quiet NaN. A type of *NaN* that can propagate through most arithmetic operations without signaling exceptions. A quiet NaN is used to represent the results of certain invalid operations, such as invalid arithmetic operations on infinities or on NaNs, when invalid. See *Signaling NaN*.

R

rA. The rA instruction field is used to specify a GPR to be used as a source or destination.

rB. The rB instruction field is used to specify a GPR to be used as a source.

rD. The rD instruction field is used to specify a GPR to be used as a destination.

rS. The rS instruction field is used to specify a GPR to be used as a source.

Real address mode. An MMU mode when no address translation is performed and the *effective address* specified is the same as the physical address. The processor's MMU is operating in real address mode if its ability to perform address translation has been disabled through the MSR registers IR and/or DR bits.

Record bit. Bit 31 (or the Rc bit) in the instruction encoding. When it is set, updates the condition register (CR) to reflect the result of the operation.

Referenced bit. One of two *page history bits* found in each *page table entry*. The processor sets the *referenced bit* whenever the page is accessed for a read or write. See also *Page access history bits*.

Register indirect addressing. A form of addressing that specifies one GPR that contains the address for the load or store.

Register indirect with immediate index addressing. A form of addressing that specifies an immediate value to be added to the contents of a specified GPR to form the target address for the load or store.

Register indirect with index addressing. A form of addressing that specifies that the contents of two GPRs be added together to yield the target address for the load or store.

Rename register. Temporary buffers used by instructions that have finished execution but have not completed.

Reservation. The processor establishes a reservation on a *cache block* of memory space when it executes an **lwarx** instruction to read a memory semaphore into a GPR.

Reservation station. A buffer between the dispatch and execute stages that allows instructions to be dispatched even though the results of instructions on which the dispatched instruction may depend are not available.

Retirement. Removal of the completed instruction from the CQ.

RISC (reduced instruction set computing). An *architecture* characterized by fixed-length instructions with nonoverlapping functionality and by a separate set of load and store instructions that perform memory accesses.

S

Scan interface. The G2 test interface.

Secondary cache. A cache memory that is typically larger and has a longer access time than the primary cache. A secondary cache may be shared by multiple devices. Also referred to as L2, or level-2, cache.

Set (*v*). To write a nonzero value to a bit or bit field; the opposite of *clear*. The term ‘set’ may also be used to generally describe the updating of a bit or bit field.

Set (*n*). A subdivision of a *cache*. Cacheable data can be stored in a given location in one of the sets, typically corresponding to its lower-order address bits. Because several memory locations can map to the same location, cached data is typically placed in the set whose *cache block* corresponding to that address was used least recently. See *Set-associative*.

Set-associative. Aspect of cache organization in which the cache space is divided into sections, called *sets*. The cache controller associates a particular main memory address with the contents of a particular set, or region, within the cache.

Shadowing. Shadowing allows a register to be updated by instructions that are executed out of order without destroying machine state information.

Signaling NaN. A type of *NaN* that generates an invalid operation program exception when it is specified as arithmetic operands. See *Quiet NaN*.

Significand. The component of a binary floating-point number that consists of an explicit or implicit leading bit to the left of its implied binary point and a fraction field to the right.

Simplified mnemonics. Assembler mnemonics that represent a more complex form of a common operation.

Slave. The device addressed by a master device. The slave is identified in the address tenure and is responsible for supplying or latching the requested data for the master during the data tenure.

Snooping. Monitoring addresses driven by a bus master to detect the need for coherency actions.

Snoop push. Response to a snooped transaction that hits a modified cache block. The cache block is written to memory and made available to the snooping device.

Split-transaction. A transaction with independent request and response tenures.

Split-transaction bus. A bus that allows address and data transactions from different processors to occur independently.

Stage. The term *stage* is used in two different senses, depending on whether the pipeline is being discussed as a physical entity or a sequence of events. In the latter case, a stage is an element in the pipeline during which certain actions are performed, such as decoding the instruction, performing an arithmetic operation, or writing back the results. Typically, the latency of a stage is one processor clock cycle. Some events, such as dispatch, write-back, and completion, happen instantaneously and may be thought to occur at the end of a stage. An instruction can spend multiple cycles in one stage. An integer multiply, for example, takes multiple cycles in the execute stage. When this occurs, subsequent instructions may stall. An instruction may also occupy more than one stage simultaneously, especially in the sense that a stage can be seen as a physical resource—for example, when instructions are dispatched they are assigned a place

in the CQ at the same time they are passed to the execute stage. They can be said to occupy both the complete and execute stages in the same clock cycle.

Stall. An occurrence when an instruction cannot proceed to the next stage.

Static branch prediction. Mechanism by which software (for example, compilers) can hint to the machine hardware about the direction a branch is likely to take.

Store Queue. Holds store operations that have not been committed to memory, resulting from completed or retired instructions.

Superscalar. A superscalar processor is one that can dispatch multiple instructions concurrently from a conventional linear instruction stream. In a superscalar implementation, multiple instructions can be in the same stage at the same time.

Supervisor mode. The privileged operation state of a processor. In supervisor mode, software, typically the operating system, can access all control registers and can access the supervisor memory space, among other privileged operations.

Synchronization. A process to ensure that operations occur strictly *in order*. See *Context synchronization* and *Execution synchronization*.

Synchronous exception. An *exception* that is generated by the execution of a particular instruction or instruction sequence. There are two types of synchronous exceptions, *precise* and *imprecise*.

System memory. The physical memory available to a processor.

T

Tenure. The period of bus mastership. For the G2, there can be separate address bus tenures and data bus tenures. A tenure consists of three phases: arbitration, transfer, and termination.

TLB (translation lookaside buffer). A cache that holds recently-used *page table entries*.

Throughput. The measure of the number of instructions that are processed per clock cycle.

Transaction. A complete exchange between two bus devices. A transaction is typically comprised of an address tenure and one or more data tenures, which may overlap or occur separately from the address tenure. A transaction may be minimally comprised of an address tenure only.

Transfer termination. Signal that refers to both signals that acknowledge the transfer of individual beats (of both single-beat transfer and individual beats of a burst transfer) and to signals that mark the end of the tenure.

U

UISA (user instruction set architecture). The level of the architecture to which user-level software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, floating-point memory conventions and exception model as seen by user programs, and the memory and programming models.

Underflow. A condition that occurs during arithmetic operations when the result cannot be represented accurately in the destination register. For example, underflow can happen if two floating-point fractions are multiplied and the result requires a smaller *exponent* and/or *mantissa* than the single-precision format can provide. In other words, the result is too small to be represented accurately.

User mode. The operating state of a processor used typically by application software. In user mode, software can access only certain control registers and can access only user memory space. No privileged operations can be performed. Also referred to as problem state.

V

VEA (virtual environment architecture). The level of the *architecture* that describes the memory model for an environment in which multiple devices can access memory, defines aspects of the cache model, defines cache control instructions, and defines the time-base facility from a user-level perspective. *Implementations* that conform to the PowerPC VEA also adhere to the UISA, but may not necessarily adhere to the OEA.

Virtual address. An intermediate address used in the translation of an *effective address* to a physical address.

Virtual memory. The address space created using the memory management facilities of the processor. Program access to *virtual memory* is possible only when it coincides with *physical memory*.

W

Way. A location in the cache that holds a cache block, its tags and status bits.

Word. A 32-bit data element.

Write-back. A cache memory update policy in which processor write cycles are directly written only to the cache. External memory is updated only indirectly, for example, when a modified cache block is *cast out* to make room for newer data.

Write-through. A cache memory update policy in which all processor write cycles are written to both the cache and memory.

Index

A

$\overline{\text{AACK}}$ signal, 8-26
 $\overline{\text{ABB}}$ signal, 8-12, 9-7
 ABE (address broadcast enable) bit, 4-23
 Active-low signals, 8-1
 Address breakpoint register, 2-10
 Address broadcast enable, 2-13
 Address bus
 address transfer attribute
 A_n , 8-15
 $\overline{\text{APE}}$, 8-18, 9-12
 AP_n , 8-17
 $\overline{\text{CI}}$, 8-24
 CSE_n , 8-25
 $\overline{\text{GBL}}$, 8-25
 $\overline{\text{TBST}}$, 8-23, 9-13
 TC_n , 8-24, 9-19
 TSIZ_n , 8-22, 9-13
 TT_n , 8-19, 9-13
 $\overline{\text{WT}}$, 8-24
 address transfer start
 $\overline{\text{TS}}$, 8-14, 9-11
 address transfer termination
 $\overline{\text{AACK}}$, 8-26
 $\overline{\text{ARTRY}}$, 4-21, 8-26
 terminating address transfer, 9-19
 arbitration signals, 8-11, 9-6
 bus arbitration
 $\overline{\text{ABB}}$, 8-12, 9-7
 $\overline{\text{BG}}$, 8-11, 9-6
 $\overline{\text{BR}}$, 8-11, 9-6
 bus parking, 9-11
 tenure, 9-6
 Address bus parity signals, 8-1
 Address calculation
 branch instructions, 3-26
 effective address, 3-9
 floating-point load and store, 3-24
 integer load and store, 3-19
 Address matching, 11-5
 Address queue, 4-2
 Address translation, *see* Memory management unit

Addressing conventions
 addressing modes, 3-8
 alignment, 3-2
 Aligned data transfer, 3-1, 9-14, 9-18
 Alignment, 5-4
 data transfers, 3-1, 9-14
 exception, 5-28, 6-15
 rules, 3-2
 A_n signals, 8-15
 AND , 11-5, 11-6
 $\overline{\text{APE}}$ signal, 8-18, 9-12
 AP_n signals, 8-17
 Arbitration, system bus, 9-9, 9-21
 $\overline{\text{ARTRY}}$ signal, 4-21, 8-26
 Asserted, 8-1
 Asynchronous
 maskable, 5-3
 nonmaskable, 5-3
 Atomic memory references
 stwcx., 3-28
 using **lwarx/stwcx.**, 4-20
 Automatic power reduction mode, 10-1

B

Base/decrementer registers, 10-2
 BAT, 1-3, 4-11
 BAT registers
 G2_LE only (BAT4–BAT7), 2-18
 BE, 5-13
 IABR, 11-4
 $\overline{\text{IABR2}}$, 11-4
 $\overline{\text{BG}}$ signal, 8-11, 9-6
 bidirectional signals, 8-3
 BIU, 4-2, 4-8
 Block address translation, 4-34, 4-39, 6-20
 BAT registers
 implementation of BAT array, 2-18
 block address translation flow, 6-11
 lower, 4-35, 4-39
 selection of block address translation, 6-9
 upper, 4-35, 4-39
 Block size mask, 2-19

- Boundedly undefined, definition, 3-6
 - BPU, 1-1
 - \overline{BR} signal, 8-11, 9-6
 - Branch folding, 7-2, 7-17
 - Branch instructions
 - address calculation, 3-26
 - branch instructions, 3-26, A-22
 - condition register logical, 3-27, A-22
 - system linkage, 3-33, A-23
 - trap, 3-27, A-23
 - Branch prediction, 7-1, 7-18
 - Branch processing unit, 7-4
 - branch instruction timing, 7-20
 - execution timing, 7-16
 - latency, branch instructions, 7-26
 - overview, 1-9
 - Branch resolution definition, 7-1
 - Branch trace enable (BE), 2-7, 11-3, 11-5
 - Breakpoint
 - condition, 11-2
 - enabled, 11-4
 - exception, 11-2
 - registers, 11-1
 - Burst data transfers
 - 32-bit data bus, 9-14
 - 64-bit data bus, 9-14
 - transfers with data delays, timing, 9-35
 - Burst transactions, 4-9
 - Bus arbitration, *see* Data bus
 - Bus configurations, 9-37, 9-39
 - Bus interface unit (BIU), 4-2
 - Bus snooping, 10-2
 - Byte ordering
 - considerations, 5-21
 - default, 3-1, 3-9
 - Byte-reverse instructions, 3-21, A-20
- C**
- C bit, 6-39
 - Cache
 - cache locking
 - address translation
 - data cache locking, 4-34
 - instruction cache locking, 4-39
 - BAT examples, 4-34
 - data cache locking
 - address translation, 4-34
 - disabling exceptions, 4-35
 - enabling, 4-34
 - entire cache locking, 4-37
 - invalidation, 4-36
 - invalidation (if locked), 4-38
 - loading, 4-37
 - locking, 4-34
 - MSR bits, 4-35
 - way-locking, 4-37
 - disabling exceptions
 - data cache locking, 4-35
 - instruction cache locking, 4-40
 - enabling
 - data cache, 4-34
 - instruction cache, 4-38
 - entire cache locking definition, 4-32
 - instruction cache locking
 - address translation, 4-39
 - disabling, 4-40
 - enabling, 4-38
 - entire cache locking, 4-42
 - invalidating instruction cache (if locked), 4-43
 - MSR bits, 4-40
 - preloading instructions, 4-40
 - way-locking, 4-42
 - invalidation
 - data cache, 4-36
 - data cache (if locked), 4-38
 - instruction cache (if locked), 4-43
 - loading
 - data cache, 4-37
 - instruction cache preloading, 4-40
 - MSR bits
 - disabling exceptions, data cache locking, 4-35
 - disabling instruction cache locking, 4-40
 - organization, 4-32
 - procedures, 4-33
 - register summary, 4-32
 - terminology, 4-32
 - way-locking definition, 4-32
- cache miss, 7-13
- characteristics, 4-1
- instructions, 3-31, 3-35, 4-22, A-23
- MEI state definition, 4-16
- organization, instruction/data, 4-3-4-8
- overview, 1-24
- Cache arbitration, 7-10
- Cache block push operation, 4-9
- Cache block, definition, 4-1
- Cache cast-out operation, 4-9
- Cache coherency
 - actions on load operations, 4-19
 - actions on store operations, 4-19
 - copy-back operation, 4-12
 - in single-processor systems, 4-19
 - MEI protocol, 4-15
 - out-of-order execution, 4-14
 - overview, 4-3
 - protocol, 4-3
 - reaction to bus operations, 4-20

- WIMG bits, 4-10, 4-14, 9-29
 - write-back mode, 4-12
- Cache hit, 7-10
- Cache locking, 4-31
- Cache management instructions, 3-31, 3-35, 4-22, A-23
- Cache operations
 - basic data cache operations, 4-8
 - data cache transactions, 4-9
 - instruction cache fill operations, 4-4
 - overview, 1-13, 4-1
 - response to bus transactions, 4-20
- Cache unit
 - memory performance, 7-22
 - operation of the cache, 9-2
 - overview, 4-1
- Cache-inhibited, 6-16
- Cache-inhibited accesses (I bit)
 - cache interactions, 4-10
 - I-bit setting, 4-12
 - timing considerations, 7-23
- CE, 5-13
- DABR, 11-2
- DABR2, 11-2
- Changed (C) bit, 6-11, 6-21
- Changed (C) bit maintenance
 - recording, 6-21-6-24
- Changed (C) bit maintenance recording, 6-11
- Checkstop
 - signal, 8-41, 9-41
 - state, 5-24
- Checkstop high-impedance enable (core_ckstp_tre)
 - input, 8-42
- Checkstop output enable (core_ckstp_oe) output, 8-42
- \overline{CI} signal, 8-24
- Classes of instructions, 3-6
- Clean block operation, 4-20
- Clock signals
 - CLK_OUT, 8-54
 - PLL_CFG n , 8-55
 - SYSCLK, 8-53
- CMOS, 10-1
- Combinational matching, 11-5
- Company or manufacturer ID number, 2-4
- Compare and match type conditions, 11-2
- Compare instructions, 3-17, A-16
- Compare type and match type conditions, 11-3
- Completion
 - considerations, 7-13
 - definition, 7-1
 - unit, 11-3
- Completion queue, 7-1, Glossar y-3
- Context synchronization, 3-10
- Control bits, 10-2
- Conventions, xxxv, xli, 3-1

- COP/scan interface, 8-47
- COP_SVR instruction, 11-1
- Copy-back mode, 7-22
- core_cint signal, 8-39
- core_dbwo signal, 9-43
- CR logical instructions, 3-27
- Critical interrupt, 5-5, 5-16
 - exception enable (G2_LE only), 2-7
 - registers (G2_LE only), 2-10
- CSE n signals, 8-25
- CSRR0, 5-9, 5-11, 5-15, 5-16, 5-17
- CSRR1, 5-9, 5-11, 5-15, 5-16, 5-17, 5-18

D

- DABR, 11-1, 11-3
- DSISR, 11-3
- DABR{BT}, 11-3
- DABR2, 11-1, 11-3
- DABR2{BT}, 11-3
- DAR, 5-24, 11-2, 11-3
- Data accesses, 6-1
- Data address breakpoint
 - control register, 2-10
 - match, 11-2
 - registers, 11-1
 - registers (DABR, DABR2), 11-2
- Data address control register (DBCR), 11-3
- Data address register, 5-26
- Data address translation, 2-8
- Data block address translation, 4-35, 4-39
- Data breakpoint registers, 2-10
- Data bus
 - 32-bit data bus mode, 9-37
 - arbitration signals, 8-29, 9-7
 - bus arbitration, 9-21
 - data tenure, 9-6
 - data transfer, 8-31, 9-23
 - data transfer termination, 8-37, 9-24
- Data bus high, 8-32
- Data bus in (core_dh_in{0-31}, core_dl_in{0-31}), 8-32
- Data bus input enable (core_dh_ien, core_dl_ien) output, 8-32
- Data bus low, 8-32
- Data bus out (core_dh_out{0-31}, core_dl_out{0-31})
 - output, 8-33
- Data cache, 4-2
 - basic operations, 4-8
 - broadcasting, 4-7
 - bus transactions, 4-9
 - cache control, 4-6
 - configuration, 4-1
 - DCFI, DCE, DLOCK bits, 4-6
 - disabling, 4-7

- enable, 2-12
- fill operations, 4-8
- flash invalidate, 2-13
- lock, 2-12
- locking, 4-7
- organization, 4-6
- touch load operations, 4-8
- touch load support, 4-8
- way-lock, 2-16
- Data cache enable, 4-7
- Data cache flash invalidate, 4-6
- data coherency, 10-4
- Data load translation miss, 5-5
- Data storage interrupt (DSI), *see* DSI exception
- Data store translation miss, 5-5
- Data TLB miss on load exception, 5-36
- Data TLB miss on store exception, 5-37
- Data transfers
 - alignment, 3-1, 9-14
 - burst ordering, 9-14
 - eciwx and ecowx instructions, alignment, 9-18
 - signals, 9-23
- DBAT, 1-3
- \overline{DBB} signal, 8-30, 9-7, 9-22
- DBCR, 11-1
- \overline{DBDIS} signal, 8-36
- \overline{DBG} signal, 8-29, 9-7
- \overline{DBWO} signal, 8-29, 9-7, 9-23
- DCFI, 4-38
- DCMP, 6-34, 6-37
- DCMP and ICMP registers, 6-34
- Debug control registers, 11-1
- Debug control signals, 8-51
- Decrementer, 5-5
 - exception, 10-2
 - interrupt, 5-32, 10-2
 - timer, 10-2
- Default power state, 10-2
- Defined instruction class, 3-7
- Destination registers, 11-2
- DHn/DLn signals, 8-32
- Direct address translation (translation disabled)
 - data accesses, 4-11, 6-9, 6-11, 6-19
 - instruction accesses, 4-11, 6-9, 6-11, 6-19
- Direct-store access on the 603e, 4-10
- Dispatch considerations, 7-13
- DLOCK, 4-7, 4-37
- DMAISS, 6-37
- DMAISS/IMISS registers, 6-34
- DMMU, 6-25
- DBCR, 11-6
- Double-word, 4-2
- Doze mode, 10-2
- \overline{DPE} signal, 8-35

- DPn signals, 8-34
- DR, 5-14
- \overline{DRTRY} signal, 8-38, 9-24, 9-27
- DSI, 5-4, 11-2
- DSI exception, 5-24, 11-2, 11-3, 11-4
- DSISR, 5-1, 11-2
- DTLB, 1-3
- Dynamic power management, 10-1
 - enable (DPM), 2-12
 - modes, 10-2

E

- ECC errors, 9-28
- Effective address, 11-2
- Effective address calculation
 - address translation, 6-3
 - branches, 3-9, 3-26
 - loads and stores, 3-9, 3-19, 3-24
- Error termination, 9-27
- Exception, 11-2
- Exception little-endian mode, 2-7
- Exception prefix, 2-7
- Exception vector, 11-8
- Exception vector range, 11-2
- Exception vectors and priority, 11-6
- Exceptions
 - alignment exception, 5-28
 - classifications, 5-2
 - critical interrupt, 5-33
 - data TLB miss on load, 5-36
 - data TLB miss on store, 5-37
 - decrementer interrupt, 5-32
 - DSI, 5-24
 - enabling and disabling, 5-14
 - external interrupt, 5-27
 - FP unavailable, 5-32
 - instruction address breakpoint, 5-6, 5-37
 - instruction related, 3-10
 - instruction TLB miss, 5-36
 - machine check, 5-22
 - overview, 1-26
 - processing, 5-9, 5-15
 - program, 5-31
 - register settings
 - FPSCR, 5-31
 - MSR, 5-18
 - SRR0/SRR1, 5-10
 - reset, 5-19
 - returning from an exception handler, 5-16
 - summary, 3-10
 - system call, 5-34
 - system management interrupt, 5-39
 - trace, 5-34

- Execution synchronization, 3-10
- Execution unit, 1-9, 10-1
- Expanded debugging facilities in breakpoint registers, 11-4
- External asynchronous interrupts, 10-2
- External control instructions, 3-32, 9-18, A-24
- External interrupt, 5-4
- External interrupt enable, 2-7
- External system logic, 10-2

F

- FE0, 5-13, 5-14
- FE1, 5-13, 5-14
- Features list
 - G2
 - G2_LE, 1-3
- Finish cycle, definition, 7-2
- Floating-point available, 2-7
- Floating-point exception mode 0, 2-7
- Floating-point exception mode 1, 2-7
- Floating-point model
 - FE0/FE1 bits, 5-14
 - FP arithmetic instructions, 3-16, A-18
 - FP compare instructions, 3-18, A-19
 - FP execution models, 3-4
 - FP load instructions, 3-24, A-21
 - FP move instructions, 3-18, A-22
 - FP multiply-add instructions, 3-16, A-18
 - FP rounding/conversion instructions, 3-17, A-18
 - FP store instructions, 3-25, A-21
 - FP unavailable exception, 5-32
 - FPSCR instructions, 3-18, A-19
- Floating-point unavailable, 5-5
- Floating-point unit, 7-4
 - execution timing, 7-21
 - latency, FP instructions, 7-29
 - overview, 1-9
- Flow control instructions
 - branch instruction address calculation, 3-26
 - branch instructions, 3-26
 - condition register logical, 3-27
- Flush block operation, 4-20
- Force branch indirect on bus, 2-13
- Force-single-step operation instruction, 11-1
- FP, 5-13
- FPR, 5-1
- FPR0–FPR31, 2-2
- FPSCR, 5-1
- FPSCR instructions, 3-18, A-19
- FPU, 1-1
- Full-power mode, 10-2
 - with DPM disabled, 10-3
- Fully static, 10-1

G

- G (guarded memory), 4-3
- G2
 - features not present on PID6-603e, 1-5
 - overview, 1-1, 1-16
- G2_LE
 - overview, 1-1
- G2_LE-specific instructions, 3-37
- $\overline{\text{GBL}}$ signal, 8-25
- GPR, 5-1
- GPR0–GPR31, 2-2
- Guarded memory, 4-14
- Guarded memory bit (G bit)
 - cache interactions, 4-10
 - G-bit setting, 4-13

H

- Half-word, 4-2
- Handling, 5-2
- Hard reset and machine check, 5-17
- Hard reset sequence, 10-1
- Hardstop, 11-4
- Hardware handshake, 10-4
- HASH1/HASH2 registers, 6-35
- Hashing functions
 - primary PTEG, 6-30
 - secondary PTEG, 6-31
- HID0 register
 - DCFI, DCE, DLOCK bits, 4-6
 - doze bit, 10-4
 - DPM enable bit, 10-3
 - ICFI, ICE, ILOCK bits, 4-5
 - nap bit, 10-4
- HID0(DPM), 10-1
- HID1 register
 - bit settings, 2-14
 - PLL configuration, 2-14, 8-55
- High BAT enable, 2-15
- High-impedance control signal, 8-1
- $\overline{\text{HRESET}}$ signal, 8-42

I

- I (caching-inhibited), 4-3
- IABR, 11-1
- IABR2, 11-1
- IBAT, 1-3
- IBCR, 11-1
- IBCR(DNS), 11-6
- ICFI, 4-43
- ICMP, 6-34
- IEEE 1149.1-compliant interface, 9-42
- ILE, 5-13

- Illegal instruction class, 3-7
- ILOCK, 4-5
- ILOCK control bit, 4-5
- IMMU, 6-25
- Indeterminate processor core state, 11-8
- Input/output enable and high-impedance control signals, 8-3
- Instruction accesses, 6-1
- Instruction address breakpoint
 - control register (IBCR), 2-10
 - examples, 11-6
 - exception, 5-37, 11-2, 11-3, 11-4
 - exception handler, 11-2, 11-8
 - registers, 11-1
- Instruction address translation, 2-7
- instruction block address translation, 4-35, 4-39
- Instruction cache
 - cache control bits, 4-4
 - cache fill operations, 4-4
 - configuration, 4-2
 - organization, 4-4
- Instruction cache enable, 2-12
- Instruction cache flash invalidate, 2-12
- Instruction cache lock, 2-12
- Instruction cache way-lock, 2-15
- Instruction queue, 7-8
- Instruction timing
 - examples
 - cache hit, 7-11, 7-14
 - execution unit, 7-16
 - instruction flow, 7-8
 - memory performance considerations, 7-22
 - overview, 1-32, 7-3
 - terminology, 7-1
- Instruction TLB miss exception, 5-36
- Instruction translation miss, 5-5
- Instruction unit, 1-8
- Instructional address control register (IBCR), 11-2
- Instructions
 - branch address calculation, 3-26
 - branch instructions, 3-26, A-22
 - cache management instructions, 3-31, 3-35, 4-22, A-23
 - classes, 3-6
 - condition register logical, 3-27, A-22
 - defined instructions, 3-7
 - external control, 3-32, A-24
 - floating-point
 - arithmetic, 3-16, A-18
 - compare, 3-17, A-19
 - FP load instructions, 3-24, A-21
 - FP move instructions, 3-18, A-22
 - FP status and control register, 3-18
 - FP store instructions, 3-25, A-21
 - FPSCR instructions, 3-18, A-19
 - multiply-add, 3-16, A-18
 - rounding and conversion, 3-17, A-18
 - G2-specific instructions, 3-37
 - illegal instructions, 3-7
 - integer
 - arithmetic, 3-12, A-15
 - compare, 3-13, A-16
 - load, A-19
 - logical, 3-13, A-16
 - multiple, 3-22, A-20
 - rotate and shift, 3-14, A-17
 - store, 3-20, A-20
 - latency summary, 7-26
 - load and store
 - address generation, floating-point, 3-24
 - address generation, integer, 3-19
 - byte-reverse instructions, 3-21, A-20
 - integer load, 3-20
 - integer multiple instructions, 3-22, A-20
 - integer store, 3-20
 - string instructions, 3-23, A-21
 - memory control, 3-31, 3-35, 4-22, A-23
 - memory synchronization, 3-28, 3-30, A-21
 - PowerPC instructions, list
 - form (format), A-25
 - function, A-15
 - legend, A-36
 - mnemonic, A-1
 - opcode, A-8
 - processor control, 3-28, 3-30, 3-33, A-23
 - reserved instructions, 3-8
 - segment register manipulation, 3-36, A-24
 - simplified mnemonics, 3-37
 - supervisor-level cache management, 3-36
 - system linkage, 3-33, A-23
 - TLB management instructions, 3-36, A-24
 - trap instructions, 3-27, A-23
- INT signal, 8-39, 9-41
- Integer arithmetic instructions, 3-12, A-15
- Integer compare instructions, 3-13, A-16
- Integer load instructions, 3-20, A-19
- Integer logical instructions, 3-13, A-16
- Integer multiple instructions, 3-22, A-20
- Integer rotate and shift instructions, 3-14, A-17
- Integer store instructions, 3-20, A-20
- Integer unit, 7-4
 - execution timing, 7-20
 - latency, integer instructions, 7-27
 - overview, 1-9
- Interrupt and checkstop signals, 8-39
- Interrupt vector, 10-2
- Interrupt, critical, 5-33
- Interrupt, external, 5-27

Interrupt, *see* Exceptions

IP, 5-13, 5-15

ISI, 5-4

ITLB, 1-3

IU, 1-1

J

JTAG/COP

interface, 11-1

interface signals, 8-2

K

Kill block operation, 4-20

L

Latency, 7-2, 7-26, 9-24

LE, 5-14

Little-endian mode enable, 2-8

Load operations, memory coherency actions, 4-19

Load/store

address generation, 3-19, 3-24

byte-reverse instructions, 3-21, A-20

floating-point load instructions, 3-24, A-21

floating-point move instructions, 3-18, A-22

floating-point store instructions, 3-25, A-21

integer load instructions, 3-20, A-19

integer store instructions, 3-20, A-20

load/store multiple instructions, 3-22, A-20

memory synchronization instructions, 3-28, 3-30, A-21

string instructions, 3-23, A-21

Load/store unit, 7-4

execution timing, 7-21

latency, load and store instructions, 7-30

Logical addresses

translation into physical addresses, 6-1

Low-power operation, 10-1

LRU algorithm, 4-5

LSU, 1-1

lwarx/stwax.

support, 9-42

lwarx/stwax.

atomic memory references, 4-20

M

Machine check, 5-4

Machine check enable, 2-7

Machine check exception

checkstop state, 5-24

register settings

enabled, 5-23

SRR1 bit settings, 5-10

Machine state register, 4-35

Major processor design revision indicator, 2-5

Manufacturing revision, 2-5

Maskable asynchronous, 5-6

MCP signal, 8-40

ME, 5-13

MEI (modified, exclusive, or invalid, 4-3

MEI protocol

definition, MEI states, 4-16

enforcing memory coherency, 9-29

hardware considerations, 4-17

Memory accesses, 9-4

Memory coherency bit (M bit)

cache interactions, 4-10

I-bit setting, 4-12

M-bit setting, 4-12

timing considerations, 7-22

Memory control instructions

segment register manipulation, 3-36

TLB management, 3-36

user-level cache, 3-31, 3-35, 4-22

Memory management unit

address translation flow, 6-11

address translation mechanisms, 6-8, 6-11

block address translation, 6-9, 6-11, 6-20

block diagram, 6-5-6-7

data cache locking, 4-34

direct address translation, 4-11, 6-9, 6-11, 6-19

exceptions, 6-14

features summary, 6-2

general, 6-1

instruction cache locking, 4-39

instructions and registers, 6-17

memory protection, 6-10

overview, 1-31

page address translation, 6-8, 6-11, 6-27

page history status, 6-11, 6-21-6-24

page table search operation, 6-27

segment model, 6-21

software table search operation, 6-31, 6-36, 6-38

Memory reservation, 4-2

Memory synchronization

instructions, 3-28, 3-30, A-21

stwx., 3-28

Memory/cache access modes

performance impact of copy-back mode, 7-22

see also WIMG bits

Memory-coherent bit (M), 7-22

MESI, 1-3

- Minor processor design revision indicator, 2-5
- Misaligned accesses, 3-2
- Misaligned data transfer, 9-16
- MMU, 1-3, 4-8
- Mode control bits, 4-2
- Move instructions, 3-18
- MSR (machine state register)
 - bit settings, 2-6, 5-12
 - DR/IR bit, 2-7, 5-14
 - EE bit, 2-7, 5-13
 - FE0/FE1 bits, 5-14
 - POW bit, 2-6, 5-12
 - RI bit, 5-16
 - settings due to exception, 5-18
 - TGPR bit, 2-6, 5-13
- MSR{EE}, 10-6
- MSR{POW}, 10-6
- MSR{SE}, 11-4

N

- Nap mode, 10-2
- No-DRTRY mode, 9-39
- Nondenormalized mode, support, 3-15
- Nonmaskable, asynchronous, 5-6
- NOOPTI, 2-13, 4-8

O

- Operand conventions, 3-1
- Operand placement and performance, 3-4
- Operating environment architecture (OEA), xxxii, 3-32, 6-1
- Optional instructions, A-36
- OR
 - condition, 11-7
 - function, 11-5
 - operation, 11-6
- Other debug resources, 11-3
- Out-of-Order Data Accesses, 4-14
- Output enable signals, 8-1

P

- Page address translation
 - page address translation flow, 6-27
 - page size, 6-21
 - selection of page address translation, 6-8, 6-14
 - table search operation, 6-27
 - TLB organization, 6-26
- Page history status
 - R and C bit recording, 6-11, 6-21-6-24
- Page tables
 - resources for table search operations, 6-32
 - software table search operation, 6-31, 6-36

- SPRG(4-7) registers, 2-8
 - table search for PTE, 6-27
- Performance considerations, memory, 7-22
- Performance transparent functionality, 10-3
- Phase-locked loop, 10-2
- Physical address generation
 - memory management unit, 6-1
- Physical block number, 2-19
- Pipeline
 - instruction timing, definition, 7-2
 - pipeline stages, 7-7
 - superscalar/pipeline diagram, 7-5
- Pipelined execution unit, 7-4
- Power management
 - doze mode, 10-3
 - doze, nap, sleep, DPM bits, 2-14
 - full-power mode, 10-3
 - nap mode, 10-4
 - programmable power modes, 10-3
 - sleep mode, 10-5
 - software considerations, 10-6
- Power management modes, 1-14, 10-3
- Power-on reset, 5-8, 5-19
- Power-on reset settings, 5-20
- PowerPC architecture
 - instruction list, A-1, A-8, A-15
 - levels of implementation, 1-15
 - operating environment architecture (OEA), xxxii, 3-32
 - user instruction set architecture (UISA), xxxi, 2-1
 - virtual environment architecture (VEA), xxxi, 3-30
- Power-saving mode, 10-1
- PR, 5-13
- Privilege level (PR), 2-7, 11-3
- Privilege levels, supervisor-level cache instruction, 3-36
- Privileged state, *see* Supervisor mode
- Problem state, *see* User mode
- Process revision, 2-5
- Processor control instructions, 3-28, 3-30, 3-33, A-23
- Processor ID type, 2-4
- Processor identification, 2-5
- Program, 5-5
- Program exception, 5-31
- Program order, definition, 7-2
- Program-controllable power reduction mode, 10-1
- Programmable power modes, 10-2
- Programmable power states
 - doze mode, 10-3
 - full-power mode (DPM enabled/disabled), 10-3
 - nap mode, 10-4
 - sleep mode, 10-5
- Protection of memory areas
 - no-execute protection, 6-12

- options available, 6-10
- protection violations, 6-14
- PTEGs (PTE groups), 6-27
- PTEs (page table entries), 6-27

Q

- \overline{QACK} signal, 8-45, 9-37, 9-40
- \overline{QREQ} signal, 8-46, 9-41
- Qualified bus grant, 9-6
- Qualified data bus grant, 9-22
- Quiesce acknowledge signal, 10-4
- Quiesce request signal, 10-4
- Quiescent state, 10-4

R

- R bit, 6-39
- Read atomic operation, 4-20
- Read operation, 4-20
- Read with intent to modify operation, 4-20
- read-with-intent-to-modify (RWITM), 4-8
- Read-with-intent-to-modify (RWITM) EXAMPLES, 4-19
- Real address (RA), *see* Physical address generation
- Real addressing mode, *see* Direct address translation
- Recognition, 5-2
- Recoverable exception, 2-8
- Reduced-pinout mode, 9-40
- Referenced (R) bit, 6-11, 6-22
 - maintenance recording, 6-11, 6-21-6-24, 6-31
- Registers
 - cache locking register summary, 4-32
 - cache locking registers
 - HID0, 4-33
 - HID2, 4-33
 - MSR, 4-33
 - configuration registers
 - MSR, 2-6
 - PVR, 2-4
 - exception handling registers
 - DAR, 2-8
 - DSISR, 2-9
 - SPRG0–SPRG3, 2-20, 5-11
 - SRR0, 2-9
 - SRR0/SRR1, 2-19
 - implementation-specific registers
 - DCMP/ICMP, 2-16
 - DMAISS/IMISS, 2-16
 - HASH1/HASH2, 2-17
 - HID0/HID1, 2-10, 2-14
 - IABR, 2-21
 - RPA, 2-17

- memory management registers
 - BAT, 2-8
 - SDR1, 2-8
 - SR, 2-8
- supervisor-level
 - BAT, 2-8, 2-19
 - DAR, 2-8
 - DCMP/ICMP, 2-16, 6-34
 - DEC, 2-9
 - DMAISS/IMISS, 2-16, 6-34
 - DSISR, 2-9
 - EAR, 2-9
 - HASH1/HASH2, 2-17, 6-35
 - HID0/HID1, 2-10, 2-14
 - IABR, 2-21
 - MSR, 2-6
 - PVR, 2-4
 - RPA, 2-17
 - SDR1, 2-8
 - SPRG0–SPRG3, 2-20, 5-11
 - SR, 2-8
 - SRR0, 2-9
 - SRR0/SRR1, 2-19
 - TB, 2-9
- user-level
 - CR, 2-2
 - CTR, 2-4
 - FPR0–FPR31, 2-2
 - FPSCR, 2-2
 - GPR0–GPR31, 2-2
 - LR, 2-4
 - TB, 2-4
 - TGPR0–TGPR3, 6-33
 - XER, 2-4
- Rename buffer, 7-2
- Rename register operation, 7-15
- Reservation station, 7-2
- Reserved instruction class, 3-8
- Reset
 - \overline{HRESET} signal, 8-42
 - reset exception, 5-19
 - settings caused by hard reset, 5-20
 - \overline{SRESET} signal, 8-43, 9-41
- Reset configuration signals, 8-43
- Reset signals, 8-42
- Resource ID, 9-18
- Retirement, definition, 7-2
- RI, 5-14
- RISC, 1-1
- Rotate and shift instructions, 3-14, A-17
- RPA (required physical address), 6-36
- \overline{RSRV} signal, 8-45, 8-46, 9-42

S

SE, 5-13

Segment registers

SR manipulation instructions, 3-36, A-24

Segmented memory model, *see* Memory management unit

Self-modifying code, 3-19

Serializing instructions, 7-15

Signal groupings

Address arbitration signals, 8-1

Address transfer signals, 8-1

Address transfer start signals, 8-1

Address transfer termination signals, 8-1

Clock signals, 8-2

Data arbitration signals, 8-2

Data transfer signals, 8-2

Data transfer termination signals, 8-2

Debug control, 8-2

High-impedance control signals, 8-2

Input enable signals, 8-2

Output enable signals, 8-2

Processor status, 8-2

Reset configuration signals, 8-2

System status signals, 8-2

Test interface signals, 8-2

Transfer attribute signals, 8-1

Signals

 \overline{AACK} , 8-26 \overline{ABB} , 8-12, 9-7

address arbitration, 8-11, 9-6

address transfer, 9-11

address transfer attribute, 9-12

 A_n , 8-15 \overline{APE} , 8-18 AP_n , 8-17 \overline{ARTRY} , 8-26 \overline{BG} , 8-11, 9-6 \overline{BR} , 8-11, 9-6

checkstop, 9-41

 \overline{CI} , 8-24 $\overline{CKSTP_IN}$, 8-41 $\overline{CKSTP_OUT}$, 8-41

CLK_OUT, 8-54

COP/scan interface, 8-47

 $\overline{core_cint}$, 8-39 $\overline{core_dbwo}$, 9-43CSE $_n$, 8-25

data arbitration, 9-7, 9-21

data transfer termination, 9-24

 \overline{DBB} , 8-30, 9-7, 9-22 \overline{DBDIS} , 8-36 \overline{DBG} , 8-29, 9-7 \overline{DBWO} , 8-29, 9-7, 9-23 $\overline{DHn/DLn}$, 8-32 \overline{DPE} , 8-35 DP_n , 8-34 \overline{DRTRY} , 8-38, 9-24, 9-27 \overline{GBL} , 8-25 \overline{HRESET} , 8-42 \overline{INT} , 8-39, 9-41 \overline{MCP} , 8-40

non-protocol specific

TCK (JTAG test clock), 8-48

TDI (JTAG test data input), 8-48

TDO (JTAG test data output), 8-49

TMS (JTAG test mode select), 8-49

 \overline{TRST} (JTAG test reset), 8-49 $\overline{PLL_CFG_n}$, 8-55 \overline{QACK} , 8-45, 9-37, 9-40 \overline{QREQ} , 8-46, 9-41 \overline{RSRV} , 8-45, 8-46, 9-42 \overline{SMI} , 5-39, 8-40 \overline{SRESET} , 8-43, 9-41 \overline{TA} , 8-37 \overline{TBEN} , 8-46 \overline{TBST} , 8-23, 9-24 TC_n , 8-24, 9-19 \overline{TEA} , 8-38, 9-27 $\overline{TLBISYNC}$, 8-46 \overline{TS} , 8-14TSIZ $_n$, 8-22, 9-13 TT_n , 8-19, 9-13 \overline{WT} , 8-24

Single-beat reads with data delays, timing, 9-34

Single-beat transactions, 4-9

Single-beat transfer

reads with data delays, timing, 9-33

reads, timing, 9-31

termination, 9-25

writes, timing, 9-32

Single-step

enabled, 11-4

functions, 11-1

trace enable (SE), 2-7, 11-3

Sleep mode, 10-2

 \overline{SMI} signal, 5-39, 8-40

Snoop operation, 4-20, 7-22

Soft reset, 5-17

Softstop, 11-4

Software debug features, 11-3

Software debugging, 11-4

Software programming model interface, 11-3

Software table search, SPRG(4-7), 2-8

Split-bus transaction, 9-7

SPRG, 1-3, 5-11

 $\overline{SPRG0-SPRG3}$, conventional uses, 5-12 \overline{SRESET} signal, 8-43

SRR0, 5-9, 5-11, 5-15, 5-16, 5-17, 11-2
 SRR0/SRR1 (status save/restore registers)
 bit settings for machine check exception, 5-10
 bit settings for table search operations, 5-10
 format, 2-19, 5-11
 SRR1, 5-9, 5-14, 5-15, 5-16, 5-18
 SRU, 1-1
 Stall, definition, 7-3
 Static branch prediction, 7-18
 Static design, 10-5
 Status save/restore register 0, 5-15
 Store operations
 memory coherency actions, 4-19
 single-beat writes, 9-32
 String instructions, 3-23, A-21
 Superscalar, 7-3
 Supervisor mode, *see* Privilege levels
 Supervisor-level programs, 11-1
 Supervisor-level registers summary, 2-4
 Supervisor-level SPR, 11-2, 11-3
 Sync operation, 4-20
 Synchronization
 context/execution synchronization, 3-10
 execution of rfi, 5-16
 memory synchronization instructions, 3-28, 3-30, A-21
 requirements, 11-8
 Synchronous
 imprecise, 5-2
 precise, 5-2, 5-6
 SYSCLK, 8-53
 SYSCLK signal, 8-53
 System call, 5-5
 System call exception, 5-34
 System interface
 overview, 1-34
 System linkage instructions, 3-33, A-23
 System management interrupt, 5-6, 5-39, 10-2
 System memory base address, 2-10
 System quiesce control signals, 9-41
 System register unit, 7-4
 execution timing, 7-21
 latency, CR logical instructions, 7-27
 latency, system register instructions, 7-26
 System reset, 5-4
 System status
 $\overline{\text{CKSTP_IN}}$, 8-41
 $\overline{\text{CKSTP_OUT}}$, 8-41
 core_cint, 8-39
 $\overline{\text{HRESET}}$, 8-42
 $\overline{\text{INT}}$, 8-39
 $\overline{\text{MCP}}$, 8-40
 $\overline{\text{QACK}}$, 8-45
 $\overline{\text{QREQ}}$, 8-46

$\overline{\text{RSRV}}$, 8-45, 8-46
 $\overline{\text{SMI}}$, 8-40
 $\overline{\text{SRESET}}$, 8-43
 TBEN, 8-46
 $\overline{\text{TLBISYNC}}$, 8-46
 System version register, 2-10

T

$\overline{\text{TA}}$ signal, 8-37
 Table search operations
 algorithm, 6-27
 software routines, 6-31, 6-36-6-48
 SRR1 bit settings, 5-10
 table search flow (primary and secondary), 6-29
 Taken, 5-2
 TBEN signal, 8-46
 $\overline{\text{TBST}}$ signal, 8-23, 9-13, 9-24
 TCK (JTAG test clock) signal, 8-48
 TC n signals, 8-24, 9-19
 TDI (JTAG test data input) signal, 8-48
 TDO (JTAG test data output) signal, 8-49
 $\overline{\text{TEA}}$ signal, 8-38, 9-27
 Termination, 9-19, 9-24
 Test interface, 8-50
 TGPR0–GPR3 registers, 6-33
 Throughput, 7-3
 Time base
 lower, 2-9
 register, 10-2
 upper, 2-9
 Time-of-day maintenance, 10-5
 Timing diagrams, interface
 address transfer signals, 9-11
 burst transfers with data delays, 9-35
 single-beat reads, 9-31
 single-beat reads with data delays, 9-33
 single-beat writes, 9-32
 single-beat writes with data delays, 9-34
 use of $\overline{\text{TEA}}$, 9-36
 using $\overline{\text{DBWO}}$, 9-43
 Timing, instruction
 BPU execution timing, 7-16
 branch timing example, 7-20
 cache arbitration, 7-10
 cache hit, 7-10, 7-11, 7-14
 FPU execution timing, 7-21
 instruction dispatch, 7-13
 instruction flow, 7-8
 instruction scheduling guidelines, 7-23
 IU execution timing, 7-20
 latency summary, 7-26
 load/store unit execution timing, 7-21
 overview, 7-3

- SRU execution timing, 7-21
- stage, definition, 7-2
- TLB
 - description, 6-25
 - invalidate, A-24
 - invalidate (tlbie instruction), 6-26, 6-48
 - TLB management instructions, 3-37, A-24
- $\overline{\text{TLBISYNC}}$ signal, 8-46
- TMS (JTAG test mode select) signal, 8-49
- Trace, 5-5
- Trace exception, 5-34, 11-4, 11-5
- Trace facility, 11-1
- Transactions, data cache, 4-9
- Transfer, 9-11, 9-23
- Transfer type signals, 8-1
- Translation control bit, 6-8
- Trap instructions, 3-27
- $\overline{\text{TRST}}$ (JTAG test reset) signal, 8-49
- True little-endian, 2-15
- $\overline{\text{TS}}$ signal, 8-14, 9-11
- TSIZ n signals, 8-22, 9-13
- TT n signals, 8-19, 9-13

U

- Unidirectional/bidirectional signals, 8-5
- Unrecoverable state, 11-2, 11-8
- Use of $\overline{\text{TEA}}$, timing, 9-36
- User mode, 5-1
- User instruction set architecture (UISA), xxxi, 2-1
- User-level registers summary, 2-2
- Using $\overline{\text{DBWO}}$, timing, 9-43

V

- Virtual page number, 6-29
- Virtual environment architecture (VEA), xxxi, 3-30

W

- W (write-through), 4-3
- Watchpoint signaling, 11-5
- Watchpoint/breakpoint indication signals, 11-1
- WIMG, 4-6
- WIMG bits, 4-10, 9-29
- Wire-ORing, 8-5
- Word, 4-2
- Write with atomic operation, 4-20
- Write with flush operation, 4-20
- Write with kill operation, 4-20
- Write-back mode, 4-12, 7-3
- Write-through memory area, 6-16
- Write-through mode (W bit)
 - cache interactions, 4-10
 - timing considerations, 7-23
- W-bit setting, 4-11
- $\overline{\text{WT}}$ signal, 8-24



Overview	1
Register Model	2
Instruction Set Model	3
Instruction and Data Cache Operation	4
Exceptions	5
Memory Management	6
Instruction Timing	7
Signal Descriptions	8
Core Interface Operation	9
Power Management	10
Debug Features	11
PowerPC Instruction Set Listings	A
Revision History	B
Glossary of Terms and Abbreviations	GLO
Index	IND



1 Overview

2 Register Model

3 Instruction Set Model

4 Instruction and Data Cache Operation

5 Exceptions

6 Memory Management

7 Instruction Timing

8 Signal Descriptions

9 Core Interface Operation

10 Power Management

11 Debug Features

A PowerPC Instruction Set Listings

B Revision History

GLO Glossary of Terms and Abbreviations

IND Index