

Oracle SPARC Architecture 2015

*One Architecture
... Multiple Innovative Implementations*

Draft D1.0.0, 12 Jan 2016

*Privilege Levels: Privileged
and Nonprivileged*

Distribution: Public

Some portions of this specification are not yet finalized; please check monthly to see if an updated revision is available for download.

Copyright © Oracle and/or its affiliates. All rights reserved.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd..

Comments and "bug reports" regarding this document are welcome; they should be submitted to email address: David.Weaver@Oracle.com

Contents

Preface	i
1 Document Overview	1
1.1 Navigating <i>Oracle SPARC Architecture 2015</i>	1
1.2 Fonts and Notational Conventions	2
1.2.1 Implementation Dependencies	3
1.2.2 Notation for Numbers	3
1.2.3 Informational Notes	3
1.3 Reporting Errors in this Specification	4
2 Definitions	5
3 Architecture Overview	13
3.1 The Oracle SPARC Architecture 2015	13
3.1.1 Features	13
3.1.2 Attributes	14
3.1.2.1 Design Goals	15
3.1.2.2 Register Windows	15
3.1.3 System Components	15
3.1.3.1 Binary Compatibility	15
3.1.3.2 Oracle SPARC Architecture 2015 MMU	15
3.1.3.3 Privileged Software	15
3.1.4 Architectural Definition	16
3.1.5 Oracle SPARC Architecture 2015 Compliance with SPARC V9 Architecture ..	16
3.1.6 Implementation Compliance with Oracle SPARC Architecture 2015	16
3.2 Processor Architecture	16
3.2.1 Integer Unit (IU)	16
3.2.2 Floating-Point Unit (FPU)	17
3.3 Instructions	17
3.3.1 Memory Access	17
3.3.1.1 Memory Alignment Restrictions	18
3.3.1.2 Addressing Conventions	18
3.3.1.3 Addressing Range	18
3.3.1.4 Load/Store Alternate	18
3.3.1.5 Separate Instruction and Data Memories	19
3.3.1.6 Input/Output (I/O)	19
3.3.1.7 Memory Synchronization	19
3.3.2 Integer Arithmetic / Logical / Shift Instructions	19
3.3.3 Control Transfer	20
3.3.4 State Register Access	20
3.3.4.1 Ancillary State Registers	20
3.3.4.2 PR State Registers	20

3.3.5	Floating-Point Operate	21
3.3.6	Conditional Move	21
3.3.7	Register Window Management	21
3.3.8	SIMD	21
3.4	Traps	21
4	Data Formats	23
4.1	Integer Data Formats	24
4.1.1	Signed Integer Data Types	24
4.1.1.1	Signed Integer Byte, Halfword, and Word	25
4.1.1.2	Signed Integer Doubleword (64 bits)	25
4.1.1.3	Signed Integer Extended-Word (64 bits)	25
4.1.2	Unsigned Integer Data Types	25
4.1.2.1	Unsigned Integer Byte, Halfword, and Word	26
4.1.2.2	Unsigned Integer Doubleword (64 bits)	26
4.1.2.3	Unsigned Extended Integer (64 bits)	26
4.1.3	Tagged Word (32 bits)	26
4.2	Floating-Point Data Formats	27
4.2.1	Floating Point, Single Precision (32 bits)	27
4.2.2	Floating Point, Double Precision (64 bits)	27
4.2.3	Floating Point, Quad Precision (128 bits)	28
4.2.4	Floating-Point Data Alignment in Memory and Registers	29
4.3	SIMD Data Formats	29
4.3.1	Uint8 SIMD Data Format	30
4.3.2	Int16 SIMD Data Formats	30
4.3.3	Int32 SIMD Data Format	30
5	Registers	31
5.1	Reserved Register Fields	32
5.2	General-Purpose R Registers	32
5.2.1	Global R Registers	33
5.2.2	Windowed R Registers	34
5.2.3	Special R Registers	37
5.3	Floating-Point Registers	38
5.3.1	Floating-Point Register Number Encoding	40
5.3.2	Double and Quad Floating-Point Operands	41
5.4	Floating-Point State Register (FSR)	42
5.4.1	Floating-Point Condition Codes (fcc0, fcc1, fcc2, fcc3)	42
5.4.2	Rounding Direction (rd)	43
5.4.3	Trap Enable Mask (tem)	43
5.4.4	Nonstandard Floating-Point (ns)	43
5.4.5	FPU Version (ver)	43
5.4.6	Floating-Point Trap Type (ftt)	44
5.4.7	Accrued Exceptions (aexc)	46
5.4.8	Current Exception (cexc)	46
5.4.9	Floating-Point Exception Fields	47
5.4.10	FSR Conformance	48
5.5	Ancillary State Registers	48
5.5.1	32-bit Multiply/Divide Register (Y) (ASR 0)	50
5.5.2	Integer Condition Codes Register (CCR) (ASR 2)	50
5.5.2.1	Condition Codes (CCR.xcc and CCR.icc)	50
5.5.3	Address Space Identifier (ASI) Register (ASR 3)	51
5.5.4	Tick (TICK) Register (ASR 4)	52
5.5.5	Program Counters (PC, NPC) (ASR 5)	53
5.5.6	Floating-Point Registers State (FPRS) Register (ASR 6)	53
5.5.7	General Status Register (GSR) (ASR 19)	54

5.5.8	SOFTINT ^P Register (ASRs 20, 21, 22)	55
5.5.8.1	SOFTINT_SET ^P Pseudo-Register (ASR 20)	55
5.5.8.2	SOFTINT_CLR ^P Pseudo-Register (ASR 21)	56
5.5.9	System Tick (STICK) Register (ASR 24)	56
5.5.10	System Tick Compare (STICK_CMPR ^P) Register (ASR 25)	57
5.5.11	Compatibility Feature Register (CFR) (ASR 26)	59
5.5.12	Pause Count (PAUSE) Register (ASR 27)	60
5.5.13	Mwait Count (MWAIT) Register (ASR 28)	61
5.6	Register-Window PR State Registers	62
5.6.1	Current Window Pointer (CWP ^P) Register (PR 9)	63
5.6.2	Savable Windows (CANSAVE ^P) Register (PR 10)	63
5.6.3	Restorable Windows (CANRESTORE ^P) Register (PR 11)	64
5.6.4	Clean Windows (CLEANWIN ^P) Register (PR 12)	64
5.6.5	Other Windows (OTHERWIN ^P) Register (PR 13)	65
5.6.6	Window State (WSTATE ^P) Register (PR 14)	65
5.6.7	Register Window Management	65
5.6.7.1	Register Window State Definition	66
5.6.7.2	Register Window Traps	66
5.7	Non-Register-Window PR State Registers	66
5.7.1	Trap Program Counter (TPC ^P) Register (PR 0)	67
5.7.2	Trap Next PC (TNPC ^P) Register (PR 1)	67
5.7.3	Trap State (TSTATE ^P) Register (PR 2)	68
5.7.4	Trap Type (TT ^P) Register (PR 3)	69
5.7.5	Tick (TICK) Register (PR 4)	69
5.7.6	Trap Base Address (TBA ^P) Register (PR 5)	69
5.7.7	Processor State (PSTATE ^P) Register (PR 6)	69
5.7.8	Trap Level Register (TL ^P) (PR 7)	73
5.7.9	Processor Interrupt Level (PIL ^P) Register (PR 8)	74
5.7.10	Global Level Register (GL ^P) (PR 16)	75
6	Instruction Set Overview	77
6.1	Instruction Execution	77
6.2	Instruction Formats	78
6.3	Instruction Categories	78
6.3.1	Memory Access Instructions	79
6.3.1.1	Memory Alignment Restrictions	79
6.3.1.2	Addressing Conventions	80
6.3.1.3	Address Space Identifiers (ASIs)	83
6.3.1.4	Separate Instruction Memory	84
6.3.2	Memory Synchronization Instructions	85
6.3.3	Integer Arithmetic and Logical Instructions	85
6.3.3.1	Setting Condition Codes	85
6.3.3.2	Shift Instructions	85
6.3.3.3	Set High 22 Bits of Low Word	85
6.3.3.4	Integer Multiply/Divide	85
6.3.3.5	Tagged Add/Subtract	85
6.3.4	Control-Transfer Instructions (CTIs)	86
6.3.4.1	Conditional Branches	87
6.3.4.2	Unconditional Branches	88
6.3.4.3	CALL and JMPL Instructions	88
6.3.4.4	RETURN Instruction	88
6.3.4.5	DONE and RETRY Instructions	88
6.3.4.6	Trap Instruction (Tcc)	88
6.3.4.7	DCTI Couples	89
6.3.5	Conditional Move Instructions	89
6.3.6	Register Window Management Instructions	90
6.3.6.1	SAVE Instruction	90

6.3.6.2	RESTORE Instruction	90
6.3.6.3	SAVED Instruction	91
6.3.6.4	RESTORED Instruction	91
6.3.6.5	Flush Windows Instruction	91
6.3.7	Ancillary State Register (ASR) Access	92
6.3.8	Privileged Register Access	92
6.3.9	Floating-Point Operate (FPop) Instructions	92
6.3.10	Implementation-Dependent Instructions	92
6.3.11	Reserved Opcodes and Instruction Fields	93
7	Instructions	95
7.46.1	FMUL8x16 Instruction	186
7.46.2	FMUL8x16AU Instruction	187
7.46.3	FMUL8x16AL Instruction	187
7.46.4	FMUL8SUX16 Instruction	188
7.46.5	FMUL8ULx16 Instruction	188
7.46.6	FMULD8SUX16 Instruction	189
7.46.7	FMULD8ULx16 Instruction	189
7.52.1	FPACK16	198
7.52.2	FPACK32	199
7.52.3	FPACKFIX	200
7.90.1	Memory Synchronization	270
7.90.2	Synchronization of the Virtual Processor	271
7.90.3	TSO Ordering Rules affecting Use of MEMBAR	271
7.108.1	Exceptions	305
7.108.2	Weak versus Strong Prefetches	306
7.108.3	Prefetch Variants	306
7.108.3.1	Prefetch for Several Reads (f _{cn} = 0, 20(14 ₁₆))	307
7.108.3.2	Prefetch for One Read (f _{cn} = 1, 21(15 ₁₆))	307
7.108.3.3	Prefetch for Several Writes (and Possibly Reads) (f _{cn} = 2, 22(16 ₁₆))	307
7.108.3.4	Prefetch for One Write (f _{cn} = 3, 23(17 ₁₆))	307
7.108.3.5	Prefetch Page (f _{cn} = 4)	308
7.108.3.6	Prefetch to Nearest Unified Cache (f _{cn} = 17(11 ₁₆))	308
7.108.4	Implementation-Dependent Prefetch Variants (f _{cn} = 16, 18, 19, and 24–31)	308
7.108.5	Additional Notes	308
8	IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015	391
8.1	Traps Inhibiting Results	391
8.2	Underflow Behavior	392
8.2.1	Trapped Underflow Definition (uf _m = 1)	393
8.2.2	Untrapped Underflow Definition (uf _m = 0)	393
8.3	Integer Overflow Definition	393
8.4	Floating-Point Nonstandard Mode	393
8.5	Arithmetic Result Tables	394
8.5.1	Floating-Point Add (FADD) <i>and Add and Halve (FHADD)</i>	395
8.5.2	Floating-Point Negative Add (FNADD) <i>and Negative Add and Halve (FNHADD)</i>	396
8.5.3	Floating-Point Subtract (FSUB) <i>and Subtract and Halve (FHSUB)</i>	396
8.5.4	Floating-Point Multiply	397
8.5.5	Floating-Point Negative Multiply (FNMUL)	397
8.5.6	Floating-Point Multiply-Add (FMAf)	398
8.5.7	Floating-Point Negative Multiply-Add (FNMADD)	398
8.5.8	Floating-Point Multiply-Subtract (FMSUB)	399
8.5.9	Floating-Point Negative Multiply-Subtract (FNMSUB)	400
8.5.10	Floating-Point Divide (FDIV)	402

8.5.11	Floating-Point Square Root (FSQRT)	402
8.5.12	Floating-Point Compare (FCMP, FCMPE)	403
8.5.13	Floating-Point Lexicographic Compare (FLCMP)	403
8.5.14	Floating-Point to Floating-Point Conversions (F<s d q>TO<s d q>)	404
8.5.15	Floating-Point to Integer Conversions (F<s d q>TO<i x>)	405
8.5.16	Integer to Floating-Point Conversions (F<i x>TO<s d q>)	405
9	Memory	407
9.1	Memory Location Identification	407
9.2	Memory Accesses and Cacheability	407
9.2.1	Coherence Domains	408
9.2.1.1	Cacheable Accesses	408
9.2.1.2	Noncacheable Accesses	408
9.2.1.3	Noncacheable Accesses with Side-Effect	408
9.3	Memory Addressing and Alternate Address Spaces	409
9.3.1	Memory Addressing Types	410
9.3.2	Memory Address Spaces	410
9.3.3	Address Space Identifiers	410
9.4	SPARC V9 Memory Model	412
9.4.1	SPARC V9 Program Execution Model	412
9.4.2	Virtual Processor/Memory Interface Model	414
9.5	The Oracle SPARC Architecture Memory Model — TSO	415
9.5.1	Memory Model Selection	415
9.5.2	Programmer-Visible Properties of the Oracle SPARC Architecture TSO Model	416
9.5.3	TSO Ordering Rules	417
9.5.4	Hardware Primitives for Mutual Exclusion	418
9.5.4.1	Compare-and-Swap (CASA, CASXA)	418
9.5.4.2	Swap (SWAP)	418
9.5.4.3	Load Store Unsigned Byte (LDSTUB)	418
9.5.5	Memory Ordering and Synchronization	418
9.5.5.1	Ordering MEMBAR Instructions	419
9.5.5.2	Sequencing MEMBAR Instructions	419
9.5.5.3	Synchronizing Instruction and Data Memory	420
9.6	Nonfaulting Load	421
9.7	Store Coalescing	422
10	Address Space Identifiers (ASIs)	423
10.1	Address Space Identifiers and Address Spaces	423
10.2	ASI Values	423
10.3	ASI Assignments	424
10.3.1	Supported ASIs	424
10.4	Special Memory Access ASIs	432
10.4.1	ASIs 10 ₁₆ , 11 ₁₆ , 16 ₁₆ , and 17 ₁₆ (ASI_*AS_IF_USER_*)	432
10.4.2	ASIs 18 ₁₆ , 19 ₁₆ , 1E ₁₆ , and 1F ₁₆ (ASI_*AS_IF_USER_*_LITTLE)	433
10.4.3	ASI 14 ₁₆ (ASI_REAL)	433
10.4.4	ASI 15 ₁₆ (ASI_REAL_IO)	434
10.4.5	ASI 1C ₁₆ (ASI_REAL_LITTLE)	434
10.4.6	ASI 1D ₁₆ (ASI_REAL_IO_LITTLE)	434
10.4.7	ASIs 22 ₁₆ , 23 ₁₆ , 26 ₁₆ , 27 ₁₆ , 2A ₁₆ , 2B ₁₆ , 2E ₁₆ 2F ₁₆ (Privileged Load Integer Twin Extended Word)	434
10.4.8	ASIs 26 ₁₆ and 2E ₁₆ (Privileged Load Integer Twin Extended Word, Real Addressing)	435
10.4.9	ASIs E2 ₁₆ , E3 ₁₆ , EA ₁₆ , EB ₁₆ (Nonprivileged Load Integer Twin Extended Word)	436
10.4.10	Block Load and Store ASIs	436
10.4.11	Partial Store ASIs	437
10.4.12	Short Floating-Point Load and Store ASIs	437

10.4.13	Load-Monitor ASIs	437
10.5	ASI-Accessible Registers	437
10.5.1	Privileged Scratchpad Registers (ASI_SCRATCHPAD)	438
10.6	ASI Changes in the Oracle SPARC Architecture	438
11	Performance Instrumentation	439
11.1	High-Level Requirements	439
11.1.1	Usage Scenarios	439
11.1.2	Metrics	440
11.1.3	Accuracy Requirements	440
11.2	Performance Counters and Controls	441
11.2.1	Counter Overflow	441
12	Traps	443
12.1	Virtual Processor Privilege Modes	444
12.2	Virtual Processor States and Traps	445
12.2.0.1	Usage of Trap Levels	445
12.3	Trap Categories	445
12.3.1	Precise Traps	446
12.3.2	Deferred Traps	446
12.3.3	Disrupting Traps	447
12.3.3.1	Disrupting versus Precise and Deferred Traps	447
12.3.3.2	Causes of Disrupting Traps	448
12.3.3.3	Conditioning of Disrupting Traps	448
12.3.3.4	Trap Handler Actions for Disrupting Traps	449
12.3.4	Uses of the Trap Categories	449
12.4	Trap Control	449
12.4.1	PIL Control	450
12.4.2	FSR.tem Control	450
12.5	Trap-Table Entry Addresses	450
12.5.1	Trap-Table Entry Address to Privileged Mode	450
12.5.2	Privileged Trap Table Organization	451
12.5.3	Trap Type (TT)	451
12.5.3.1	Trap Type for Spill/Fill Traps	458
12.5.4	Trap Priorities	458
12.6	Trap Processing	458
12.6.1	Normal Trap Processing	459
12.7	Exception and Interrupt Descriptions	460
12.7.1	SPARC V9 Traps Not Used in Oracle SPARC Architecture 2015	464
12.8	Register Window Traps	465
12.8.1	Window Spill and Fill Traps	465
12.8.2	<i>clean_window</i> Trap	465
12.8.3	Vectoring of Fill/Spill Traps	465
12.8.4	CWP on Window Traps	466
12.8.5	Window Trap Handlers	466
13	Interrupt Handling	467
13.1	Interrupt Packets	467
13.2	Software Interrupt Register (SOFTINT)	468
13.2.1	Setting the Software Interrupt Register	468
13.2.2	Clearing the Software Interrupt Register	468
13.3	Interrupt Queues	468
13.3.1	Interrupt Queue Registers	469
13.4	Interrupt Traps	470

Preface

First came the 32-bit SPARC Version 7 (V7) architecture, publicly released in 1987. Shortly after, the SPARC V8 architecture was announced and published in book form. The 64-bit SPARC V9 architecture was released in 1994. The UltraSPARC Architecture 2005 specification (predecessor to this Oracle SPARC Architecture 2015) provided the first significant update in over 10 years to Oracle's SPARC processor architecture.

What's New?

Oracle SPARC Architecture 2015 pulls together in one document all parts of the architecture:

- the nonprivileged (Level 1) architecture from SPARC V9
- most of the privileged (Level 2) architecture from SPARC V9
- more in-depth coverage of all SPARC V9 features

Plus, it includes all of Sun's now-standard architectural extensions (beyond SPARC V9), developed through the processor generations of UltraSPARC III, IV, IV+, and T1:

- the VIS™ 1 and VIS 2 instruction set extensions and the associated GSR register
- multiple levels of global registers, controlled by the GL register
- Sun's 64-bit MMU architecture
- privileged instructions ALLCLEAN, OTHERW, NORMALW, and INVALIDW
- access to the VER register is now hyperprivileged
- the SIR instruction is now hyperprivileged

Oracle SPARC Architecture 2011 includes the following changes since UltraSPARC Architecture 2007:

- the VIS 3 instructions set extensions

In addition, architectural features are now tagged with Software Classes and Implementation Classes¹. Software Classes provide a new, high-level view of the expected architectural longevity and portability of software that references those features. Implementation Classes give an indication of how efficiently each feature is likely to be implemented across current and future Oracle SPARC Architecture processor implementations. This information provides guidance that should be particularly helpful to programmers who write in assembly language or those who write tools that generate SPARC instructions. It also provides the infrastructure for defining clear procedures for adding and removing features from the architecture over time, with minimal software disruption.

¹ although most features in this specification are already tagged with Software Classes, the full description of those Classes does not appear in this version of the specification. Please check back (<http://opensparc.sunsource.net/nonav/opensparct1.html>) for a later release of this document, which *will* include that description

Acknowledgements

This specification builds upon all previous SPARC specifications — SPARC V7, V8, and especially, SPARC V9. It therefore owes a debt to all the pioneers who developed those architectures.

SPARC V7 was developed by the SPARC (“Sunrise”) architecture team at Sun Microsystems, with special assistance from Professor David Patterson of University of California at Berkeley.

The enhancements present in SPARC V8 were developed by the nine member companies of the SPARC International Architecture Committee: Amdahl Corporation, Fujitsu Limited, ICL, LSI Logic, Matsushita, Philips International, Ross Technology, Sun Microsystems, and Texas Instruments.

SPARC V9 was also developed by the SPARC International Architecture Committee, with key contributions from the individuals named in the Editor’s Notes section of *The SPARC Architecture Manual-Version 9*.

The voluminous enhancements and additions present in this *Oracle SPARC Architecture 2015* specification are the result of **years** of deliberation, review, and feedback from readers of earlier revisions. I would particularly like to acknowledge the following people for their key contributions:

- The Oracle SPARC Architecture working group, who reviewed dozens of drafts of this specification and strived for the highest standards of accuracy and completeness; its active members included: Hendrik-Jan Agterkamp, Paul Caprioli, Steve Chessin, Hunter Donahue, Greg Grohoski, John (JJ) Johnson, Paul Jordan, Jim Laudon, Jim Lewis, Bob Maier, Wayne Mesard, Greg Onufer, Seongbae Park, Joel Storm, David Weaver, and Tom Webber.
- Robert (Bob) Maier, for his work on UltraSPARC Architecture 2005, including expansion of exception descriptions in every page of the Instructions chapter, major re-writes of several chapters and appendices (including *Memory*, *Memory Management*, *Performance Instrumentation*, and *Interrupt Handling*), significant updates to 5 other chapters, and tireless efforts to infuse commonality wherever possible across implementations.
- Steve Chessin and Joel Storm, “ace” reviewers — the two of them spotted more typographical errors and small inconsistencies than all other reviewers combined
- Jim Laudon (an UltraSPARC T1 architect and author of that processor’s implementation specification), for numerous descriptions of new features which were merged into the UltraSPARC Architecture 2005 specification
- The working group responsible for developing the system of Software Classes and Implementation Classes, comprising: Steve Chessin, Yuan Chou, Peter Damron, Q. Jacobson, Nicolai Kosche, Bob Maier, Ashley Saulsbury, Lawrence Spracklen, and David Weaver.
- Lawrence Spracklen, for his advice and numerous contributions regarding descriptions of VIS instructions
- Tom Webber, for providing descriptions of several new features in Oracle SPARC Architecture 2015
- Al Martin, for providing meticulously detailed floating-point exception tables, which have been integrated into Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

I hope you find the *Oracle SPARC Architecture 2015* specification more complete, accurate, and readable than its predecessors.

— *David Weaver*
Oracle SPARC Architecture Sr. Principal Engineer and specification editor

Corrections and other comments regarding this specification can be emailed to:
David.Weaver@Oracle.com

Document Overview

This chapter discusses:

- **Navigating Oracle SPARC Architecture 2015** on page 1.
- **Fonts and Notational Conventions** on page 2.
- **Reporting Errors in this Specification** on page 4.

1.1 Navigating *Oracle SPARC Architecture 2015*

If you are new to the SPARC architecture, read Chapter 3, *Architecture Overview*, study the definitions in Chapter 2, *Definitions*, then look into the subsequent sections and appendixes for more details in areas of interest to you.

If you are familiar with the SPARC V9 architecture but not Oracle SPARC Architecture 2015, note that Oracle SPARC Architecture 2015 conforms to the SPARC V9 Level 1 architecture (and most of Level 2), with numerous extensions — particularly with respect to VIS instructions.

This specification is structured as follows:

- Chapter 2, *Definitions*, which defines key terms used throughout the specification
- Chapter 3, *Architecture Overview*, provides an overview of Oracle SPARC Architecture 2015
- Chapter 4, *Data Formats*, describes the supported data formats
- Chapter 5, *Registers*, describes the register set
- Chapter 6, *Instruction Set Overview*, provides a high-level description of the Oracle SPARC Architecture 2015 instruction set
- Chapter 7, *Instructions*, describes the Oracle SPARC Architecture 2015 instruction set in great detail
- Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*, describes the trap model
- Chapter 9, *Memory* describes the supported memory model
- Chapter 10, *Address Space Identifiers (ASIs)*, provides a complete list of supported ASIs
- Chapter 11, *Performance Instrumentation* describes the architecture for performance monitoring hardware
- Chapter 12, *Traps*, describes the trap model
- Chapter 13, *Interrupt Handling*, describes how interrupts are handled
- Chapter 14, *Memory Management*, describes MMU operation
- Appendix A, *Opcodes Maps*, provides the overall picture of how the instruction set is mapped into opcodes
- Appendix B, *Implementation Dependencies*, describes all implementation dependencies

- Appendix C, *Assembly Language Syntax*, describes extensions to the SPARC assembly language syntax; in particular, synthetic instructions are documented in this appendix

1.2 Fonts and Notational Conventions

Fonts are used as follows:

- *Italic* font is used for emphasis, book titles, and the first instance of a word that is defined.
- *Italic* font is also used for terms where substitution is expected, for example, “*eccn*”, “virtual processor *n*”, or “*reg_plus_imm*”.
- *Italic sans serif* font is used for exception and trap names. For example, “The *privileged_action* exception...”
- lowercase helvetica font is used for register field names (named bits) and instruction field names, for example: “The *rs1* field contains...”
- UPPERCASE HELVETICA font is used for register names; for example, *FSR*.
- TYPEWRITER (Courier) font is used for literal values, such as code (assembly language, C language, ASI names) and for state names. For example: `%f0`, `ASI_PRIMARY`, `execute_state`.
- When a register field is shown along with its containing register name, they are separated by a period (‘.’), for example, “*FSR.cexc*”.
- UPPERCASE words are acronyms or instruction names. Some common acronyms appear in the glossary in Chapter 2, *Definitions*. **Note:** Names of some instructions contain both upper- and lower-case letters.
- An underscore character joins words in register, register field, exception, and trap names. **Note:** Such words may be split across lines at the underbar without an intervening hyphen. For example: “This is true whenever the *integer_condition_code* field...”

The following notational conventions are used:

- The left arrow symbol (←) is the assignment operator. For example, “`PC ← PC + 1`” means that the Program Counter (PC) is incremented by 1.
- Square brackets ([]) are used in two different ways, distinguishable by the context in which they are used:
 - Square brackets indicate indexing into an array. For example, `TT[TL]` means the element of the Trap Type (TT) array, as indexed by the contents of the Trap Level (TL) register.
 - Square brackets are also used to indicate optional additions/extensions to symbol names. For example, “`ST[D|Q]F`” expands to all three of “`STF`”, “`STDF`”, and “`STQF`”. Similarly, `ASI_PRIMARY[_LITTLE]` indicates two related address space identifiers, `ASI_PRIMARY` and `ASI_PRIMARY_LITTLE`. (Contrast with the use of angle brackets, below)
- Angle brackets (< >) indicate mandatory additions/extensions to symbol names. For example, “`ST<D|Q>F`” expands to mean “`STDF`” and “`STQF`”. (Contrast with the second use of square brackets, above)
- Curly braces ({ }) indicate a bit field within a register or instruction. For example, `CCR{4}` refers to bit 4 in the Condition Code Register.
- A consecutive set of values is indicated by specifying the upper and lower limit of the set separated by a colon (:), for example, `CCR{3:0}` refers to the set of four least significant bits of register CCR. (Contrast with the use of double periods, below)

- A double period (..) indicates any *single* intermediate value between two given inclusive end values is possible. For example, NAME[2..0] indicates four forms of NAME exist: NAME, NAME2, NAME1, and NAME0; whereas NAME<2..0> indicates that three forms exist: NAME2, NAME1, and NAME0. (Contrast with the use of the colon, above)
- A vertical bar (|) separates mutually exclusive alternatives inside square brackets ([]), angle brackets (< >), or curly braces ({ }). For example, “NAME[A | B]” expands to “NAME, NAMEA, NAMEB” and “NAME<A | B>” expands to “NAMEA, NAMEB”.
- An asterisk (*) is used as a wild card, encompassing the full set of valid values. For example, FCMP* refers to FCMP with all valid suffixes (in this case, FCMP<s | d | q> and FCMPE<s | d | q>). An asterisk is typically used when the full list of valid values either is not worth listing (because it has little or no relevance in the given context) or the valid values are too numerous to list in the available space.
- A slash (/) is used to separate paired or complementary values in a list, for example, “the LDBLOCKF/STBLOCKF instruction pair”
- The double colon (::) is an operator that indicates concatenation (typically, of bit vectors). Concatenation strictly strings the specified component values into a single longer string, in the order specified. The concatenation operator performs no arithmetic operation on any of the component values.

1.2.1 Implementation Dependencies

Implementors of Oracle SPARC Architecture 2015 processors are allowed to resolve some aspects of the architecture in machine-dependent ways.

The *definition* of each implementation dependency is indicated by the notation “**IMPL. DEP. #nn-XX:** Some descriptive text”. The number *nn* provides an index into the complete list of dependencies in Appendix B, *Implementation Dependencies*.

A *reference* to (but not definition of) an implementation dependency is indicated by the notation “(impl. dep. #nn)”.

1.2.2 Notation for Numbers

Numbers throughout this specification are decimal (base-10) unless otherwise indicated. Numbers in other bases are followed by a numeric subscript indicating their base (for example, 1001₂, FFFF 0000₁₆). Long binary and hexadecimal numbers within the text have spaces inserted every four characters to improve readability. Within C language or assembly language examples, numbers may be preceded by “0x” to indicate base-16 (hexadecimal) notation (for example, 0xFFFF0000).

1.2.3 Informational Notes

This guide provides several different types of information in notes, as follows:

Note	General notes contain incidental information relevant to the paragraph preceding the note.
Programming Note	Programming notes contain incidental information about how software can use an architectural feature.
Implementation Note	An Implementation Note contains incidental information, describing how an Oracle SPARC Architecture 2015 processor might implement an architectural feature.

V9 Compatibility Note	Note containing information about possible differences between Oracle SPARC Architecture 2015 and SPARC V9 implementations. Such information is relevant to Oracle SPARC Architecture 2015 implementations and might not apply to other SPARC V9 implementations.
Forward Compatibility Note	Note containing information about how the Oracle SPARC Architecture is expected to evolve in the future. Such notes are not intended as a guarantee that the architecture will evolve as indicated, but as a guide to features that should not be depended upon to remain the same, by software intended to run on both current and future implementations.

1.3 Reporting Errors in this Specification

This specification has been reviewed for completeness and accuracy. Nonetheless, as with any document this size, errors and omissions may occur, and reports of such are welcome. Please send "bug reports" and other comments on this document to email address: David.Weaver@Oracle.com

Definitions

This chapter defines concepts and terminology common to all implementations of Oracle SPARC Architecture 2015.

- address space** A range of 2^{64} locations that can be addressed by instruction fetches and load, store, or load-store instructions. See also **address space identifier (ASI)**.
- address space identifier (ASI)** An 8-bit value that identifies a particular address space. An ASI is (implicitly or explicitly) associated with every instruction access or data access. See also **implicit ASI**.
- aliased** Said of each of two virtual or real addresses that refer to the same underlying memory location.
- application program** A program executed with the virtual processor in nonprivileged mode. **Note:** Statements made in this specification regarding application programs may not be applicable to programs (for example, debuggers) that have access to privileged virtual processor state (for example, as stored in a memory-image dump).
- ASI** Address space identifier.
- ASR** Ancillary State register.
- big-endian** An addressing convention. Within a multiple-byte integer, the byte with the smallest address is the most significant; a byte's significance decreases as its address increases.
- byte** Eight consecutive bits of data, aligned on an 8-bit boundary.
- CCR** Abbreviation for Condition Codes Register.
- clean window** A register window in which each of the registers contain 0, a valid address from the current address space, or valid data from the current address space.
- coherence** A set of protocols guaranteeing that all memory accesses are globally visible to all caches on a shared-memory bus.
- completed (memory operation)** Said of a memory transaction when an idealized memory has executed the transaction with respect to all processors. A load is considered completed when no subsequent memory transaction can affect the value returned by the load. A store is considered completed when no subsequent load can return the value that was overwritten by the store.
- context** A set of translations that defines a particular address space. See also **Memory Management Unit (MMU)**.
- context ID** A numeric value that uniquely identifies a particular context.
- copyback** The process of sending a copy of the data from a cache line owned by a physical processor core, in response to a snoop request from another device.
- CPI** Cycles per instruction. The number of clock cycles it takes to execute an instruction.
- cross-call** An interprocessor call in a system containing multiple virtual processors.

CTI	Abbreviation for control-transfer instruction .
current window	The block of 24 R registers that is presently in use. The Current Window Pointer (CWP) register points to the current window.
cycle	The atomic unit of time in a physical implementation of a processor core. The duration of a cycle is its period, and the inverse of the period is the physical processor core's operating frequency (typically measured in gigaHertz, in contemporary implementations). The physical processor core divides the work of managing instructions and data and executing instructions into multiple cycles. This division of processing steps into cycles is implementation-dependent. The operating frequency is implementation-dependent and potentially varying in time for a given implementation.
data access (instruction)	A load, store, load-store, or FLUSH instruction.
DCTI	Delayed control transfer instruction.
denormalized number	Synonym for subnormal number .
deprecated	The term applied to an architectural feature (such as an instruction or register) for which an Oracle SPARC Architecture implementation provides support <i>only</i> for compatibility with previous versions of the architecture. Use of a deprecated feature must generate correct results but may compromise software performance. Deprecated features should not be used in new Oracle SPARC Architecture software and may not be supported in future versions of the architecture.
doubleword	An 8-byte datum. Note: The definition of this term is architecture dependent and may differ from that used in other processor architectures.
even parity	The mode of parity checking in which each combination of data bits plus a parity bit contains an even number of '1' bits.
exception	A condition that makes it impossible for the processor to continue executing the current instruction stream. Some exceptions may be masked (that is, trap generation disabled — for example, floating-point exceptions masked by FSR.tem) so that the decision on whether or not to apply special processing can be deferred and made by software at a later time. See also trap .
explicit ASI	An ASI that that is provided by a load, store, or load-store alternate instruction (either from its imm_asi field or from the ASI register).
extended word	An 8-byte datum, nominally containing integer data. Note: The definition of this term is architecture dependent and may differ from that used in other processor architectures.
fccn	One of the floating-point condition code fields fcc0, fcc1, fcc2, or fcc3.
FGU	Floating-point and Graphics Unit (which most implementations specify as a superset of FPU).
floating-point exception	An exception that occurs during the execution of a floating-point operate (FPop) instruction. The exceptions <i>fp_exception_ieee_754</i> and <i>fp_exception_other</i> .
F register	A floating-point register. The SPARC V9 architecture includes single-, double-, and quad-precision F registers.
floating-point operate instructions	Instructions that perform floating-point calculations, as defined in <i>Floating-Point Operate (FPop) Instructions</i> on page 92. FPop instructions do not include FBfcc instructions, loads and stores between memory and the F registers, or non-floating-point operations that read or write F registers.
floating-point trap type	The specific type of a floating-point exception, encoded in the FSR.ftt field.

floating-point unit	A processing unit that contains the floating-point registers and performs floating-point operations, as defined by this specification.
FPop	Abbreviation for floating-point operate (instructions).
FPRS	Floating-Point Register State register.
FPU	Floating-Point Unit.
FSR	Floating-Point Status register.
GL	Global Level register.
GSR	General Status register.
halfword	A 2-byte datum. Note: The definition of this term is architecture dependent and may differ from that used in other processor architectures.
hyperprivileged	An adjective that describes: (1) the state of the processor when the processor is in hyperprivileged mode; (2) processor state that is only accessible to software while the processor is in hyperprivileged mode
IEEE 754	IEEE Standard 754-1985, the IEEE Standard for Binary Floating-Point Arithmetic.
IEEE-754 exception	A floating-point exception, as specified by IEEE Std 754-1985. Listed within this specification as IEEE_754_exception.
implementation	Hardware or software that conforms to all of the specifications of an instruction set architecture (ISA).
implementation dependent	An aspect of the Oracle SPARC Architecture that can legitimately vary among implementations. In many cases, the permitted range of variation is specified. When a range is specified, compliant implementations must not deviate from that range.
implicit ASI	An address space identifier that is implicitly supplied by the virtual processor on all instruction accesses and on data accesses that do not explicitly provide an ASI value (from either an imm_asi instruction field or the ASI register).
initiated	Synonym for issued .
instruction field	A bit field within an instruction word.
instruction group	One or more independent instructions that can be dispatched for simultaneous execution.
instruction set architecture	A set that defines instructions, registers, instruction and data memory, the effect of executed instructions on the registers and memory, and an algorithm for controlling instruction execution. Does not define clock cycle times, cycles per instruction, data paths, etc. This specification defines the Oracle SPARC Architecture 2015 instruction set architecture.
integer unit	A processing unit that performs integer and control-flow operations and contains general-purpose integer registers and virtual processor state registers, as defined by this specification.
interrupt request	A request for service presented to a virtual processor by an external device.
inter-strand	Describes an operation that crosses virtual processor (strand) boundaries.
intra-strand	Describes an operation that occurs entirely within one virtual processor (strand).
invalid (ASI or address)	Undefined, reserved, or illegal.
ISA	Instruction set architecture.

issued A memory transaction (load, store, or atomic load-store) is said to be “issued” when a virtual processor has sent the transaction to the memory subsystem and the completion of the request is out of the virtual processor’s control. Synonym for **initiated**.

IU Integer Unit.

little-endian An addressing convention. Within a multiple-byte integer, the byte with the smallest address is the least significant; a byte’s significance increases as its address increases.

load An instruction that reads (but does not write) memory or reads (but does not write) location(s) in an alternate address space. Some examples of *Load* includes loads into integer or floating-point registers, block loads, and alternate address space variants of those instructions. See also **load-store** and **store**, the definitions of which are mutually exclusive with *load*.

load-store An instruction that explicitly both reads and writes memory or explicitly reads and writes location(s) in an alternate address space. *Load-store* includes instructions such as *CASA*, *CASXA*, *LDSTUB*, and the deprecated *SWAP* instruction. See also **load** and **store**, the definitions of which are mutually exclusive with *load-store*.

may A keyword indicating flexibility of choice with no implied preference. **Note:** “may” indicates that an action or operation is allowed; “can” indicates that it is possible.

Memory Management

Unit (MMU) The address translation hardware in an Oracle SPARC Architecture implementation that translates 64-bit virtual address into underlying hardware addresses. The MMU is composed of the ASRs and ASI registers used to manage address translation. See also **context real address**, and **virtual address**.

memory version

mismatch See *version mismatch*.

microcore Synonym for **pipeline**.

MMU Abbreviation for **Memory Management Unit**.

multiprocessor system A system containing more than one processor.

must A keyword indicating a mandatory requirement. Designers must implement all such mandatory requirements to ensure interoperability with other Oracle SPARC Architecture-compliant products. Synonym for **shall**.

next program counter Conceptually, a register that contains the address of the instruction to be executed next if a trap does not occur.

NFO Nonfault access only.

nonfaulting load A load operation that behaves identically to a normal load operation, except when supplied an invalid effective address by software. In that case, a regular load triggers an exception whereas a nonfaulting load appears to ignore the exception and loads its destination register with a value of zero (on an Oracle SPARC Architecture processor, hardware treats regular and nonfaulting loads identically; the distinction is made in trap handler software). Contrast with **speculative load**.

nonprivileged An adjective that describes

- (1) the state of the virtual processor when `PSTATE.priv = 0`, that is, when it is in nonprivileged mode;
- (2) virtual processor state information that is accessible to software regardless of the current privilege mode; for example, nonprivileged registers, nonprivileged ASRs, or, in general, nonprivileged state;
- (3) an instruction that can be executed in any privilege mode (privileged or nonprivileged).

nonprivileged mode The mode in which a virtual processor is operating when executing application software (at the lowest privilege level). Nonprivileged mode is defined by `PSTATE.priv = 0`. See also **privileged** and **hyperprivileged**.

nontranslating ASI	An ASI that does not refer to memory (for example, refers to control/status register(s)) and for which the MMU does not perform address translation.
NPC	Next program counter.
nucleus software	Privileged software running at a trap level greater than 0 (TL > 0).
NUMA	Nonuniform memory access.
N_REG_WINDOWS	The number of register windows present in a particular implementation.
octlet	Eight bytes (64 bits) of data. Not to be confused with “octet,” which has been commonly used to describe eight bits of data. In this document, the term <i>byte</i> , rather than octet, is used to describe eight bits of data.
odd parity	The mode of parity checking in which each combination of data bits plus a parity bit together contain an odd number of ‘1’ bits.
opcode	A bit pattern that identifies a particular instruction.
optional	A feature not required for Oracle SPARC Architecture 2015 compliance.
PC	Program counter.
physical processor	<i>Synonym for processor</i> ; used when an explicit contrast needs to be drawn between processor and virtual processor. See also processor and virtual processor .
PIL	Processor Interrupt Level register.
pipeline	Refers to an execution pipeline, the basic collection of hardware needed to execute instructions. <i>Synonym for microcore</i> . See also processor , strand , thread , and virtual processor .
prefetchable	(1) An attribute of a memory location that indicates to an MMU that PREFETCH operations to that location may be applied. (2) A memory location condition for which the system designer has determined that no undesirable effects will occur if a PREFETCH operation to that location is allowed to succeed. Typically, normal memory is prefetchable. Nonprefetchable locations include those that, when read, change state or cause external events to occur. For example, some I/O devices are designed with registers that clear on read; others have registers that initiate operations when read. See also side effect .
privileged	An adjective that describes: (1) the state of the virtual processor when PSTATE.priv = 1, that is, when the virtual processor is in privileged mode; (2) processor state that is only accessible to software while the virtual processor is in privileged mode; for example, privileged registers, privileged ASRs, or, in general, privileged state; (3) an instruction that can be executed only when the virtual processor is in privileged mode.
privileged mode	The mode in which a processor is operating when PSTATE.priv = 1. See also nonprivileged and hyperprivileged .
processor	The unit on which a shared interface is provided to control the configuration and execution of a collection of strands; a physical module that plugs into a system. <i>Synonym for processor module</i> . See also pipeline , strand , thread , and virtual processor .
processor core	<i>Synonym for physical core</i> .
processor module	<i>Synonym for processor</i> .
program counter	A register that contains the address of the instruction currently being executed.
quadword	A 16-byte datum. Note: The definition of this term is architecture dependent and may be different from that used in other processor architectures.

- R register** An integer register. Also called a general-purpose register or working register.
- RA** Real address.
- RAS** Reliability, Availability, and Serviceability
- RAW** Read After Write (hazard)
- rd** Rounding direction.
- real address** An address produced by a virtual processor that refers to a particular software-visible memory location, as viewed from privileged mode. Virtual addresses are usually translated by a combination of hardware and software to real addresses, which can be used to access real memory. See also **virtual address**.
- reserved** Describing an instruction field, certain bit combinations within an instruction field, or a register field that is reserved for definition by future versions of the architecture.
- A reserved instruction field must read as 0, unless the implementation supports extended instructions within the field. The behavior of an Oracle SPARC Architecture 2015 virtual processor when it encounters a nonzero value in a reserved instruction field is as defined in *Reserved Opcodes and Instruction Fields* on page 93.*
- A reserved bit combination within an instruction field is defined in Chapter 7, *Instructions*. In all cases, an Oracle SPARC Architecture 2015 processor must decode and trap on such reserved bit combinations.*
- A reserved field within a register reads as 0 in current implementations and, when written by software, should always be written with values of that field previously read from that register or with the value zero (as described in *Reserved Register Fields* on page 32).*
- Throughout this specification, figures and tables illustrating registers and instruction encodings indicate reserved fields and reserved bit combinations with a wide (“em”) dash (—).
- restricted** Describes an address space identifier (ASI) that may be accessed only while the virtual processor is operating in privileged mode.
- retired** An instruction is said to be “retired” when one of the following two events has occurred:
- (1) A precise trap has been taken, with TPC containing the instruction's address (the instruction has not changed architectural state in this case).
 - (2) The instruction's execution has progressed to a point at which architectural state affected by the instruction has been updated such that all three of the following are true:
 - The PC has advanced beyond the instruction.
 - Except for deferred trap handlers, no consumer in the same instruction stream can see the old values and all consumers in the same instruction stream will see the new values.
 - Stores are visible to all loads in the same instruction stream, including stores to noncacheable locations.
- RMO** Abbreviation for Relaxed Memory Order (a memory model).
- RTO** Read to Own (a type of transaction, used to request ownership of a cache line).
- RTS** Read to Share (a type of transaction, used to request read-only access to a cache line).
- shall** Synonym for **must**.
- should** A keyword indicating flexibility of choice with a strongly preferred implementation. Synonym for **it is recommended**.
- side effect** The result of a memory location having additional actions beyond the reading or writing of data. A side effect can occur when a memory operation on that location is allowed to succeed. Locations with side effects include those that, when accessed, change state or cause external events to occur. For example, some I/O devices contain registers that clear on read; others have registers that initiate operations when read. See also **prefetchable**.

- SIMD** Single Instruction/Multiple Data; a class of instructions that perform identical operations on multiple data contained (or “packed”) in each source operand.
- speculative load** A load operation that is issued by a virtual processor speculatively, that is, before it is known whether the load will be executed in the flow of the program. Speculative accesses are used by hardware to speed program execution and are transparent to code. An implementation, through a combination of hardware and system software, must nullify speculative loads on memory locations that have side effects; otherwise, such accesses produce unpredictable results. Contrast with **nonfaulting load**.
- store** An instruction that writes (but does not explicitly read) memory or writes (but does not explicitly read) location(s) in an alternate address space. Some examples of *Store* includes stores from either integer or floating-point registers, block stores, Partial Store, and alternate address space variants of those instructions. See also **load** and **load-store**, the definitions of which are mutually exclusive with *store*.
- strand** The hardware state that must be maintained in order to execute a software thread. See also **pipeline**, **processor**, **thread**, and **virtual processor**.
- subnormal number** A nonzero floating-point number, the exponent of which has a value of zero. A more complete definition is provided in IEEE Standard 754-1985.
- superscalar** An implementation that allows several instructions to be issued, executed, and committed in one clock cycle.
- supervisor software** Software that executes when the virtual processor is in privileged mode.
- synchronization** An operation that causes the processor to wait until the effects of all previous instructions are completely visible before any subsequent instructions are executed.
- system** A set of virtual processors that share a common hardware memory address space.
- taken** A control-transfer instruction (CTI) is *taken* when the CTI writes the target address value into NPC.
A trap is *taken* when the control flow changes in response to an exception, reset, Tcc instruction, or interrupt. An exception must be detected and recognized before it can cause a trap to be taken.
- TBA** Trap base address.
- thread** A software entity that can be executed on hardware. See also **pipeline**, **processor**, **strand**, and **virtual processor**.
- TNPC** Trap-saved next program counter.
- TPC** Trap-saved program counter.
- trap** The action taken by a virtual processor when it changes the instruction flow in response to the presence of an exception, reset, a Tcc instruction, or an interrupt. The action is a vectored transfer of control to more-privileged software through a table, the address of which is specified by the privileged Trap Base Address (TBA) register. See also **exception**.
- TSB** Translation storage buffer. A table of the address translations that is maintained by software in system memory and that serves as a cache of virtual-to-real address mappings.
- TSO** Total Store Order (a memory model).
- TTE** Translation Table Entry. Describes the virtual-to-real translation and page attributes for a specific page in the page table. In some cases, this term is explicitly used to refer to entries in the TSB.
- UA-2015** Oracle SPARC Architecture 2015
- unassigned** A value (for example, an ASI number), the semantics of which are not architecturally mandated and which may be determined independently by each implementation within any guidelines given.

- undefined** An aspect of the architecture that has deliberately been left unspecified. Software should have no expectation of, nor make any assumptions about, an undefined feature or behavior. Use of such a feature can deliver unexpected results and may or may not cause a trap. An undefined feature may vary among implementations, and may also vary over time on a given implementation.
- Notwithstanding any of the above, undefined aspects of the architecture shall not cause security holes (such as changing the privilege state or allowing circumvention of normal restrictions imposed by the privilege state), put a virtual processor into a more-privileged mode, or put the virtual processor into an unrecoverable state.
- unimplemented** An architectural feature that is not directly executed in hardware because it is optional or is emulated in software.
- unpredictable** Synonym for **undefined**.
- uniprocessor system** A system containing a single virtual processor.
- unrestricted** Describes an address space identifier (ASI) that can be used in all privileged modes; that is, regardless of the value of PSTATE.priv.
- user application program** Synonym for **application program**.
- VA** Abbreviation for **virtual address**.
- virtual address** An address produced by a virtual processor that refers to a particular software-visible memory location. Virtual addresses usually are translated by a combination of hardware and software to real addresses, which can be used to access real memory. See also **real address**.
- virtual core, virtual processor core** Synonyms for **virtual processor**.
- virtual processor** The term *virtual processor*, or *virtual processor core*, is used to identify each strand in a processor. At any given time, an operating system can have a different thread scheduled on each virtual processor. See also **pipeline**, **processor**, **strand**, and **thread**.
- VIS** Abbreviation for VIS™ Instruction Set.
- VP** Abbreviation for **virtual processor**.
- word** A 4-byte datum. **Note:** The definition of this term is architecture dependent and may differ from that used in other processor architectures.

Architecture Overview

The Oracle SPARC Architecture supports 32-bit and 64-bit integer and 32-bit, 64-bit, and 128-bit floating-point as its principal data types. The 32-bit and 64-bit floating-point types conform to IEEE Std 754-1985. The 128-bit floating-point type conforms to IEEE Std 1596.5-1992. The architecture defines general-purpose integer, floating-point, and special state/status register instructions, all encoded in 32-bit-wide instruction formats. The load/store instructions address a linear, 2^{64} -byte virtual address space.

The *Oracle SPARC Architecture 2015* specification describes a processor architecture to which Sun Microsystem's SPARC processor implementations (beginning with UltraSPARC T1) comply. Future implementations are expected to comply with either this document or a later revision of this document.

The Oracle SPARC Architecture 2015 is a descendant of the SPARC V9 architecture and complies fully with the "Level 1" (nonprivileged) SPARC V9 specification.

Nonprivileged (application) software that is intended to be portable across all SPARC V9 processors should be written to adhere to *The SPARC Architecture Manual-Version 9*.

Material in this document specific to Oracle SPARC Architecture 2015 processors may not apply to SPARC V9 processors produced by other vendors.

In this specification, the word *architecture* refers to the processor features that are visible to an assembly language programmer or to a compiler code generator. It does not include details of the implementation that are not visible or easily observable by software, nor those that only affect timing (performance).

3.1 The Oracle SPARC Architecture 2015

This section briefly describes features, attributes, and components of the Oracle SPARC Architecture 2015 and, further, describes correct implementation of the architecture specification and SPARC V9-compliance levels.

3.1.1 Features

The Oracle SPARC Architecture 2015, like its ancestor SPARC V9, includes the following principal features:

- **A linear 64-bit address space** with 64-bit addressing.
- **32-bit wide instructions** — These are aligned on 32-bit boundaries in memory. Only load and store instructions access memory and perform I/O.

- **Few addressing modes** — A memory address is given as either “register + register” or “register + immediate”.
- **Triadic register addresses** — Most computational instructions operate on two register operands or one register and a constant and place the result in a third register.
- **A large windowed register file** — At any one instant, a program sees 8 global integer registers plus a 24-register window of a larger register file. The windowed registers can be used as a cache of procedure arguments, local values, and return addresses.
- **Floating point** — The architecture provides an IEEE 754-compatible floating-point instruction set, operating on a separate register file that provides 32 single-precision (32-bit), 32 double-precision (64-bit), and 16 quad-precision (128-bit) overlaid registers.
- **Fast trap handlers** — Traps are vectored through a table.
- **Multiprocessor synchronization instructions** — Multiple variations of atomic load-store memory operations are supported.
- **Predicted branches** — The branch with prediction instructions allows the compiler or assembly language programmer to give the hardware a hint about whether a branch will be taken.
- **Branch elimination instructions** — Several instructions can be used to eliminate branches altogether (for example, Move on Condition). Eliminating branches increases performance in superscalar and superpipelined implementations.
- **Hardware trap stack** — A hardware trap stack is provided to allow nested traps. It contains all of the machine state necessary to return to the previous trap level. The trap stack makes the handling of faults and error conditions simpler, faster, and safer.

In addition, Oracle SPARC Architecture 2015 includes the following features that were not present in the SPARC V9 specification:

- **Hyperprivileged mode**, which simplifies porting of operating systems, supports far greater portability of operating system (privileged) software, and supports the ability to run multiple simultaneous guest operating systems. (hyperprivileged mode is described in detail in the Hyperprivileged version of this specification)
- **Multiple levels of global registers** — Instead of the two 8-register sets of global registers specified in the SPARC V9 architecture, Oracle SPARC Architecture 2015 provides multiple sets; typically, one set is used at each trap level.
- **Extended instruction set** — Oracle SPARC Architecture 2015 provides many instruction set extensions, including the VIS instruction set for “vector” (SIMD) data operations.
- **More detailed, specific instruction descriptions** — Oracle SPARC Architecture 2015 provides many more details regarding what exceptions can be generated by each instruction and the specific conditions under which those exceptions can occur. Also, detailed lists of valid ASIs are provided for each load/store instruction from/to alternate space.
- **Detailed MMU architecture** — Oracle SPARC Architecture 2015 provides a blueprint for the software view of the UltraSPARC MMU (TTEs and TSBs).

3.1.2 Attributes

Oracle SPARC Architecture 2015 is a processor *instruction set architecture* (ISA) derived from SPARC V8 and SPARC V9, which in turn come from a reduced instruction set computer (RISC) lineage. As an architecture, Oracle SPARC Architecture 2015 allows for a spectrum of processor and system *implementations* at a variety of price/performance points for a range of applications, including scientific/engineering, programming, real-time, and commercial applications.

3.1.2.1 Design Goals

The Oracle SPARC Architecture 2015 architecture is designed to be a target for optimizing compilers and high-performance hardware implementations. This specification documents the Oracle SPARC Architecture 2015 and provides a design spec against which an implementation can be verified, using appropriate verification software.

3.1.2.2 Register Windows

The Oracle SPARC Architecture 2015 architecture is derived from the SPARC architecture, which was formulated at Sun Microsystems in 1984 through 1987. The SPARC architecture is, in turn, based on the RISC I and II designs engineered at the University of California at Berkeley from 1980 through 1982. The SPARC “register window” architecture, pioneered in the UC Berkeley designs, allows for straightforward, high-performance compilers and a reduction in memory load/store instructions.

Note that privileged software, not user programs, manages the register windows. Privileged software can save a minimum number of registers (approximately 24) during a context switch, thereby optimizing context-switch latency.

3.1.3 System Components

The Oracle SPARC Architecture 2015 allows for a spectrum of subarchitectures, such as cache system.

3.1.3.1 Binary Compatibility

The most important mandate for the Oracle SPARC Architecture is compatibility across implementations of the architecture for application (nonprivileged) software, down to the binary level. Binaries executed in nonprivileged mode should behave identically on all Oracle SPARC Architecture systems when those systems are running an operating system known to provide a standard execution environment. One example of such a standard environment is the SPARC V9 Application Binary Interface (ABI).

Although different Oracle SPARC Architecture 2015 systems can execute nonprivileged programs at different rates, they will generate the same results as long as they are run under the same memory model. See Chapter 9, *Memory*, for more information.

Additionally, Oracle SPARC Architecture 2015 is binary upward-compatible from SPARC V9 for applications running in nonprivileged mode that conform to the SPARC V9 ABI and upward-compatible from SPARC V8 for applications running in nonprivileged mode that conform to the SPARC V8 ABI.

3.1.3.2 Oracle SPARC Architecture 2015 MMU

Although the SPARC V9 architecture allows its implementations freedom in their MMU designs, Oracle SPARC Architecture 2015 defines a common MMU architecture (see Chapter 14, *Memory Management*) with some specifics left to implementations (see processor implementation documents).

3.1.3.3 Privileged Software

Oracle SPARC Architecture 2015 does not assume that all implementations must execute identical privileged software (operating systems). Thus, certain traits that are visible to privileged software may be tailored to the requirements of the system.

3.1.4 Architectural Definition

The Oracle SPARC Architecture 2015 is defined by the chapters and appendixes of this specification. A correct implementation of the architecture interprets a program strictly according to the rules and algorithms specified in the chapters and appendixes.

Oracle SPARC Architecture 2015 defines a set of implementations that conform to the SPARC V9 architecture, Level 1.

3.1.5 Oracle SPARC Architecture 2015 Compliance with SPARC V9 Architecture

Oracle SPARC Architecture 2015 fully complies with SPARC V9 Level 1 (nonprivileged). It partially complies with SPARC V9 Level 2 (privileged).

3.1.6 Implementation Compliance with Oracle SPARC Architecture 2015

Compliant implementations must not add to or deviate from this standard except in aspects described as implementation dependent. Appendix B, *Implementation Dependencies* lists all Oracle SPARC Architecture 2015, SPARC V9, and SPARC V8 implementation dependencies. Documents for specific Oracle SPARC Architecture 2015 processor implementations describe the manner in which implementation dependencies have been resolved in those implementations.

IMPL. DEP. #1-V8: Whether an instruction complies with Oracle SPARC Architecture 2015 by being implemented directly by hardware, simulated by software, or emulated by firmware is implementation dependent.

3.2 Processor Architecture

An Oracle SPARC Architecture processor logically consists of an integer unit (IU) and a floating-point unit (FPU), each with its own registers. This organization allows for implementations with concurrent integer and floating-point instruction execution. Integer registers are 64 bits wide; floating-point registers are 32, 64, or 128 bits wide. Instruction operands are single registers, register pairs, register quadruples, or immediate constants.

An Oracle SPARC Architecture virtual processor can run in *nonprivileged mode*, *privileged mode*, or in mode(s) of greater privilege. In privileged mode, the processor can execute nonprivileged and privileged instructions. In nonprivileged mode, the processor can only execute nonprivileged instructions. In nonprivileged or privileged mode, an attempt to execute an instruction requiring greater privilege than the current mode causes a trap.

3.2.1 Integer Unit (IU)

An Oracle SPARC Architecture 2015 implementation's integer unit contains the general-purpose registers and controls the overall operation of the virtual processor. The IU executes the integer arithmetic instructions and computes memory addresses for loads and stores. It also maintains the program counters and controls instruction execution for the FPU.

IMPL. DEP. #2-V8: An Oracle SPARC Architecture implementation may contain from 72 to 640 general-purpose 64-bit R registers. This corresponds to a grouping of the registers into $MAXPGL + 1$ sets of global R registers plus a circular stack of $N_REG_WINDOWS$ sets of 16 registers each, known as register windows. The number of register windows present ($N_REG_WINDOWS$) is implementation dependent, within the range of 3 to 32 (inclusive).

3.2.2 Floating-Point Unit (FPU)

An Oracle SPARC Architecture 2015 implementation's FPU has thirty-two 32-bit (single-precision) floating-point registers, thirty-two 64-bit (double-precision) floating-point registers, and sixteen 128-bit (quad-precision) floating-point registers, some of which overlap.

If no FPU is present, then it appears to software as if the FPU is permanently disabled.

If the FPU is not enabled, then an attempt to execute a floating-point instruction generates an *fp_disabled* trap and the *fp_disabled* trap handler software must either

- Enable the FPU (if present) and reexecute the trapping instruction, or
- Emulate the trapping instruction in software.

3.3 Instructions

Instructions fall into the following basic categories:

- Memory access
- Integer arithmetic / logical / shift
- Control transfer
- State register access
- Floating-point operate
- Conditional move
- Register window management
- SIMD (single instruction, multiple data) instructions

These classes are discussed in the following subsections.

3.3.1 Memory Access

Load, store, load-store, and PREFETCH instructions are the only instructions that access memory. They use two R registers or an R register and a signed 13-bit immediate value to calculate a 64-bit, byte-aligned memory address. The Integer Unit appends an ASI to this address.

The destination field of the load/store instruction specifies either one or two R registers or one, two, or four F registers that supply the data for a store or that receive the data from a load.

Integer load and store instructions support byte, halfword (16-bit), word (32-bit), and extended-word (64-bit) accesses. There are versions of integer load instructions that perform either sign-extension or zero-extension on 8-bit, 16-bit, and 32-bit values as they are loaded into a 64-bit destination register. Floating-point load and store instructions support word, doubleword, and quadword¹ memory accesses.

¹ No Oracle SPARC Architecture processor currently implements the LDQF instruction in hardware; it generates an exception and is emulated in software running at a higher privilege level.

CASA, CASXA, and LDSTUB are special atomic memory access instructions that concurrent processes use for synchronization and memory updates.

Note | The SWAP instruction is also specified, but it is deprecated and should not be used in newly developed software.

The (nonportable) LDTXA instruction supplies an atomic 128-bit (16-byte) load that is important in certain system software applications.

3.3.1.1 Memory Alignment Restrictions

A memory access on an Oracle SPARC Architecture virtual processor must typically be aligned on an address boundary greater than or equal to the size of the datum being accessed. An improperly aligned address in a load, store, or load-store instruction may trigger an exception and cause a subsequent trap. For details, see *Memory Alignment Restrictions* on page 79.

3.3.1.2 Addressing Conventions

The Oracle SPARC Architecture uses big-endian byte order by default: the address of a quadword, doubleword, word, or halfword is the address of its most significant byte. Increasing the address means decreasing the significance of the unit being accessed. All instruction accesses are performed using big-endian byte order.

The Oracle SPARC Architecture also supports little-endian byte order for data accesses only: the address of a quadword, doubleword, word, or halfword is the address of its least significant byte. Increasing the address means increasing the significance of the data unit being accessed.

Addressing conventions are illustrated in FIGURE 6-2 on page 81 and FIGURE 6-3 on page 83.

3.3.1.3 Addressing Range

IMPL. DEP. #405-S10: An Oracle SPARC Architecture implementation may support a full 64-bit virtual address space or a more limited range of virtual addresses. In an implementation that does not support a full 64-bit virtual address space, the supported range of virtual addresses is restricted to two equal-sized ranges at the extreme upper and lower ends of 64-bit addresses; that is, for n -bit virtual addresses, the valid address ranges are 0 to $2^{n-1} - 1$ and $2^{64} - 2^{n-1}$ to $2^{64} - 1$.

3.3.1.4 Load/Store Alternate

Variants of load/store instructions, the *load/store alternate* instructions, can specify an arbitrary 8-bit address space identifier for the load/store data access.

Access to alternate spaces 00_{16} – $2F_{16}$ is restricted to privileged software, access to alternate spaces 30_{16} – $7F_{16}$ is restricted to hyperprivileged software, and access to alternate spaces 80_{16} – FF_{16} is unrestricted. Some of the ASIs are available for implementation-dependent uses. Privileged software can use the implementation-dependent ASIs to access special protected registers, such as cache control registers, virtual processor state registers, and other processor-dependent or system-dependent values. See *Address Space Identifiers (ASIs)* on page 83 for more information.

Alternate space addressing is also provided for the atomic memory access instructions LDSTUBA, CASA, and CASXA.

Note | The SWAPA instruction is also specified, but it is deprecated and should not be used in newly developed software.

3.3.1.5 Separate Instruction and Data Memories

The interpretation of addresses can be unified, in which case the same translations and caching are applied to both instructions and data. Alternatively, addresses can be “split”, in which case instruction references use one caching and translation mechanism and data references use another, although the same underlying main memory is shared.

In such split-memory systems, the coherency mechanism may be split, so a write¹ into data memory is not immediately reflected in instruction memory. For this reason, programs that modify their own instruction stream (self-modifying code²) and that wish to be portable across all Oracle SPARC Architecture (and SPARC V9) processors must issue FLUSH instructions, or a system call with a similar effect, to bring the instruction and data caches into a consistent state.

An Oracle SPARC Architecture virtual processor may or may not have coherent instruction and data caches. Even if an implementation does have coherent instruction and data caches, a FLUSH instruction is required for self-modifying code — not for cache coherency, but to flush pipeline instruction buffers that contain unmodified instructions which may have been subsequently modified.

3.3.1.6 Input/Output (I/O)

The Oracle SPARC Architecture assumes that input/output registers are accessed through load/store alternate instructions, normal load/store instructions, or read/write Ancillary State Register instructions (RDAsr, WRAsr).

IMPL. DEP. #123-V9: The semantic effect of accessing input/output (I/O) locations is implementation dependent.

IMPL. DEP. #6-V8: Whether the I/O registers can be accessed by nonprivileged code is implementation dependent.

IMPL. DEP. #7-V8: The addresses and contents of I/O registers are implementation dependent.

3.3.1.7 Memory Synchronization

Two instructions are used for synchronization of memory operations: FLUSH and MEMBAR. Their operation is explained in *Flush Instruction Memory* on page 171 and *Memory Barrier* on page 269, respectively.

Note | STBAR is also available, but it is deprecated and should not be used in newly developed software.

3.3.2 Integer Arithmetic / Logical / Shift Instructions

The arithmetic/logical/shift instructions perform arithmetic, tagged arithmetic, logical, and shift operations. With one exception, these instructions compute a result that is a function of two source operands; the result is either written into a destination register or discarded. The exception, SETHI, can be used in combination with other arithmetic and/or logical instructions to create a constant in an R register.

Shift instructions shift the contents of an R register left or right by a given number of bits (“shift count”). The shift distance is specified by a constant in the instruction or by the contents of an R register.

¹ this includes use of store instructions (executed on the same or another virtual processor) that write to instruction memory, or any other means of writing into instruction memory (for example, DMA)

² this is practiced, for example, by software such as debuggers and dynamic linkers

3.3.3 Control Transfer

Control-transfer instructions (CTIs) include PC-relative branches and calls, register-indirect jumps, and conditional traps. Most of the control-transfer instructions are delayed; that is, the instruction immediately following a control-transfer instruction in logical sequence is dispatched before the control transfer to the target address is completed. Note that the next instruction in logical sequence may not be the instruction following the control-transfer instruction in memory.

The instruction following a delayed control-transfer instruction is called a *delay* instruction. Setting the *annul bit* in a conditional delayed control-transfer instruction causes the delay instruction to be annulled (that is, to have no effect) if and only if the branch is not taken. Setting the annul bit in an *unconditional* delayed control-transfer instruction (“branch always”) causes the delay instruction to be always annulled.

Note | The SPARC V8 architecture specified that the delay instruction was always fetched, even if annulled, and that an annulled instruction could not cause any traps. The SPARC V9 architecture does not require the delay instruction to be fetched if it is annulled.

Branch and CALL instructions use PC-relative displacements. The jump and link (JMWL) and return (RETURN) instructions use a register-indirect target address. They compute their target addresses either as the sum of two R registers or as the sum of an R register and a 13-bit signed immediate value. The “branch on condition codes without prediction” instruction provides a displacement of ± 8 Mbytes; the “branch on condition codes with prediction” instruction provides a displacement of ± 1 Mbyte; the “branch on register contents” instruction provides a displacement of ± 128 Kbytes; and the CALL instruction’s 30-bit word displacement allows a control transfer to any address within ± 2 gigabytes ($\pm 2^{31}$ bytes).

Note | The return from privileged trap instructions (DONE and RETRY) get their target address from the appropriate TPC or TNPC register.

3.3.4 State Register Access

3.3.4.1 Ancillary State Registers

The read and write ancillary state register instructions read and write the contents of ancillary state registers visible to nonprivileged software (Y, CCR, ASI, PC, TICK, and FPRS) and some registers visible only to privileged software (SOFTINT and STICK_CMPR).

IMPL. DEP. #8-V8-Cs20: Ancillary state registers (ASRs) in the range 0–27 that are not defined in Oracle SPARC Architecture 2015 are reserved for future architectural use.

IMPL. DEP. #9-V8-Cs20: The privilege level required to execute each of the implementation-dependent read/write ancillary state register instructions (for ASRs 28–31) is implementation dependent.

3.3.4.2 PR State Registers

The read and write privileged register instructions (RDPR and WRPR) read and write the contents of state registers visible only to privileged software (TPC, TNPC, TSTATE, TT, TICK, TBA, PSTATE, TL, PIL, CWP, CANSERVE, CANRESTORE, CLEANWIN, OTHERWIN, and WSTATE).

3.3.5 Floating-Point Operate

Floating-point operate (FPop) instructions perform all floating-point calculations; they are register-to-register instructions that operate on the floating-point registers. FPops compute a result that is a function of one, two, or three source operands. The groups of instructions that are considered FPops are listed in *Floating-Point Operate (FPop) Instructions* on page 92.

3.3.6 Conditional Move

Conditional move instructions conditionally copy a value from a source register to a destination register, depending on an integer or floating-point condition code or on the contents of an integer register. These instructions can be used to reduce the number of branches in software.

3.3.7 Register Window Management

Register window instructions manage the register windows. SAVE and RESTORE are nonprivileged and cause a register window to be pushed or popped. FLUSHW is nonprivileged and causes all of the windows except the current one to be flushed to memory. SAVED and RESTORED are used by privileged software to end a window spill or fill trap handler.

3.3.8 SIMD

Oracle SPARC Architecture 2015 includes SIMD (single instruction, multiple data) instructions, also known as "vector" instructions, which allow a single instruction to perform the same operation on multiple data items, totalling 64 bits, such as eight 8-bit, four 16-bit, or two 32-bit data items. These operations are part of the "VIS" extensions.

3.4 Traps

A *trap* is a vectored transfer of control to privileged software through a trap table that may contain the first 8 instructions (32 for some frequently used traps) of each trap handler. The base address of the table is established by software in a state register (the Trap Base Address register, TBA). The displacement within the table is encoded in the type number of each trap and the level of the trap. Part of the trap table is reserved for hardware traps, and part of it is reserved for software traps generated by trap (Tcc) instructions.

A trap causes the current PC and NPC to be saved in the TPC and TNPC registers. It also causes the CCR, ASI, PSTATE, and CWP registers to be saved in TSTATE. TPC, TNPC, and TSTATE are entries in a hardware trap stack, where the number of entries in the trap stack is equal to the number of supported trap levels. A trap also sets bits in the PSTATE register and typically increments the GL register. Normally, the CWP is not changed by a trap; on a window spill or fill trap, however, the CWP is changed to point to the register window to be saved or restored.

A trap can be caused by a Tcc instruction, an asynchronous exception, an instruction-induced exception, or an interrupt request not directly related to a particular instruction. Before executing each instruction, a virtual processor determines if there are any pending exceptions or interrupt requests. If any are pending, the virtual processor selects the highest-priority exception or interrupt request and causes a trap.

See Chapter 12, *Traps*, for a complete description of traps.

Data Formats

The Oracle SPARC Architecture recognizes these fundamental data types:

- Signed integer: 8, 16, 32, and 64 bits
- Unsigned integer: 8, 16, 32, and 64 bits
- SIMD data formats: Uint8 SIMD (32 bits), Int16 SIMD (64 bits), and Int32 SIMD (64 bits)
- Floating point: 32, 64, and 128 bits

The widths of the data types are as follows:

- Byte: 8 bits
- Halfword: 16 bits
- Word: 32 bits
- Tagged word: 32 bits (30-bit value plus 2-bit tag)
- Doubleword/Extended-word: 64 bits
- Quadword: 128 bits

The signed integer values are stored as two's-complement numbers with a width commensurate with their range. Unsigned integer values, bit vectors, Boolean values, character strings, and other values representable in binary form are stored as unsigned integers with a width commensurate with their range. The floating-point formats conform to the IEEE Standard for Binary Floating-point Arithmetic, IEEE Std 754-1985. In tagged words, the least significant two bits are treated as a tag; the remaining 30 bits are treated as a signed integer.

Data formats are described in these sections:

- **Integer Data Formats** on page 24.
- **Floating-Point Data Formats** on page 27.
- **SIMD Data Formats** on page 29.

Names are assigned to individual subwords of the multiword data formats as described in these sections:

- **Signed Integer Doubleword (64 bits)** on page 25.
- **Unsigned Integer Doubleword (64 bits)** on page 26.
- **Floating Point, Double Precision (64 bits)** on page 27.
- **Floating Point, Quad Precision (128 bits)** on page 28.

4.1 Integer Data Formats

TABLE 4-1 describes the width and ranges of the signed, unsigned, and tagged integer data formats.

TABLE 4-1 Signed Integer, Unsigned Integer, and Tagged Format Ranges

Data Type	Width (bits)	Range
Signed integer byte	8	-2^7 to $2^7 - 1$
Signed integer halfword	16	-2^{15} to $2^{15} - 1$
Signed integer word	32	-2^{31} to $2^{31} - 1$
Signed integer doubleword/extended-word	64	-2^{63} to $2^{63} - 1$
Unsigned integer byte	8	0 to $2^8 - 1$
Unsigned integer halfword	16	0 to $2^{16} - 1$
Unsigned integer word	32	0 to $2^{32} - 1$
Unsigned integer doubleword/extended-word	64	0 to $2^{64} - 1$
Integer tagged word	32	0 to $2^{30} - 1$

TABLE 4-2 describes the memory and register alignment for multiword integer data. All registers in the integer register file are 64 bits wide, but can be used to contain smaller (narrower) data sizes. Note that there is no difference between integer extended-words and doublewords in memory; the only difference is how they are represented in registers.

TABLE 4-2 Integer Doubleword/Extended-word Alignment

Subformat Name	Subformat Field	Memory Address		Register Number	
		Required Alignment	Address (big-endian) ¹	Required Alignment	Register Number
SD-0	signed_dbl_integer{63:32}	$n \bmod 8 = 0$	n	$r \bmod 2 = 0$	r
SD-1	signed_dbl_integer{31:0}	$(n + 4) \bmod 8 = 4$	$n + 4$	$(r + 1) \bmod 2 = 1$	$r + 1$
SX	signed_ext_integer{63:0}	$n \bmod 8 = 0$	n	—	r
UD-0	unsigned_dbl_integer{63:32}	$n \bmod 8 = 0$	n	$r \bmod 2 = 0$	r
UD-1	unsigned_dbl_integer{31:0}	$(n + 4) \bmod 8 = 4$	$n + 4$	$(r + 1) \bmod 2 = 1$	$r + 1$
UX	unsigned_ext_integer{63:0}	$n \bmod 8 = 0$	n	—	r

1. The Memory Address in this table applies to big-endian memory accesses. Word and byte order are reversed when little-endian accesses are used.

The data types are illustrated in the following subsections.

4.1.1 Signed Integer Data Types

Figures in this section illustrate the following signed data types:

- Signed integer byte
- Signed integer halfword
- Signed integer word
- Signed integer doubleword
- Signed integer extended-word

4.1.1.1 Signed Integer Byte, Halfword, and Word

FIGURE 4-1 illustrates the signed integer byte, halfword, and word data formats.




FIGURE 4-1 Signed Integer Byte, Halfword, and Word Data Formats

4.1.1.2 Signed Integer Doubleword (64 bits)

FIGURE 4-2 illustrates both components (SD-0 and SD-1) of the signed integer double data format.




FIGURE 4-2 Signed Integer Double Data Format

4.1.1.3 Signed Integer Extended-Word (64 bits)

FIGURE 4-3 illustrates the signed integer extended-word (SX) data format.




FIGURE 4-3 Signed Integer Extended-Word Data Format

4.1.2 Unsigned Integer Data Types

Figures in this section illustrate the following unsigned data types:

- Unsigned integer byte
- Unsigned integer halfword
- Unsigned integer word
- Unsigned integer doubleword
- Unsigned integer extended-word

4.1.2.1 Unsigned Integer Byte, Halfword, and Word

FIGURE 4-4 illustrates the unsigned integer byte data format.




FIGURE 4-4 Unsigned Integer Byte, Halfword, and Word Data Formats

4.1.2.2 Unsigned Integer Doubleword (64 bits)

FIGURE 4-5 illustrates both components (UD-0 and UD-1) of the unsigned integer double data format.




FIGURE 4-5 Unsigned Integer Double Data Format

4.1.2.3 Unsigned Extended Integer (64 bits)

FIGURE 4-6 illustrates the unsigned extended integer (UX) data format.




FIGURE 4-6 Unsigned Extended Integer Data Format

4.1.3 Tagged Word (32 bits)

FIGURE 4-7 illustrates the tagged word data format.




FIGURE 4-7 Tagged Word Data Format

4.2 Floating-Point Data Formats

Single-precision, double-precision, and quad-precision floating-point data types are described below.

4.2.1 Floating Point, Single Precision (32 bits)

FIGURE 4-8 illustrates the floating-point single-precision data format, and TABLE 4-3 describes the formats.




FIGURE 4-8 Floating-Point Single-Precision Data Format

TABLE 4-3 Floating-Point Single-Precision Format Definition

s = sign (1 bit)	
e = biased exponent (8 bits)	
f = fraction (23 bits)	
u = undefined	
Normalized value ($0 < e < 255$):	$(-1)^s \times 2^{e-127} \times 1.f$
Subnormal value ($e = 0$):	$(-1)^s \times 2^{-126} \times 0.f$
Zero ($e = 0, f = 0$)	$(-1)^s \times 0$
Signalling NaN	s = u; e = 255 (max); f = .0uu--uu (At least one bit of the fraction must be nonzero)
Quiet NaN	s = u; e = 255 (max); f = .1uu--uu
$-\infty$ (negative infinity)	s = 1; e = 255 (max); f = .000--00
$+\infty$ (positive infinity)	s = 0; e = 255 (max); f = .000--00

4.2.2 Floating Point, Double Precision (64 bits)

FIGURE 4-9 illustrates both components (FD-0 and FD-1) of the floating-point double-precision data format, and TABLE 4-4 describes the formats.




FIGURE 4-9 Floating-Point Double-Precision Data Format

TABLE 4-4 Floating-Point Double-Precision Format Definition

s	= sign (1 bit)
e	= biased exponent (11 bits)
f	= fraction (52 bits)
u	= undefined
<hr/>	
Normalized value ($0 < e < 2047$):	$(-1)^s \times 2^{e-1023} \times 1.f$
Subnormal value ($e = 0$):	$(-1)^s \times 2^{-1022} \times 0.f$
Zero ($e = 0, f = 0$)	$(-1)^s \times 0$
Signalling NaN	$s = u; e = 2047$ (max); $f = .0uu--uu$ (At least one bit of the fraction must be nonzero)
Quiet NaN	$s = u; e = 2047$ (max); $f = .1uu--uu$
$-\infty$ (negative infinity)	$s = 1; e = 2047$ (max); $f = .000--00$
$+\infty$ (positive infinity)	$s = 0; e = 2047$ (max); $f = .000--00$

4.2.3 Floating Point, Quad Precision (128 bits)

FIGURE 4-10 illustrates all four components (FQ-0 through FQ-3) of the floating-point quad-precision data format, and TABLE 4-5 describes the formats.

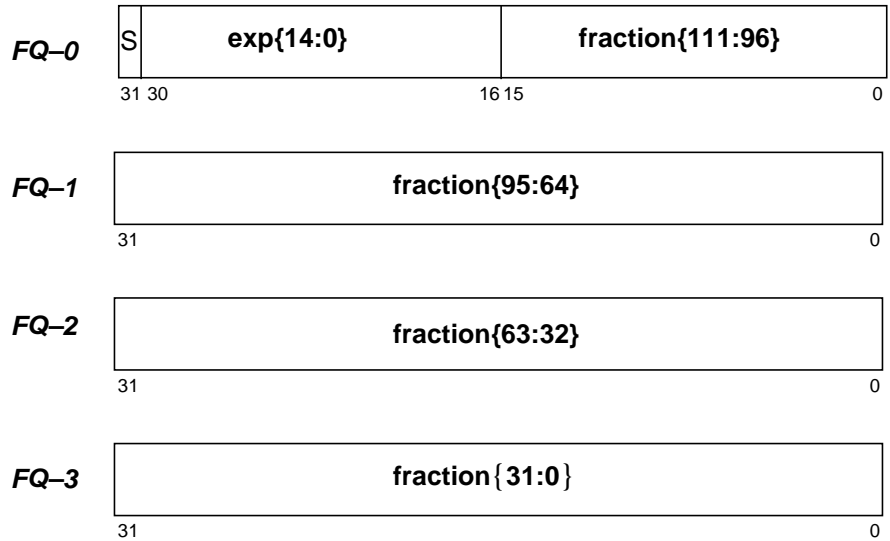


FIGURE 4-10 Floating-Point Quad-Precision Data Format

TABLE 4-5 Floating-Point Quad-Precision Format Definition

s	= sign (1 bit)
e	= biased exponent (15 bits)
f	= fraction (112 bits)
u	= undefined
<hr/>	
Normalized value ($0 < e < 32767$):	$(-1)^s \times 2^{e-16383} \times 1.f$
Subnormal value ($e = 0$):	$(-1)^s \times 2^{-16382} \times 0.f$
Zero ($e = 0, f = 0$)	$(-1)^s \times 0$
Signalling NaN	$s = u; e = 32767$ (max); $f = .0uu--uu$ (At least one bit of the fraction must be nonzero)

TABLE 4-5 Floating-Point Quad-Precision Format Definition (*Continued*)

s = sign (1 bit)	
e = biased exponent (15 bits)	
f = fraction (112 bits)	
u = undefined	
Quiet NaN	s = u; e = 32767 (max); f = .1uu--uu
- ∞ (negative infinity)	s = 1; e = 32767 (max); f = .000--00
+ ∞ (positive infinity)	s = 0; e = 32767 (max); f = .000--00

4.2.4 Floating-Point Data Alignment in Memory and Registers

TABLE 4-6 describes the address and memory alignment for floating-point data.

TABLE 4-6 Floating-Point Doubleword and Quadword Alignment

Subformat Name	Subformat Field	Memory Address		Register Number	
		Required Alignment	Address (big-endian)*	Required Alignment	Register Number
FD-0	s:exp{10:0}:fraction{51:32}	0 mod 4 †	n	0 mod 2	f
FD-1	fraction{31:0}	0 mod 4 †	n + 4	1 mod 2	f + 1 [◇]
FQ-0	s:exp{14:0}:fraction{111:96}	0 mod 4 ‡	n	0 mod 4	f
FQ-1	fraction{95:64}	0 mod 4 ‡	n + 4	1 mod 4	f + 1 [◇]
FQ-2	fraction{63:32}	0 mod 4 ‡	n + 8	2 mod 4	f + 2
FQ-3	fraction{31:0}	0 mod 4 ‡	n + 12	3 mod 4	f + 3 [◇]

* The memory Address in this table applies to big-endian memory accesses. Word and byte order are reversed when little-endian accesses are used.

† Although a floating-point doubleword is required only to be word-aligned in memory, it is recommended that it be doubleword-aligned (that is, the address of its FD-0 word should be 0 mod 8 so that it can be accessed with doubleword loads/stores instead of multiple singleword loads/stores).

‡ Although a floating-point quadword is required only to be word-aligned in memory, it is recommended that it be quadword-aligned (that is, the address of its FQ-0 word should be 0 mod 16).

◇ Note that this 32-bit floating-point register is only directly addressable in the lower half of the register file (that is, if its register number is ≤ 31).

4.3 SIMD Data Formats

SIMD (single instruction/multiple data) instructions perform identical operations on multiple data contained (“packed”) in each source operand. This section describes the data formats used by SIMD instructions.

Conversion between the different SIMD data formats can be achieved through SIMD multiplication or by the use of the SIMD data formatting instructions.

Programming Note The SIMD data formats can be used in graphics calculations to represent intensity values for an image (e.g., α , B, G, R). Intensity values are typically grouped in one of two ways, when using SIMD data formats:

- Band interleaved images, with the various color components of a point in the image stored together, and
- Band sequential images, with all of the values for one color component stored together.

4.3.1 Uint8 SIMD Data Format

The Uint8 SIMD data format consists of four unsigned 8-bit integers contained in a 32-bit word (see FIGURE 4-11).




FIGURE 4-11 Uint8 SIMD Data Format

4.3.2 Int16 SIMD Data Formats

The Int16 SIMD data format consists of four signed 16-bit integers contained in a 64-bit word (see FIGURE 4-12).




FIGURE 4-12 Int16 SIMD Data Format

4.3.3 Int32 SIMD Data Format

The Int32 SIMD data format consists of two signed 32-bit integers contained in a 64-bit word (see FIGURE 4-13).




FIGURE 4-13 Int32 SIMD Data Format

Programming Note The integer SIMD data formats can be used to hold fixed-point data. The position of the binary point in a SIMD datum is implied by the programmer and does not influence the computations performed by instructions that operate on that SIMD data format.

Registers

The following registers are described in this chapter:

- **General-Purpose R Registers on page 32.**
- **Floating-Point Registers on page 38.**
- **Floating-Point State Register (FSR) on page 42.**
- **Ancillary State Registers on page 48.** The following registers are included in this category:
 - **32-bit Multiply/Divide Register (y) (ASR 0) on page 50.**
 - **Integer Condition Codes Register (ccr) (ASR 2) on page 50.**
 - **Address Space Identifier (asi) Register (ASR 3) on page 51.**
 - **Tick (tick) Register (ASR 4) on page 52.**
 - **Program Counters (pc, npc) (ASR 5) on page 53.**
 - **Floating-Point Registers State (fprs) Register (ASR 6) on page 53.**
 - **General Status Register (gsr) (ASR 19) on page 54.**
 - **softintP Register (ASRs 20, 21, 22) on page 55.**
 - **softint_setP Pseudo-Register (ASR 20) on page 55.**
 - **softint_clrP Pseudo-Register (ASR 21) on page 56.**
 - **System Tick (stick) Register (ASR 24) on page 56.**
 - **System Tick Compare (stick_cmprP) Register (ASR 25) on page 57.**
 - **Compatibility Feature Register (cfr) (ASR 26) on page 59.**
 - **Compatibility Feature Register (cfr) (ASR 26) on page 59.**
 - **Mwait Count (mwait) Register (ASR 28) on page 61.**
- **Register-Window PR State Registers on page 62.** The following registers are included in this subcategory:
 - **Current Window Pointer (cwpp) Register (PR 9) on page 63.**
 - **Savable Windows (cansaveP) Register (PR 10) on page 63.**
 - **Restorable Windows (canrestoreP) Register (PR 11) on page 64.**
 - **Clean Windows (cleanwinP) Register (PR 12) on page 64.**
 - **Other Windows (otherwinP) Register (PR 13) on page 65.**
 - **Window State (wstateP) Register (PR 14) on page 65.**
- **Non-Register-Window PR State Registers on page 66.** The following registers are included in this subcategory:
 - **Trap Program Counter (tpcP) Register (PR 0) on page 67.**
 - **Trap Next PC (tnpcP) Register (PR 1) on page 67.**
 - **Trap State (tstateP) Register (PR 2) on page 68.**
 - **Trap Type (ttP) Register (PR 3) on page 69.**
 - **Trap Base Address (tbaP) Register (PR 5) on page 69.**
 - **Processor State (pstateP) Register (PR 6) on page 69.**
 - **Trap Level Register (tlP) (PR 7) on page 73.**
 - **Processor Interrupt Level (pilP) Register (PR 8) on page 74.**
 - **Global Level Register (glP) (PR 16) on page 75.**

There are additional registers that may be accessed through ASIs; those registers are described in Chapter 10, *Address Space Identifiers (ASIs)*.

5.1 Reserved Register Fields

Some register bit fields in this specification are explicitly marked as "reserved". In addition, for convenience, some registers in this chapter are illustrated as fewer than 64 bits wide. Any bits not illustrated are implicitly reserved and treated as if they were explicitly marked as reserved.

Reserved bits, whether explicitly or implicitly reserved, may be assigned meaning in future versions of the architecture.

To ensure that existing software will continue to operate correctly, software must take into account that reserved register bits may be used in the future. The following Programming and Implementation Notes support that intent.

Programming Notes	<p>Software should ensure that when a reserved register field is written, it is only written with (1) the value zero or (2) a value previously read from that field.</p> <p>If software writes a reserved register field to any value other than (1) zero or (2) a value previously read from that field, it is considered a software error. Such an error:</p> <ul style="list-style-type: none">• may or may not be detected or reported (for example, by a trap) by Oracle SPARC Architecture 2015 processors (and software should not expect that it will be)• may cause a trap or cause other unintended behavior when executed on future Oracle SPARC Architecture processors <p>When a register is read, software should not assume that register fields reserved in Oracle SPARC Architecture 2015 will read as 0 or any other particular value, either now or in the future.</p>
--------------------------	--

Implementation Notes	<p>When a register is read by software, an Oracle SPARC Architecture 2015 virtual processor should return a value of zero for any bits reserved in Oracle SPARC Architecture 2015</p> <p>When software attempts to change the contents of a register field that is reserved in UltraSPARC Architecture 200x by writing a value to that field that differs from the current contents of that field, an UltraSPARC Architecture 200x virtual processor will either ignore the write to that field or cause an exception. "Current contents" means the contents that software would observe if it read that field (nominally zero).</p>
-----------------------------	--

5.2 General-Purpose R Registers

An Oracle SPARC Architecture virtual processor contains an array of general-purpose 64-bit R registers. The array is partitioned into $MAXPGL + 1$ sets of eight *global* registers, plus $N_REG_WINDOWS$ groups of 16 registers each. The value of $N_REG_WINDOWS$ in an Oracle SPARC Architecture implementation falls within the range 3 to 32 (inclusive).

One set of 8 global registers is always visible. At any given time, a group of 24 registers, known as a *register window*, is also visible. A register window comprises the 16 registers from the current 16-register group (referred to as 8 *in* registers and 8 *local* registers), plus half of the registers from the next 16-register group (referred to as 8 *out* registers). See FIGURE 5-1.

SPARC instructions use 5-bit fields to reference R registers. That is, 32 R registers are visible to software at any moment. Which 32 out of the full set of R registers are visible is described in the following sections. The visible 32 R registers are named R[0] through R[31], illustrated in FIGURE 5-1.

R[31]	i7	- - - - -
R[30]	i6	
R[29]	i5	
R[28]	i4	
R[27]	i3	<i>ins</i>
R[26]	i2	
R[25]	i1	
R[24]	i0	- - - - -
R[23]	l7	
R[22]	l6	
R[21]	l5	
R[20]	l4	<i>locals</i>
R[19]	l3	
R[18]	l2	
R[17]	l1	
R[16]	l0	- - - - -
R[15]	o7	
R[14]	o6	
R[13]	o5	
R[12]	o4	<i>outs</i>
R[11]	o3	
R[10]	o2	
R[9]	o1	
R[8]	o0	- - - - -
R[7]	g7	
R[6]	g6	
R[5]	g5	
R[4]	g4	<i>globals</i>
R[3]	g3	
R[2]	g2	
R[1]	g1	
R[0]	g0	- - - - -

FIGURE 5-1 General-Purpose Registers (as Visible at Any Given Time)

5.2.1 Global R Registers (A1)

Registers R[0]–R[7] refer to a set of eight registers called the *global* registers (labelled g0 through g7). At any time, one of $MAXPGL + 1$ sets of eight registers is enabled and can be accessed as the current set of global registers. The currently enabled set of global registers is selected by the GL register. See *Global Level Register (gLP) (PR 16)* on page 75.

Global register zero (G0) always reads as zero; writes to it have no software-visible effect.

5.2.2 Windowed R Registers (A1)

A set of 24 R registers that is visible as R[8]–R[31] at any given time is called a “register window”. The registers that become R[8]–R[15] in a register window are called the *out* registers of the window. Note that the *in* registers of a register window become the *out* registers of an adjacent register window. See TABLE 5-1 and FIGURE 5-2.

The names *in*, *local*, and *out* originate from the fact that the *out* registers are typically used to pass parameters from (out of) a calling routine and that the called routine receives those parameters as its *in* registers.

TABLE 5-1 Window Addressing

Windowed Register Address	R Register Address
<i>in</i> [0] – <i>in</i> [7]	R[24] – R[31]
<i>local</i> [0] – <i>local</i> [7]	R[16] – R[23]
<i>out</i> [0] – <i>out</i> [7]	R[8] – R[15]
<i>global</i> [0] – <i>global</i> [7]	R[0] – R[7]

V9 Compatibility Notes In the SPARC V9 architecture, the number of 16-register windowed register sets, *N_REG_WINDOWS*, ranges from 3[†] to 32 (impl. dep. #2-V8).

The maximum global register set index in the Oracle SPARC Architecture, *MAXPGL*, ranges from 2 to 15. The number of implemented global register sets is *MAXPGL* + 1.

The total number of R registers in a given Oracle SPARC Architecture implementation is:

$$(N_REG_WINDOWS \times 16) + ((MAXPGL + 1) \times 8)$$

Therefore, an Oracle SPARC Architecture processor may contain from 72 to 640 R registers.

†. The controlling equation for register window operation, as described in 5.6.7.1 on page 66, is:

$$CANSAVE + CANRESTORE + OTHERWIN = N_REG_WINDOWS - 2$$

Since *N_REG_WINDOWS* cannot be negative, the minimum number of implemented register windows is “2”. However, since the *SAVED* and *RESTORED* instructions increment *CANSAVE* and *CANRESTORE*, the minimum value of *N_REG_WINDOWS* in practice increases to “3”. An implementation with *N_REG_WINDOWS* = 2 would not be able to support use of the *SAVED* and *RESTORED* instructions — in such an implementation, a spill trap handler would have to emulate the *SAVE* instruction (the one that caused the spill trap) in its entirety (including its addition semantics) and the spill handler would have to end with a *DONE* instruction instead of *RETRY*.

Because the windows overlap, the number of windows available to software is 1 less than the number of implemented windows; that is, $N_REG_WINDOWS - 1$. When the register file is full, the *outs* of the newest window are the *ins* of the oldest window, which still contains valid data.


Window overflow is detected by the CANSAVE register, and window underflow is detected by the CANRESTORE register, both of which are controlled by privileged software. A window overflow (underflow) condition causes a window spill (fill) trap.

When a new register window is made visible through use of a SAVE instruction, the *local* and *out* registers are guaranteed to contain either zeroes or valid data from the current context. If software executes a RESTORE and later executes a SAVE, then the contents of the resulting window's *local* and *out* registers are not guaranteed to be preserved between the RESTORE and the SAVE¹. Those registers may even have been written with "dirty" data, that is, data created by software running in a different context. However, if the *clean_window* protocol is being used, system software must guarantee that registers in the current window after a SAVE always contains only zeroes or valid data from that context. See *Clean Windows (cleanwinP) Register (PR 12)* on page 64, *Savable Windows (cansaveP) Register (PR 10)* on page 63, and *Restorable Windows (canrestoreP) Register (PR 11)* on page 64.

Implementation	An Oracle SPARC Architecture virtual processor supports the
Note	guarantee in the preceding paragraph of "either zeroes or valid data from the current context"; it may do so either in hardware or in a combination of hardware and system software.

Register Window Management Instructions on page 90 describes how the windowed integer registers are managed.

¹. For example, any of those 16 registers might be altered due to the occurrence of a trap between the RESTORE and the SAVE, or might be altered during the RESTORE operation due to the way that register windows are implemented. After a RESTORE instruction executes, software must assume that the values of the affected 16 registers from before the RESTORE are unrecoverable.



$$CANSERVE + CANRESTORE + OTHERWIN = N_REG_WINDOWS - 2$$

The current window (window 0) and the overlap window (window 5) account for the two windows in the right side of the equation. The “overlap window” is the window that must remain unused because its *ins* and *outs* overlap two other valid windows.

FIGURE 5-3 Windowed R Registers for $N_REG_WINDOWS = 8$

In FIGURE 5-3, $N_REG_WINDOWS = 8$. The eight *global* registers are not illustrated. $CWP = 0$, $CANSERVE = 4$, $OTHERWIN = 1$, and $CANRESTORE = 1$. If the procedure using window w0 executes a RESTORE, then window w7 becomes the current window. If the procedure using window w0 executes a SAVE, then window w1 becomes the current window.

5.2.3 Special R Registers

The use of two of the R registers is fixed, in whole or in part, by the architecture:

- The value of R[0] is always zero; writes to it have no program-visible effect.
- The CALL instruction writes its own address into register R[15] (*out* register 7).

Register-Pair Operands. LDTW, LDTWA, STTW, and STTWA instructions access a pair of words (“twin words”) in adjacent R registers and require even-odd register alignment. The least significant bit of an R register number in these instructions is unused and must always be supplied as 0 by software.

When the R[0]–R[1] register pair is used as a destination in LDTW or LDTWA, only R[1] is modified. When the R[0]–R[1] register pair is used as a source in STTW or STTWA, 0 is read from R[0], so 0 is written to the 32-bit word at the lowest address, and the least significant 32 bits of R[1] are written to the 32-bit word at the highest address.

An attempt to execute an LDTW, LDTWA, STTW, or STTWA instruction that refers to a misaligned (odd) destination register number causes an *illegal_instruction* trap.

5.3 Floating-Point Registers A1

The floating-point register set consists of sixty-four 32-bit registers, which may be accessed as follows:

- Sixteen 128-bit quad-precision registers, referenced as $F_Q[0]$, $F_Q[4]$, ..., $F_Q[60]$
- Thirty-two 64-bit double-precision registers, referenced as $F_D[0]$, $F_D[2]$, ..., $F_D[62]$
- Thirty-two 32-bit single-precision registers, referenced as $F_S[0]$, $F_S[1]$, ..., $F_S[31]$ (only the lower half of the floating-point register file can be accessed as single-precision registers)

The floating-point registers are arranged so that some of them overlap, that is, are aliased. The layout and numbering of the floating-point registers are shown in TABLE 5-2. Unlike the windowed R registers, all of the floating-point registers are accessible at any time. The floating-point registers can be read and written by floating-point operate (FPop1/FPop2 format) instructions, by load/store single/double/quad floating-point instructions, by VIS™ instructions, and by block load and block store instructions.

TABLE 5-2 Floating-Point Registers, with Aliasing (1 of 3)

Single Precision (32-bit)		Double Precision (64-bit)		Quad Precision (128-bit)		
Register	Assembly Language	Bits	Register	Assembly Language	Register	Assembly Language
$F_S[0]$	%f0	63:32	$F_D[0]$	%d0	127:64	
$F_S[1]$	%f1	31:0			$F_Q[0]$	%q0
$F_S[2]$	%f2	63:32	$F_D[2]$	%d2		
$F_S[3]$	%f3	31:0			$F_Q[4]$	%q4
$F_S[4]$	%f4	63:32	$F_D[4]$	%d4		
$F_S[5]$	%f5	31:0			$F_Q[8]$	%q8
$F_S[6]$	%f6	63:32	$F_D[6]$	%d6		
$F_S[7]$	%f7	31:0			$F_Q[8]$	%q8
$F_S[8]$	%f8	63:32	$F_D[8]$	%d8		
$F_S[9]$	%f9	31:0			$F_Q[10]$	%q10
$F_S[10]$	%f10	63:32	$F_D[10]$	%d10		
$F_S[11]$	%f11	31:0				

TABLE 5-2 Floating-Point Registers, with Aliasing (2 of 3)

Single Precision (32-bit)		Double Precision (64-bit)		Quad Precision (128-bit)	
Register	Assembly Language	Bits	Register	Assembly Language	Bits
F _S [12]	%f12	63:32	F _D [12]	%d12	127:64
F _S [13]	%f13	31:0			
F _S [14]	%f14	63:32	F _D [14]	%d14	63:0
F _S [15]	%f15	31:0			
F _S [16]	%f16	63:32	F _D [16]	%d16	127:64
F _S [17]	%f17	31:0			
F _S [18]	%f18	63:32	F _D [18]	%d18	63:0
F _S [19]	%f19	31:0			
F _S [20]	%f20	63:32	F _D [20]	%d20	127:64
F _S [21]	%f21	31:0			
F _S [22]	%f22	63:32	F _D [22]	%d22	63:0
F _S [23]	%f23	31:0			
F _S [24]	%f24	63:32	F _D [24]	%d24	127:64
F _S [25]	%f25	31:0			
F _S [26]	%f26	63:32	F _D [26]	%d26	63:0
F _S [27]	%f27	31:0			
F _S [28]	%f28	63:32	F _D [28]	%d28	127:64
F _S [29]	%f29	31:0			
F _S [30]	%f30	63:32	F _D [30]	%d30	63:0
F _S [31]	%f31	31:0			
		63:32	F _D [32]	%d32	127:64
		31:0			
		63:32	F _D [34]	%d34	63:0
		31:0			
		63:32	F _D [36]	%d36	127:64
		31:0			
		63:32	F _D [38]	%d38	63:0
		31:0			
		63:32	F _D [40]	%d40	127:64
		31:0			
		63:32	F _D [42]	%d42	63:0
		31:0			
		63:32	F _D [44]	%d44	127:64
		31:0			
		63:32	F _D [46]	%d46	63:0
		31:0			
					F _Q [12] %q12
					F _Q [16] %q16
					F _Q [20] %q20
					F _Q [24] %q24
					F _Q [28] %q28
					F _Q [32] %q32
					F _Q [36] %q36
					F _Q [40] %q40
					F _Q [44] %q44

TABLE 5-2 Floating-Point Registers, with Aliasing (3 of 3)

Single Precision (32-bit)		Double Precision (64-bit)		Quad Precision (128-bit)	
Register	Assembly Language	Bits	Register Assembly Language	Bits	Register Assembly Language
		63:32	F _D [48] %d48	127:64	F _Q [48] %q48
		31:0			
		63:32	F _D [50] %d50	63:0	
		31:0			
		63:32	F _D [52] %d52	127:64	F _Q [52] %q52
		31:0			
		63:32	F _D [54] %d54	63:0	
		31:0			
		63:32	F _D [56] %d56	127:64	F _Q [56] %q56
		31:0			
		63:32	F _D [58] %d58	63:0	
		31:0			
		63:32	F _D [60] %d60	127:64	F _Q [60] %q60
		31:0			
		63:32	F _D [62] %d62	63:0	
		31:0			

5.3.1 Floating-Point Register Number Encoding

Register numbers for single, double, and quad registers are encoded differently in the 5-bit register number field of a floating-point instruction. If the bits in a register number field are labelled b{4} ... b{0} (where b{4} is the most significant bit of the register number), the encoding of floating-point register numbers into 5-bit instruction fields is as given in TABLE 5-3.

TABLE 5-3 Floating-Point Register Number Encoding

Register Operand Type	Full 6-bit Register Number						Encoding in a 5-bit Register Field in an Instruction				
Single	0	b{4}	b{3}	b{2}	b{1}	b{0}	b{4}	b{3}	b{2}	b{1}	b{0}
Double	b{5}	b{4}	b{3}	b{2}	b{1}	0	b{4}	b{3}	b{2}	b{1}	b{5}
Quad	b{5}	b{4}	b{3}	b{2}	0	0	b{4}	b{3}	b{2}	0	b{5}

SPARC V8 Compatibility Note In the SPARC V8 architecture, bit 0 of double and quad register numbers encoded in instruction fields was required to be zero. Therefore, all SPARC V8 floating-point instructions can run unchanged on an Oracle SPARC Architecture virtual processor, using the encoding in TABLE 5-3.

5.3.2 Double and Quad Floating-Point Operands

A single 32-bit F register can hold one single-precision operand; a double-precision operand requires an aligned pair of F registers, and a quad-precision operand requires an aligned quadruple of F registers. At a given time, the floating-point registers can hold a maximum of 32 single-precision, 16 double-precision, or 8 quad-precision values in the lower half of the floating-point register file, plus an additional 16 double-precision or 8 quad-precision values in the upper half, or mixtures of the three sizes.

Programming Note The upper 16 double-precision (upper 8 quad-precision) floating-point registers cannot be directly loaded by 32-bit load instructions. Therefore, double- or quad-precision data that is only word-aligned in memory cannot be directly loaded into the upper registers with LDF[A] instructions. The following guidelines are recommended:

1. Whenever possible, align floating-point data in memory on proper address boundaries. If access to a datum is required to be atomic, the datum *must* be properly aligned.
2. If a double- or quad-precision datum is not properly aligned in memory or is still aligned on a 4-byte boundary, and access to the datum in memory is not required to be atomic, then software should attempt to allocate a register for it in the lower half of the floating-point register file so that the datum can be loaded with multiple LDF[A] instructions.
3. If the only available registers for such a datum are located in the upper half of the floating-point register file and access to the datum in memory is not required to be atomic, the word-aligned datum can be loaded into them by one of two methods:
 - Load the datum into an upper register by using multiple LDF[A] instructions to first load it into a double- or quad-precision register in the lower half of the floating-point register file, then copy that register to the desired destination register in the upper half.

Use an LDDF[A] or LDQF[A] instruction to perform the load directly into the upper floating-point register, understanding that use of these instructions on poorly aligned data can cause a trap (*LDDF_mem_not_aligned* or *LDQF_mem_not_aligned*) on some implementations, possibly slowing down program execution significantly.

Programming Note If an Oracle SPARC Architecture 2015 implementation does not implement a particular quad floating-point arithmetic operation in hardware and an invalid quad register operand is specified, the *illegal_instruction* trap occurs because it has higher priority.

Implementation Note Oracle SPARC Architecture 2011 implementations do not implement any quad floating-point arithmetic operations in hardware. Therefore, an attempt to execute any of them results in a trap on the *illegal_instruction* exception.

5.4 Floating-Point State Register (FSR) A1

The Floating-Point State register (FSR) fields, illustrated in FIGURE 5-4, contain FPU mode and status information. The lower 32 bits of the FSR are read and written by the (deprecated) STFSR and LDFSR instructions, respectively. The 64-bit FSR register is read by the STXFSR instruction and written by the LDXFSR instruction. The *ver*, *ftt*, *qne*, unimplemented (for example, *ns*), and reserved (“—”) fields of FSR are not modified by either LDFSR or LDXFSR. The LDXEFSR instruction can be used to write all implemented fields of FSR (notably *ftt*), but not the *ver*, *qne*, and reserved fields.




FIGURE 5-4 FSR Fields

Bits 63–38, 29–28, 21–20, and 12 of FSR are reserved. When read by an STXFSR instruction, these bits always read as zero

Programming Note For future compatibility, software should issue LDXFSR instructions only with zero values in these bits or values of these bits exactly as read by a previous STXFSR.

The subsections on pages 42 through 48 describe the remaining fields in the FSR.

5.4.1 Floating-Point Condition Codes (*fcc0*, *fcc1*, *fcc2*, *fcc3*)

The four sets of floating-point condition code fields are labelled *fcc0*, *fcc1*, *fcc2*, and *fcc3* (*fccn* refers to any of the floating-point condition code fields).

The *fcc0* field consists of bits 11 and 10 of the FSR, *fcc1* consists of bits 33 and 32, *fcc2* consists of bits 35 and 34, and *fcc3* consists of bits 37 and 36. Execution of a floating-point compare instruction (FCMP or FCMPE) updates one of the *fccn* fields in the FSR, as selected by the compare instruction. The *fccn* fields are read by STXFSR and written by LDXFSR and LDXEFSR. The *fcc0* field can also be read and written by STFSR and LDFSR, respectively. FBfcc and FBPfcc instructions base their control transfers on the content of these fields. The MOVcc and FMOVcc instructions can conditionally copy a register, based on the contents of these fields.

In TABLE 5-4, f_{rs1} and f_{rs2} correspond to the single, double, or quad values in the floating-point registers specified by a floating-point compare instruction’s *rs1* and *rs2* fields. The question mark (?) indicates an unordered relation, which is true if either f_{rs1} or f_{rs2} is a signalling NaN or a quiet NaN. If FCMP or FCMPE generates an *fp_exception_ieee_754* exception, then *fccn* is unchanged.

TABLE 5-4 Floating-Point Condition Codes (*fccn*) Fields of FSR

	Content of <i>fccn</i>			
	0	1	2	3
Indicated Relation (FCMP*, FCMPE*)	$F[rs1] = F[rs2]$	$F[rs1] < F[rs2]$	$F[rs1] > F[rs2]$	$F[rs1] ? F[rs2]$ (unordered)

TABLE 5-4 Floating-Point Condition Codes (*fccn*) Fields of FSR

	Content of <i>fccn</i>			
	0	1	2	3
Indicated Relation (FLCMP*)	$F[rs1] \geq F[rs2]$, neither operand is NaN	$F[rs1] < F[rs2]$, neither operand is NaN	$F[rs1]$ is NaN, $F[rs2]$ is not NaN regardless of $F[rs1]$	$F[rs2]$ is NaN, $F[rs1]$ is not NaN regardless of $F[rs2]$

5.4.2 Rounding Direction (*rd*)

Bits 31 and 30 select the rounding direction for floating-point results according to IEEE Std 754-1985. TABLE 5-5 shows the encodings.

TABLE 5-5 Rounding Direction (*rd*) Field of FSR

<i>rd</i>	Round Toward
0	Nearest (even, if tie)
1	0
2	$+\infty$
3	$-\infty$

If the interval mode bit of the General Status register has a value of 1 ($GSR.im = 1$), then the value of $FSR.rd$ is ignored and floating-point results are instead rounded according to $GSR.irnd$. See *General Status Register (gsr)* (ASR 19) on page 54 for further details.

5.4.3 Trap Enable Mask (*tem*)

Bits 27 through 23 are enable bits for each of the five IEEE-754 floating-point exceptions that can be indicated in the current_exception field (*cxexc*). See FIGURE 5-6 on page 47. If a floating-point instruction generates one or more exceptions and the *tem* bit corresponding to any of the exceptions is 1, then this condition causes an *fp_exception_ieee_754* trap. A *tem* bit value of 0 prevents the corresponding IEEE 754 exception type from generating a trap.

5.4.4 Nonstandard Floating-Point (*ns*)

When $FSR.ns = 1$, it causes a SPARC V9 virtual processor to produce implementation-defined results that may or may not correspond to IEEE Std 754-1985 (impl. dep. #18-V8).

For an implementation in which no nonstandard floating-point mode exists, the *ns* bit of FSR should always read as 0 and writes to it should be ignored.

For detailed requirements for the case when an Oracle SPARC Architecture processor elects to implement floating-point nonstandard mode, see *Floating-Point Nonstandard Mode* on page 393.

5.4.5 FPU Version (*ver*)

IMPL. DEP. #19-V8: Bits 19 through 17 identify one or more particular implementations of the FPU architecture.

For each SPARC V9 IU implementation, there may be one or more FPU implementations, or none. $FSR.ver$ identifies the particular FPU implementation present. The value in $FSR.ver$ for each implementation is strictly implementation dependent. Consult the appropriate document for each implementation for its setting of $FSR.ver$.

$FSR.ver = 7$ is reserved to indicate that no hardware floating-point controller is present.

The `ver` field of `FSR` is read-only; it cannot be modified by the `LDFSR`, `LDXEFSR`, or `LDXFBSR` instructions.

5.4.6 Floating-Point Trap Type (ftt)

Several conditions can cause a floating-point exception trap. When a floating-point exception trap occurs, `FSR.ftt` (`FSR{16:14}`) identifies the cause of the exception, the “floating-point trap type.” After a floating-point exception occurs, `FSR.ftt` encodes the type of the floating-point exception until it is cleared (set to 0) by execution of an `STFSR`, `STXFBSR`, or `FPop` that does not cause a trap due to a floating-point exception.

The `FSR.ftt` field can be read by a `STFSR` or `STXFBSR` instruction. The `LDFSR` and `LDXFBSR` instructions do not affect `FSR.ftt`. `FSR.ftt` can be directly written to a specific value by the `LDXEFSR` instruction.

Privileged software that handles floating-point traps must execute an `STFSR` (or `STXFBSR`) to determine the floating-point trap type. `STFSR` and `STXFBSR` set `FSR.ftt` to zero after the store completes without error. If the store generates an error and does not complete, `FSR.ftt` remains unchanged.

Programming Note Neither `LDFSR` nor `LDXFBSR` can be used for the purpose of clearing the `ftt` field, since both leave `ftt` unchanged. However, executing a nontrapping floating-point operate (`FPop`) instruction such as “`fmovs %f0,%f0`” prior to returning to nonprivileged mode will zero `FSR.ftt`. The `ftt` field remains zero until the next `FPop` instruction completes execution.

The `LDXEFSR` instruction can be used to directly write a specific value to `FSR.ftt`. (In fact, `LDXEFSR` is the *only* way to directly write a non-zero value to `FSR.ftt`.)

`FSR.ftt` encodes the primary condition (“floating-point trap type”) that caused the generation of an `fp_exception_other` or `fp_exception_ieee_754` exception. It is possible for more than one such condition to occur simultaneously; in such a case, only the highest-priority condition will be encoded in `FSR.ftt`. The conditions leading to `fp_exception_other` and `fp_exception_ieee_754` exceptions, their relative priorities, and the corresponding `FSR.ftt` values are listed in TABLE 5-6. Note that the `FSR.ftt` values 4 and 5 were defined in the SPARC V9 architecture but are not currently in use, and that the value 7 is reserved for future architectural use.

TABLE 5-6 FSR Floating-Point Trap Type (ftt) Field

Condition Detected During Execution of an FPop	Relative Priority (1 = highest)	Result	
		FSR.ftt Set to Value	Exception Generated
<code>invalid_fp_register</code>	20	6	<code>fp_exception_other</code>
<code>unfinished_FPop</code>	30	2	<code>fp_exception_other</code>
<code>IEEE_754_exception</code>	40	1	<code>fp_exception_ieee_754</code>
<i>Reserved</i>	—	3, 4, 5, 7	—
(none detected)	—	0	—

The `IEEE_754_exception` and `unfinished_FPop` conditions will likely arise occasionally in the normal course of computation and must be recoverable by system software.

When a floating-point trap occurs, the following results are observed by user software:

1. The value of `aexc` is unchanged.
2. When an `fp_exception_ieee_754` trap occurs, a bit corresponding to the trapping exception is set in `cexc`. On other traps, the value of `cexc` is unchanged.
3. The source and destination registers are unchanged.

4. The value of *fccn* is unchanged.

The foregoing describes the result seen by a user trap handler if an IEEE exception is signalled, either immediately from an *fp_exception_ieee_754* exception or after recovery from an unfinished_FPop. In either case, *cexc* as seen by the trap handler reflects the exception causing the trap.

In the cases of an *fp_exception_other* exception with a floating-point trap type of unfinished_FPop that does not subsequently generate an IEEE trap, the recovery software should set *cexc*, *aexc*, and the destination register or *fccn*, as appropriate.

ftt = 1 (IEEE_754_exception). The IEEE_754_exception floating-point trap type indicates the occurrence of a floating-point exception conforming to IEEE Std 754-1985. The IEEE 754 exception type (overflow, inexact, etc.) is set in the *cexc* field. The *aexc* and *fccn* fields and the destination F register are unchanged.

ftt = 2 (unfinished_FPop). The unfinished_FPop floating-point trap type indicates that the virtual processor was unable to generate correct results or that exceptions as defined by IEEE Std 754-1985 have occurred. In cases where exceptions have occurred, the *cexc* field is unchanged.

Implementation Note	Implementations are encouraged to support standard IEEE 754 floating-point arithmetic with reasonable performance (that is, without generating <i>fp_exception_other</i> with FSR.ftt=unfinished_FPop) in all cases, even if some cases are slower than others.
----------------------------	---

IMPL. DEP. #248-U3: The conditions under which an *fp_exception_other* exception with floating-point trap type of unfinished_FPop can occur are implementation dependent. An implementation may cause *fp_exception_other* with FSR.ftt = unfinished_FPop under a different (but specified) set of conditions.

ftt = 3 (Reserved).

SPARC V9 Compatibility Note	In SPARC V9, FSR.ftt = 3 was defined to be "unimplemented_FPop". All conditions which used to cause <i>fp_exception_other</i> with FSR.ftt = 3 now cause an <i>illegal_instruction</i> exception, instead. FSR.ftt = 3 is now reserved and available for other future uses.
------------------------------------	---

ftt = 4 (Reserved).

SPARC V9 Compatibility Note	In the SPARC V9 architecture, FSR.ftt = 4 was defined to be "sequence_error", for use with certain error conditions associated with a floating-point queue (FQ). Since Oracle SPARC Architecture implementations generate precise (rather than deferred) traps for floating-point operations, an FQ is not needed; therefore sequence_error conditions cannot occur and ftt = 4 has been returned to the pool of reserved ftt values.
------------------------------------	---

ftt = 5 (Reserved).

SPARC V9 Compatibility Note	In the SPARC V9 architecture, FSR.ftt = 5 was defined to be "hardware_error", for use with hardware error conditions associated with an external floating-point unit (FPU) operating asynchronously to the main processor (IU). Since Oracle SPARC Architecture processors are now implemented with an integral FPU, a hardware error in the FPU can generate an exception directly, rather than indirectly report the error through FSR.ftt (as was required when FPUs were external to IUs). Therefore, ftt = 5 has been returned to the pool of reserved ftt values.
------------------------------------	---

ftt = 6 (invalid_fp_register). This trap type indicates that one or more F register operands of an FPop are misaligned; that is, a quad-precision register number is not 0 **mod** 4. An implementation generates an *fp_exception_other* trap with FSR.ftt = invalid_fp_register in this case.

Implementation Note	If an Oracle SPARC Architecture 2015 processor does not implement a particular quad FPop in hardware, that FPop generates an <i>illegal_instruction</i> exception instead of <i>fp_exception_other</i> with FSR.ftt = 6 (invalid_fp_register), regardless of the specified F registers.
----------------------------	---

5.4.7 Accrued Exceptions (aexc)

Bits 9 through 5 accumulate IEEE_754 floating-point exceptions as long as floating-point exception traps are disabled through the tem field. See FIGURE 5-7 on page 47.

After an FPop completes with ftt = 0, the tem and cexc fields are logically **anded** together. If the result is nonzero, aexc is left unchanged and an *fp_exception_ieee_754* trap is generated; otherwise, the new cexc field is **ored** into the aexc field and no trap is generated. Thus, while (and only while) traps are masked, exceptions are accumulated in the aexc field.

FSR.aexc can be set to a specific value when an LDFSR, LDXEFSR, or LDXFSR instruction is executed.

5.4.8 Current Exception (cexc)

FSR.cexc (FSR{4:0}) indicates whether one or more IEEE 754 floating-point exceptions were generated by the most recently executed FPop instruction. The absence of an exception causes the corresponding bit to be cleared (set to 0). See FIGURE 5-6 on page 47.

Programming Note	If the FPop traps and software emulate or finish the instruction, the system software in the trap handler is responsible for creating a correct FSR.cexc value before returning to a nonprivileged program.
-------------------------	---

The cexc bits are set as described in *Floating-Point Exception Fields* on page 47, by the execution of an FPop that either does not cause a trap or causes an *fp_exception_ieee_754* exception with FSR.ftt = IEEE_754_exception. An IEEE 754 exception that traps shall cause exactly one bit in FSR.cexc to be set, corresponding to the detected IEEE Std 754-1985 exception.

Floating-point operations which cause an overflow or underflow condition may also cause an “inexact” condition. For overflow and underflow conditions, FSR.cexc bits are set and trapping occurs as follows:

- If an IEEE 754 overflow condition occurs:
 - if FSR.tem.ofm = 0 and tem.nxm = 0, the FSR.cexc.ofc and FSR.cexc.nxc bits are both set to 1, the other three bits of FSR.cexc are set to 0, and an *fp_exception_ieee_754* trap does *not* occur.
 - if FSR.tem.ofm = 0 and tem.nxm = 1, the FSR.cexc.nxc bit is set to 1, the other four bits of FSR.cexc are set to 0, and an *fp_exception_ieee_754* trap *does* occur.
 - if FSR.tem.ofm = 1, the FSR.cexc.ofc bit is set to 1, the other four bits of FSR.cexc are set to 0, and an *fp_exception_ieee_754* trap *does* occur.
- If an IEEE 754 underflow condition occurs:
 - if FSR.tem.ufm = 0 and FSR.tem.nxm = 0, the FSR.cexc.ufc and FSR.cexc.nxc bits are both set to 1, the other three bits of FSR.cexc are set to 0, and an *fp_exception_ieee_754* trap does *not* occur.
 - if FSR.tem.ufm = 0 and FSR.tem.nxm = 1, the FSR.cexc.nxc bit is set to 1, the other four bits of FSR.cexc are set to 0, and an *fp_exception_ieee_754* trap *does* occur.

- if FSR.tem.ufm = 1, the FSR.cexc.ufc bit is set to 1, the other four bits of FSR.cexc are set to 0, and an *fp_exception_ieee_754* trap *does* occur.

The above behavior is summarized in TABLE 5-7 (where “✓” indicates “exception was detected” and “x” indicates “don’t care”):

TABLE 5-7 Setting of FSR.cexc Bits

Conditions						Results			
Exception(s) Detected in F.p. operation			Trap Enable Mask bits (in FSR.tem)			<i>fp_exception_ieee_754</i> Trap Occurs?	Current Exception bits (in FSR.cexc)		
of	uf	nx	ofm	ufm	nxm		ofc	ufc	nxc
-	-	-	x	x	x	no	0	0	0
-	-	✓	x	x	0	no	0	0	1
-	✓ ¹	✓ ¹	x	0	0	no	0	1	1
✓ ²	-	✓ ²	0	x	0	no	1	0	1
-	-	✓	x	x	1	yes	0	0	1
-	✓ ¹	✓ ¹	x	0	1	yes	0	0	1
-	✓	-	x	1	x	yes	0	1	0
-	✓	✓	x	1	x	yes	0	1	0
✓ ²	-	✓ ²	1	x	x	yes	1	0	0
✓ ²	-	✓ ²	0	x	1	yes	0	0	1

Notes: ¹ When the underflow trap is disabled (FSR.tem.ufm = 0) underflow is always accompanied by inexact.

² Overflow is always accompanied by inexact.

If the execution of an FPop causes a trap other than *fp_exception_ieee_754*, FSR.cexc is left unchanged.

5.4.9 Floating-Point Exception Fields

The current and accrued exception fields and the trap enable mask assume the following definitions of the floating-point exception conditions (per IEEE Std 754-1985):




FIGURE 5-6 Trap Enable Mask (tem) Fields of FSR




FIGURE 5-7 Accrued Exception Bits (aexc) Fields of FSR




FIGURE 5-8 Current Exception Bits (aexc) Fields of FSR

Invalid (nvc, nva). An operand is improper for the operation to be performed. For example, $0.0 \div 0.0$ and $\infty - \infty$ are invalid; 1 = invalid operand(s), 0 = valid operand(s).

Overflow (ofc, ofa). The result, rounded as if the exponent range were unbounded, would be larger in magnitude than the destination format's largest finite number; 1 = overflow, 0 = no overflow.

Underflow (ufc, ufa). The rounded result is inexact and would be smaller in magnitude than the smallest normalized number in the indicated format; 1 = underflow, 0 = no underflow.

Underflow is never indicated when the correct unrounded result is 0. Otherwise, when the correct unrounded result is not 0:

If `FSR.tem.ufm = 0`: Underflow occurs if a nonzero result is tiny and a loss of accuracy occurs.

If `FSR.tem.ufm = 1`: Underflow occurs if a nonzero result is tiny.

The SPARC V9 architecture allows tininess to be detected either before or after rounding. However, in all cases and regardless of the setting of `FSR.tem.ufm`, an Oracle SPARC Architecture strand detects tininess before rounding (impl. dep. #55-V8-Cs10). See *Trapped Underflow Definition (ufm = 1)* on page 393 and *Untrapped Underflow Definition (ufm = 0)* on page 393 for additional details.

Division by zero (dzc, dza). An infinite result is produced exactly from finite operands. For example, $X \div 0.0$, where X is subnormal or normalized; 1 = division by zero, 0 = no division by zero.

Inexact (nxc, nxa). The rounded result of an operation differs from the infinitely precise unrounded result; 1 = inexact result, 0 = exact result.

5.4.10 FSR Conformance

An Oracle SPARC Architecture implementation implements the `tem`, `cexc`, and `aexc` fields of FSR in hardware, conforming to IEEE Std 754-1985 (impl. dep. #22-V8).

Programming Note	Privileged software (or a combination of privileged and nonprivileged software) must be capable of simulating the operation of the FPU in order to handle the <i>fp_exception_other</i> (with <code>FSR.ftt = unfinished_FPop</code>) and <i>IEEE_754_exception</i> floating-point trap types properly. Thus, a user application program always sees an FSR that is fully compliant with IEEE Std 754-1985.
-------------------------	--

5.5 Ancillary State Registers

The SPARC V9 architecture defines several optional ancillary state registers (ASRs) and allows for additional ones. Access to a particular ASR may be privileged or nonprivileged.

An ASR is read and written with the Read State Register and Write State Register instructions, respectively. These instructions are privileged if the accessed register is privileged.

The SPARC V9 architecture left ASRs numbered 16–31 available for implementation-dependent uses. Oracle SPARC Architecture virtual processors implement the ASRs summarized in TABLE 5-8 and defined in the following subsections.

Each virtual processor contains its own set of ASRs; ASRs are not shared among virtual processors.

TABLE 5-8 ASR Register Summary

ASR number	ASR name	Register	Read by Instruction(s)	Written by Instruction(s)
0	Y ^D	Y register (deprecated)	RDY ^D	WRY ^D
1	—	<i>Reserved</i>	—	—
2	CCR	Condition Codes register	RDCCR	WRCCR
3	ASI	ASI register	RDASI	WRASI
4	TICK ^{P_{dis},H_{dis}}	TICK register	RDTICK ^{P_{dis},H_{dis}} , RDPR ^P (TICK)	WRPR ^P (TICK)
5	PC	Program Counter (PC)	RDPC	(all instructions)
6	FPRS	Floating-Point Registers Status register	RDFPRS	WRFPRS
7–12 (7-0C ₁₆)	—	<i>Reserved</i>	—	—
13 (0D ₁₆)	—	<i>Reserved</i>	—	—
14 (0E ₁₆)	—	<i>Reserved</i>	—	—
15 (0F ₁₆)	—	<i>Reserved</i> (used for MEMBAR; see text under MEMBAR [p. 270] or RDasr [p. 312] instructions)	—	—
16 (10 ₁₆)	—	<i>Reserved</i>	—	—
17 (11 ₁₆)	—	<i>Reserved</i>	—	—
16-18 (10 ₁₆ - 12 ₁₆)	—	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)	—	—
19 (13 ₁₆)	GSR	General Status register (GSR)	RDGSR, FALIGNDATA _g , many VIS and floating-point instructions	WRGSR, BMASK, SIAM
20 (14 ₁₆)	SOFTINT_SET ^P	(pseudo-register, for "Write 1s Set" to SOFTINT register, ASR 22)	—	WRSOFTINT_SET ^P
21 (15 ₁₆)	SOFTINT_CLR ^P	(pseudo-register, for "Write 1s Clear" to SOFTINT register, ASR 22)	—	WRSOFTINT_CLR ^P
22 (16 ₁₆)	SOFTINT ^P	per-virtual processor Soft Interrupt register	RDSOFTINT ^P	WRSOFTINT ^P
23 (17 ₁₆)	—	<i>Reserved</i>	—	—
24 (18 ₁₆)	STICK ^{P_{dis},H_{dis}}	System Tick register	RDSTICK ^{P_{dis},H_{dis}}	—
25 (19 ₁₆)	STICK_CMPR ^P	System Tick Compare register	RDSTICK_CMPR ^P	WRSTICK_CMPR ^P
26 (1A ₁₆)	CFR	Compatibility Feature register	RDCFR	—
27 (1B ₁₆)	PAUSE	Pause Count register	—	PAUSE (WRAsr 27)
28 (1C ₁₆)	MWAIT	MWait Count register	—	MWAIT (WRAsr 28)
29 (1D ₁₆)	—	<i>Reserved</i>	—	—
30 (1E ₁₆)	—	<i>Reserved</i>	—	—
31 (1F ₁₆)	—	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)	—	—
31 (1F ₁₆)	—	<i>Reserved</i>	—	—

5.5.1 32-bit Multiply/Divide Register (Y) (ASR 0) (D3)

The Y register is deprecated; it is provided only for compatibility with previous versions of the architecture. It should not be used in new SPARC V9 software. It is recommended that all instructions that reference the Y register (that is, SMUL, SMULcc, UMUL, UMULcc, SDIV, SDIVcc, UDIV, UDIVcc, RDY, and WRY) be avoided. For suitable substitute instructions, see the following pages: for the multiply instructions, see pages 333 and page 371; for division instructions, see pages 325 and 369; for the read instruction, see page 311; and for the write instruction, see page 374.

The low-order 32 bits of the Y register, illustrated in FIGURE 5-9, contain the more significant word of the 64-bit product of an integer multiplication, as a result of a 32-bit integer multiply (SMUL, SMULcc, UMUL, UMULcc) instruction. The Y register also holds the more significant word of the 64-bit dividend for a 32-bit integer divide (SDIV, SDIVcc, UDIV, UDIVcc) instruction.




FIGURE 5-9 Y Register

Although Y is a 64-bit register, its high-order 32 bits always read as 0.

The Y register may be explicitly read and written by the RDY and WRY instructions, respectively.

5.5.2 Integer Condition Codes Register (CCR) (ASR 2) (A1)

The Condition Codes Register (CCR), shown in FIGURE 5-10, contains the integer condition codes. The CCR register may be explicitly read and written by the RDCCR and WRCCR instructions, respectively.




FIGURE 5-10 Condition Codes Register

5.5.2.1 Condition Codes (CCR.xcc and CCR.icc)

All instructions that set integer condition codes set both the xcc and icc fields. The xcc condition codes indicate the result of an operation when viewed as a 64-bit operation. The icc condition codes indicate the result of an operation when viewed as a 32-bit operation. For example, if an operation results in the 64-bit value 0000 0000 FFFF FFFF₁₆, the 32-bit result is negative (icc.n is set to 1) but the 64-bit result is nonnegative (xcc.n is set to 0).

Each of the 4-bit condition-code fields is composed of four 1-bit subfields, as shown in FIGURE 5-11.




FIGURE 5-11 Integer Condition Codes (CCR.icc and CCR.xcc)

The n bits indicate whether the two's-complement ALU result was negative for the last instruction that modified the integer condition codes; 1 = negative, 0 = not negative.

The z bits indicate whether the ALU result was zero for the last instruction that modified the integer condition codes; 1 = zero, 0 = nonzero.

The *v* bits signify whether the ALU result was within the range of (was representable in) 64-bit (*xcc*) or 32-bit (*icc*) 2’s-complement notation for the last instruction that modified the integer condition codes; 1 = overflow, 0 = no overflow. The *v* bits may be used to test for signed overflow.

The *c* bits indicate whether a 2’s-complement carry (or borrow) occurred during the last instruction that modified the integer condition codes. Carry is set on addition if there is a carry out of bit 63 (*xcc*) or bit 31 (*icc*). Carry is set on subtraction if there is a borrow into bit 63 (*xcc*) or bit 31 (*icc*); 1 = borrow, 0 = no borrow (see TABLE 5-9). The *c* bits may be used to test for unsigned overflow.

TABLE 5-9 Setting of Carry (Borrow) bits for Subtraction That Sets CCs

Unsigned Comparison of Operand Values	Setting of Carry bits in CCR
$R[rs1]\{31:0\} \geq R[rs2]\{31:0\}$	CCR.icc.c ← 0
$R[rs1]\{31:0\} < R[rs2]\{31:0\}$	CCR.icc.c ← 1
$R[rs1]\{63:0\} \geq R[rs2]\{63:0\}$	CCR.xcc.c ← 0
$R[rs1]\{63:0\} < R[rs2]\{63:0\}$	CCR.xcc.c ← 1

Both fields of CCR (*xcc* and *icc*) are modified by arithmetic and logical instructions, the names of which end with the letters “cc” (for example, ANDcc), and by the WRCCR instruction. They can also be modified by a DONE or RETRY instruction, which replaces these bits with the contents of TSTATE.ccr. The behavior of the following instructions are conditioned by the contents of CCR.icc or CCR.xcc:

- BPcc and Tcc instructions (conditional transfer of control)
- Bicc (conditional transfer of control, based on CCR.icc only)
- MOVcc instruction (conditionally move the contents of an integer register)
- FMOVcc instruction (conditionally move the contents of a floating-point register)

Extended (64-bit) integer condition codes (*xcc*). Bits 7 through 4 are the IU condition codes, which indicate the results of an integer operation, with both of the operands and the result considered to be 64 bits wide.

32-bit Integer condition codes (*icc*). Bits 3 through 0 are the IU condition codes, which indicate the results of an integer operation, with both of the operands and the result considered to be 32 bits wide.

5.5.3 Address Space Identifier (ASI) Register (ASR 3) A1

The Address Space Identifier register (FIGURE 5-12) specifies the address space identifier to be used for load and store alternate instructions that use the “rs1 + simm13” addressing form.

The ASI register may be explicitly read and written by the RDASI and WRASI instructions, respectively.

Software (executing in any privilege mode) may write any value into the ASI register. However, values in the range 00₁₆ to 7F₁₆ are “restricted” ASIs; an attempt to perform an access using an ASI in that range is restricted to software executing in a mode with sufficient privileges for the ASI. When an instruction executing in nonprivileged mode attempts an access using an ASI in the range 00₁₆ to 7F₁₆ or an instruction executing in privileged mode attempts an access using an ASI the range 30₁₆ to 7F₁₆, a *privileged_action* exception is generated. See Chapter 10, *Address Space Identifiers (ASIs)* for details.




FIGURE 5-12 Address Space Identifier Register

5.5.4 Tick (TICK) Register (ASR 4) (A1)

FIGURE 5-13 illustrates the TICK register.




FIGURE 5-13 TICK Register

The counter field of the TICK register is a 63-bit counter that counts virtual processor clock cycles.

Bit 63 of the TICK register reads as 0.

Privileged software can read the TICK register with either the RDPR or RDTICK instruction, but only when privileged access to TICK is enabled by hyperprivileged software. An attempt by privileged software to read the TICK register when privileged access is disabled causes a *privileged_action* exception.

Privileged software cannot write to the TICK register; an attempt to do so (with the WRPR instruction) results in an *illegal_instruction* exception.

Nonprivileged software can read the TICK register by using the RDTICK instruction, but only when nonprivileged access to TICK is enabled by hyperprivileged software. If nonprivileged access is disabled, an attempt by nonprivileged software to read the TICK register using the RDTICK instruction causes a *privileged_action* exception.

An attempt by nonprivileged software at any time to read the TICK register using the privileged RDPR instruction causes a *privileged_opcode* exception.

Nonprivileged software cannot write the TICK register. An attempt by nonprivileged software to write the TICK register using the privileged WRPR instruction causes a *privileged_opcode* exception.

The foregoing description regarding access to the TICK register is summarized in TABLE 5-10 and TABLE 5-11.

TABLE 5-10 Access to TICK register through ASR 4 (RDTICK instruction)

Privilege Mode	Read Access Using RDTICK instruction
Nonprivileged	if allowed by hyperprivileged software, TICK may be read with RDTICK otherwise, <i>privileged_action</i> exception
Privileged	if allowed by hyperprivileged software, TICK may be read with RDTICK otherwise <i>privileged_action</i> exception

TABLE 5-11 Access to TICK register through 105-V9h PR 4 (RDPR and WRPR instructions)

Privilege Mode	Read Access Using RDPR TICK instruction	Write Access Using WRPR TICK instruction
Nonprivileged	<i>privileged_opcode</i> exception	<i>privileged_opcode</i> exception
Privileged	if allowed by hyperprivileged software, TICK may be read with RDPR TICK otherwise <i>privileged_action</i> exception	<i>illegal_instruction</i> exception

The difference between the values read from the TICK register on two reads is intended to reflect the number of strand cycles executed between the reads.

Programming Note If a single TICK register is shared among multiple virtual processors, then the difference between subsequent reads of TICK.counter reflects a shared cycle count, not a count specific to the virtual processor reading the TICK register.

IMPL. DEP. #105-V9: (a) If an accurate count cannot always be returned when TICK is read, any inaccuracy should be small, bounded, and documented.
 (b) An implementation may implement fewer than 63 bits in TICK.counter; however, the counter as implemented must be able to count for at least 10 years without overflowing. Any upper bits not implemented must read as zero.

5.5.5 Program Counters (PC, NPC) (ASR 5) (A1)

The PC contains the address of the instruction currently being executed. The least-significant two bits of PC always contain zeroes.

The PC can be read directly with the RDPC instruction. PC cannot be explicitly written by any instruction (including Write State Register), but is implicitly written by control transfer instructions. A WRAsr to ASR 5 causes an *illegal_instruction* exception.

The Next Program Counter, NPC, is a pseudo-register that contains the address of the next instruction to be executed if a trap does not occur. The least-significant two bits of NPC always contain zeroes.

NPC is written implicitly by control transfer instructions. However, NPC cannot be read or written explicitly by any instruction.

PC and NPC can be indirectly set by privileged software that writes to TPC[TL] and/or TNPC[TL] and executes a RETRY instruction.

See Chapter 6, *Instruction Set Overview*, for details on how PC and NPC are used.

5.5.6 Floating-Point Registers State (FPRS) Register (ASR 6) (A1)

The Floating-Point Registers State (FPRS) register, shown in FIGURE 5-14, contains control information for the floating-point register file; this information is readable and writable by nonprivileged software.




FIGURE 5-14 Floating-Point Registers State Register

The FPRS register may be explicitly read and written by the RDFPRS and WRFPRS instructions, respectively.

Enable FPU (fef). Bit 2, *fef*, determines whether the FPU is enabled. If it is disabled, executing a floating-point instruction causes an *fp_disabled* trap. If this bit is set (FPRS.fef = 1) but the PSTATE.pef bit is not set (PSTATE.pef = 0), then executing a floating-point instruction causes an *fp_disabled* exception; that is, both FPRS.fef and PSTATE.pef must be set to 1 to enable floating-point operations.

Programming Note FPRS.fef can be used by application software to notify system software that the application does not require the contents of the F registers to be preserved. Depending on system software, this may provide some performance benefit, for example, the F registers would not have to be saved or restored during context switches to or from that application. Once an application sets FPRS.fef to 0, it must assume that the values in all F registers are volatile (may change at any time).

Dirty Upper Registers (du). Bit 1 is the “dirty” bit for the upper half of the floating-point registers; that is, F[32]–F[62]. It is set to 1 whenever any of the upper floating-point registers is modified. The du bit is cleared only by software.

An Oracle SPARC Architecture 2015 virtual processor may set FPRS.du pessimistically; that is, it may be set whenever an FPop executes, even though an exception may occur that prevents the instruction from completing so no destination F register was actually modified (impl. dep. #403-S10). Note that if the FPop triggers *fp_disabled*, FPRS.du is *not* modified.

Dirty Lower Registers (dl). Bit 0 is the “dirty” bit for the lower 32 floating-point registers; that is, F[0]–F[31]. It is set to 1 whenever any of the lower floating-point registers is modified. The dl bit is cleared only by software.

An Oracle SPARC Architecture 2015 virtual processor may set FPRS.dl pessimistically; that is, it may be set whenever an FPop executes, even though an exception may occur that prevents the instruction from completing so no destination F register was actually modified (impl. dep. #403-S10). Note that if the FPop triggers *fp_disabled*, FPRS.dl is *not* modified.

5.5.7 General Status Register (GSR) (ASR 19) (A1)

The General Status Register¹ (GSR) is a nonprivileged read/write register that is implicitly referenced by many VIS instructions. The GSR can be read by the RDGSR instruction (see *Read Ancillary State Register* on page 310) and written by the WRGSR instruction (see *Write Ancillary State Register* on page 373).

If the FPU is disabled (PSTATE.pef = 0 or FPRS.fef = 0), an attempt to access this register using an otherwise-valid RDGSR or WRGSR instruction causes an *fp_disabled* trap.

The GSR is illustrated in FIGURE 5-15 and described in TABLE 5-12.




FIGURE 5-15 General Status Register (GSR) (ASR 19)

TABLE 5-12 GSR Bit Description

Bit	Field	Description
63:32	mask	This 32-bit field specifies the mask used by the BSHUFFLE instruction. The field contents are set by the BMASK instruction.
31:28	—	<i>Reserved.</i>
27	im	Interval Mode: If GSR.im = 0, rounding is performed according to FSR.rd; if GSR.im = 1, rounding is performed according to GSR.irnd.
26:25	irnd	IEEE Std 754-1985 rounding direction to use in Interval Mode (GSR.im = 1), as follows:
	irnd	Round toward ...
	0	Nearest (even, if tie)
	1	0
	2	+ ∞
	3	- ∞
24:8	—	<i>Reserved.</i>
7:3	scale	5-bit shift count in the range 0–31, used by the FPACK instructions for formatting.
2:0	align	Least three significant bits of the address computed by the last-executed ALIGNADDRESS or ALIGNADDRESS_LITTLE instruction.

¹ This register was (inaccurately) referred to as the “Graphics Status Register” in early UltraSPARC implementations

5.5.8 SOFTINT^P Register (ASRs 20 (A2), 21 (A2), 22 (A1))

Software uses the privileged, read/write SOFTINT register (ASR 22) to schedule interrupts (via *interrupt_level_n* exceptions).

SOFTINT (A1) can be read with a RDSOFTINT instruction (see *Read Ancillary State Register* on page 310) and written with a WRSOFTINT, WRSOFTINT_SET, or WRSOFTINT_CLR instruction (see *Write Ancillary State Register* on page 373). An attempt to access to this register in nonprivileged mode causes a *privileged_opcode* exception.

Programming Note | To atomically modify the set of pending software interrupts, use of the SOFTINT_SET and SOFTINT_CLR ASRs is recommended.

The SOFTINT register is illustrated in FIGURE 5-16 and described in TABLE 5-13.




FIGURE 5-16 SOFTINT Register (ASR 22)

TABLE 5-13 SOFTINT Bit Description

Bit	Field	Description
16	sm	When the STICK_CMPR (ASR 25) register's <i>int_dis</i> (interrupt disable) field is 0 (that is, System Tick Compare is enabled) and its <i>stick_cmpr</i> field matches the value in the STICK register, then SOFTINT.sm ("STICK match") is set to 1 and a level 14 interrupt (<i>interrupt_level_14</i>) is generated. See <i>System Tick Compare (stick_cmprP) Register (ASR 25)</i> on page 57 for details. SOFTINT.sm can also be directly written to 1 by software.
15:1	int_level	When SOFTINT.int_level{n-1} (SOFTINT{n}) is set to 1, an <i>interrupt_level_n</i> exception is generated.
<p>Notes: A level-14 interrupt (<i>interrupt_level_14</i>) can be triggered by SOFTINT.sm or a write to SOFTINT.int_level{13} (SOFTINT{14}).</p> <p>A level-15 interrupt (<i>interrupt_level_15</i>) can be triggered by a write to SOFTINT.int_level{14} (SOFTINT{15}), or possibly by other implementation-dependent mechanisms.</p> <p>An <i>interrupt_level_n</i> exception will only cause a trap if (PIL < n) and (PSTATE.ie = 1).</p>		
0	—	Reserved.

Setting either SOFTINT.sm or SOFTINT.int_level{13} (SOFTINT{14}) to 1 causes a level-14 interrupt (*interrupt_level_14*). However, those two bits are independent; setting one of them does not affect the other.

See *Software Interrupt Register (softint)* on page 468 for additional information regarding the SOFTINT register.

5.5.8.1 SOFTINT_SET^P Pseudo-Register (ASR 20) (A2)

A Write State register instruction to ASR 20 (WRSOFTINT_SET) atomically sets selected bits in the privileged SOFTINT Register (ASR 22) (see page 55). That is, bits 16:0 of the write data are **ored** into SOFTINT; any '1' bit in the write data causes the corresponding bit of SOFTINT to be set to 1. Bits 63:17 of the write data are ignored.

Access to ASR 20 is privileged and write-only. There is no instruction to read this pseudo-register. An attempt to write to ASR 20 in non-privileged mode, using the WRasr instruction, causes a *privileged_opcode* exception.

Programming Note | There is no actual “register” (machine state) corresponding to ASR 20; it is just a programming interface to conveniently set selected bits to ‘1’ in the SOFTINT register, ASR 22.

FIGURE 5-17 illustrates the SOFTINT_SET pseudo-register.




FIGURE 5-17 SOFTINT_SET Pseudo-Register (ASR 20)

5.5.8.2 SOFTINT_CLR^P Pseudo-Register (ASR 21) (A2)

A Write State register instruction to ASR 21 (WRSOFTINT_CLR) atomically clears selected bits in the privileged SOFTINT register (ASR 22) (see page 55). That is, bits 16:0 of the write data are inverted and **anded** into SOFTINT; any ‘1’ bit in the write data causes the corresponding bit of SOFTINT to be set to 0. Bits 63:17 of the write data are ignored.

Access to ASR 21 is privileged and write-only. There is no instruction to read this pseudo-register. An attempt to write to ASR 21 in non-privileged mode, using the WRasr instruction, causes a *privileged_opcode* exception.

Programming Note | There is no actual “register” (machine state) corresponding to ASR 21; it is just a programming interface to conveniently clear (set to ‘0’) selected bits in the SOFTINT register, ASR 22.

FIGURE 5-18 illustrates the SOFTINT_CLR pseudo-register.




FIGURE 5-18 SOFTINT_CLR Pseudo-Register (ASR 21))

5.5.9 System Tick (STICK) Register (ASR 24) (A1)

The System Tick (STICK) register provides a counter that consistently measures time across all virtual processors (strands) of a system. The 63-bit counter field of the STICK register automatically increments at a fixed frequency of 1.0 GHz, therefore counter bit n will be observed to increment at a frequency of $1.0 \text{ GHz} \div 2^n$.

The STICK register is illustrated in FIGURE 5-19 and described below.




FIGURE 5-19 STICK Register

Bit 63 of the STICK register reads as 0.

The counter field spans bits 62: m of the STICK register. m is implementation-dependent (impl.dep. #442-S10(c)), but m must be less than or equal to 4 (STICK granularity of 16ns or better); m can be 0.

The counter must increment 1 billion \pm 25,000 times every second (an error rate of no more than 25 parts per million). This means that the counter must not gain or lose more than 2.16 seconds per day.

IMPL. DEP. #442-S10: (a) If an accurate count cannot always be returned when STICK is read, any inaccuracy should be small, bounded, and documented.

(b) *(no longer applies)*

(c) The bit number m of the least significant implemented bit of the counter field of the STICK register is implementation-dependent, but must be in range 4 to 0, inclusively (that is, $4 \geq m \geq 0$). Any low-order bits not implemented must read as zero.

Writes to unimplemented bits of STICK are ignored. Any implementation that does not implement all bits of STICK must ensure that comparisons to STICK_CMPR correctly account for the absence of the unimplemented bits.

At least one STICK register must be implemented per system. However, multiple STICK registers per system may be implemented (for example, STICK may be implemented per-core or per-strand). No more than one STICK may be implemented per strand. An implementation must document how many STICKs are implemented per system and their relationships to strands and other structures.

Privileged software can read the STICK register with the RDSTICK instruction, but only when privileged access to STICK is enabled by hyperprivileged software. An attempt by privileged software to read the STICK register when privileged access is disabled causes a *privileged_action* exception.

Privileged software cannot write the STICK register; an attempt to execute the WRSTICK instruction in privileged mode results in an *illegal_instruction* exception.

Nonprivileged software can read the STICK register using the RDSTICK instruction, but only when nonprivileged access to STICK is enabled by hyperprivileged software. If nonprivileged access is disabled, an attempt by nonprivileged software to read the STICK register causes a *privileged_action* exception.

Nonprivileged software cannot write the STICK register; an attempt to execute the WRSTICK instruction in nonprivileged mode results in an *illegal_instruction* exception.

The difference between two values read from the STICK register at different times reflects the amount of time that has passed between the reads; $(value1 - value0)$ yields the number of nanoseconds that elapsed between the two reads.

Note | If STICK begins counting at 0, it will overflow in approximately 34 years.

5.5.10 System Tick Compare (STICK_CMPR^P) Register (ASR 25) (A2)

The privileged STICK_CMPR register allows system software to cause an *interrupt_level_14* trap when the STICK register reaches the value specified in STICK_CMPR. Nonprivileged accesses to this register cause a *privileged_opcode* exception (see *Exception and Interrupt Descriptions* on page 460).

The System Tick Compare Register is illustrated in FIGURE 5-20 and described in TABLE 5-14.




FIGURE 5-20 STICK_CMPR Register

The value of n (as shown in FIGURE 5-20) is implementation-dependent, but must be between 0 and 8, inclusively ($0 \leq n \leq 8$).

The value of n (as shown in FIGURE 5-20) is implementation-dependent, but must be between 0 and 9, inclusively ($0 \leq n \leq 9$).

TABLE 5-14 STICK_CMPR Register Description

Bits	Field	Description
63	int_dis	Interrupt Disable. If set to 1, STICK_CMPR interrupts are disabled.
62: n	stick_cmpr	System Tick Compare Field.
$(n-1):0$	Reserved	

A *stick_match* exception occurs when all three of the following conditions are met:

- STICK_CMPR.int_dis = 0
- a transition occurs from
 - $(\text{STICK.counter} \ll m) < (\text{STICK_CMPR.stick_cmpr} \ll n)$
 - in one cycle to
 - $(\text{STICK.counter} \ll m) \geq (\text{STICK_CMPR.stick_cmpr} \ll n)$
 - in the following cycle, where:
 - m is the bit number of the least-significant implemented bit of STICK.counter (see page 56)
 - n is the bit number of the least-significant implemented bit of STICK_CMPR.stick_cmpr
 - \ll is the arithmetic left-shift operator
 - $<$ and \geq are unsigned comparisons
- the above state transition occurred due to incrementing STICK and **not** due to an explicit write to STICK or STICK_CMPR

When a *stick_match* exception occurs, SOFTINT{16} (SOFTINT.sm) is set to 1. This has the effect of posting an *interrupt_level_14* trap request to the virtual processor, which causes an *interrupt_level_14* trap when (PIL < 14) **and** (PSTATE.ie = 1). The *interrupt_level_14* trap handler must check SOFTINT{14} and SOFTINT{16} (SOFTINT.sm) to determine the cause of the *interrupt_level_14* trap.

An implementation may compare (STICK_CMPR.stick_cmpr $\ll n$) to STICK periodically, instead of comparing every cycle.

The implementation must ensure that if STICK_CMPR.int_dis = 0 and a 63-bit value written to STICK_CMPR.stick_cmpr is greater than STICK, that the implementation will detect a *stick_match* exception at a future time.

Implementation Note	<p>If an implementation does not compare against STICK_CMPR.stick_cmpr every cycle, it is recommended that the implementation not implement the less-significant bits of STICK_CMPR.stick_cmpr that correspond to the time interval between comparisons of STICK_CMPR.stick_cmpr.</p> <p>In addition, if a write to STICK_CMPR.stick_cmpr has nonzero unimplemented lower bits and if the 63-bit value written to STICK_CMPR.stick_cmpr is greater than STICK, then the implementation must ensure that the <i>stick_match</i> exception is not lost. This may be accomplished by incrementing the value of the implemented bits by an amount sufficient to ensure that a <i>stick_match</i> exception will occur at a future time, or by posting the <i>stick_match</i> exception to take effect at some later time, or by other means.</p> <p>Any alterations performed by virtual processor hardware of the values written to STICK_CMPR.stick_cmpr must be specified in that processor's Implementation Supplement to this document.</p>
----------------------------	--

5.5.11 Compatibility Feature Register (CFR) (ASR 26)

Each virtual processor has a Compatibility Feature Register (CFR). The CFR is read-only.

The format of the CFR is shown in TABLE 5-15.

TABLE 5-15 Compatibility Feature Register (CFR) Description (ASR 1A₁₆)

Bit	Field	Description
63:22	—	<i>Reserved</i>
21	—	<i>Reserved</i>
20:15	—	<i>Reserved</i>
14	xmongsqr	If set, the processor supports the XMONGSQR opcode. If not set, an attempt to execute an XMONGSQR instruction results in a <i>compatibility_feature</i> trap.
13	xmontmul	If set, the processor supports the XMONTMUL opcode. If not set, an attempt to execute an XMONTMUL instruction results in a <i>compatibility_feature</i> trap.
12	xmpmul	If set, the processor supports the XMPMUL opcode. If not set, an attempt to execute an XMPMUL instruction results in a <i>compatibility_feature</i> trap.
11	crc32c	If set, the processor supports the CRC32C opcode. If not set, an attempt to execute a CRC32C instruction results in a <i>compatibility_feature</i> trap.
10	monsqr	If set, the processor supports the MONTSQR opcode. If not set, an attempt to execute a MONTSQR instruction results in a <i>compatibility_feature</i> trap.
9	montmul	If set, the processor supports the MONTMUL opcode. If not set, an attempt to execute a MONTMUL instruction results in a <i>compatibility_feature</i> trap.
8	mpmul	If set, the processor supports the MPMUL opcode. If not set, an attempt to execute an MPMUL instruction results in a <i>compatibility_feature</i> trap.
7	sha512	If set, the processor supports the SHA512 opcode. If not set, an attempt to execute a SHA512 instruction results in a <i>compatibility_feature</i> trap.
6	sha256	If set, the processor supports the SHA256 opcode. If not set, an attempt to execute a SHA256 instruction results in a <i>compatibility_feature</i> trap.
5	sha1	If set, the processor supports the SHA1 opcode. If not set, an attempt to execute a SHA1 instruction results in a <i>compatibility_feature</i> trap.
4	md5	If set, the processor supports the MD5 opcode. If not set, an attempt to execute an MD5 instruction results in a <i>compatibility_feature</i> trap.
3	camellia	If set, the processor supports Camellia opcodes (CAMELLIA_F, CAMELLIA_FL, and CAMELLIA_FLI). If not set, an attempt to execute a Camellia instruction results in a <i>compatibility_feature</i> trap.
2	—	0 —
1	des	If set, the processor supports DES opcodes (DES_ROUND, DES_IP, DES_IIP, and DES_KEXPAND). If not set, an attempt to execute a DES instruction results in a <i>compatibility_feature</i> trap.
0	aes	If set, the processor supports AES opcodes (AES_EROUND01, AES_EROUND23, AES_DROUND01, AES_DROUND23, AES_EROUND_01_LAST, AES_EROUND_23_LAST, AES_DROUND_01_LAST, AES_DROUND_23_LAST, AES_KEXPAND0, AES_KEXPAND1, and AES_KEXPAND2). If not set, an attempt to execute an AES instruction results in a <i>compatibility_feature</i> trap.

The CFR enumerates the capabilities that the virtual processor supports. Software can use the CFR to determine whether particular features (such as instruction opcodes or registers) are available for use on the current virtual processor. If a virtual processor executes an opcode associated with a bit in CFR but that bit is 0, that feature is not present and a *compatibility_feature* trap occurs.

Programming Note For correctness and optimal performance, prior to using any feature associated with a bit in CFR, an application or library must first check the CFR to ensure that the desired feature is actually supported by the underlying hardware.

5.5.12 Pause Count (PAUSE) Register (ASR 27) C1

The nonprivileged PAUSE register provides a means for software to voluntarily and temporarily pause execution, freeing up processor resources for use by other threads of execution. When PAUSE = 0, it has no effect on execution. When PAUSE ≠ 0, execution is paused (temporarily suspended) until PAUSE = 0 again.

Software initiates a pause for n nanoseconds by writing the value n to the PAUSE register, using the WRPAUSE instruction. PAUSE is write-only and cannot be read by software; in particular, there is no RDPAUSE instruction.

When PAUSE ≠ 0, it is automatically decremented once every nanosecond. When PAUSE = 0, no decrementation occurs.




FIGURE 5-21 PAUSE Register

IMPL. DEP. #502-S30: The maximum number of nanoseconds that a virtual processor may be paused by the PAUSE instruction, MAX_PAUSE_COUNT , is $2^B - 1$. B , the number of bits used to implement the `pause_count` field, is implementation-dependent but B must be ≥ 13 . Therefore $MAX_PAUSE_COUNT \geq 2^{13} - 1$ (8191) nanoseconds. If a PAUSE instruction requests a pause of more than MAX_PAUSE_COUNT nanoseconds, the virtual processor must saturate the pause count to its maximum value, that is, write a value of MAX_PAUSE_COUNT to the PAUSE register.

Implementation Note | An implementation should implement enough bits in the PAUSE register to satisfy the equation:

$$\frac{(C \times 2^B) - (S \times 4)}{C \times 2^B} \geq 0.99$$

where:

B is the number of bits that must be implemented in the PAUSE register
 C is the minimum number of clock cycles per nanosecond (under any conditions except *Parke*d), and
 S is the number of strands that share an instruction pipeline, and
 4 represents the number of slots in the pipeline needed to execute the four instructions in a minimal busy-wait loop (load, compare, branch, PAUSE).

Solving this equation for B yields:

$$B \geq \text{ceiling}(\log_2(100 \times S \div C)) + 2$$

Implementation Note | While PAUSE ≠ 0, the strand should suspend forward progress and should release all shared hardware resources to other strands.

When the value of PAUSE reaches 0, a paused strand resumes forward progress and begins using shared processor resources again.

The PAUSE register will be zero after the number of nanoseconds specified in the PAUSE instruction have elapsed (that is, PAUSE has auto-decremented down to 0) or when any of the following events occur, all of which cause PAUSE to be zeroed:

- an unmasked (PSTATE.ie = 1) disrupting trap request
- a deferred trap request

A **masked** trap request that occurs while a strand is paused has no effect on the suspension of the strand or on the contents of the PAUSE register. In particular, a masked trap request does **not** zero PAUSE.

Implementation Note | The PAUSE instruction is implemented as a WRAsr to ASR 27. See *Pause* on page 298 and *Write Ancillary State Register* on page 373 for details.

5.5.13 Mwait Count (MWAIT) Register (ASR 28)

The nonprivileged MWAIT register provides a means for software to voluntarily and temporarily pause execution, freeing up processor resources for use by other threads of execution. When MWAIT = 0, it has no effect on execution. When MWAIT ≠ 0, execution is paused (temporarily suspended) on the virtual processor (strand) until MWAIT = 0 again. There is one MWAIT register implemented for each virtual processor.

Software initiates a pause for up to n nanoseconds by writing the value n to the MWAIT register, using the MWAIT (or, less typically, the underlying WRMWAIT) instruction. MWAIT is write-only and cannot be read by software; in particular, there is no RDMWAIT instruction.

When MWAIT ≠ 0, it is automatically decremented once every nanosecond. When MWAIT = 0, no decrementation occurs.




FIGURE 5-22 MWAIT Register

IMPL. DEP. #550-S40: The maximum number of nanoseconds that a virtual processor may be paused by the MWAIT instruction, MAX_MWAIT_COUNT , is $2^B - 1$. B , the number of bits used to implement the `mwait_count` field, is implementation-dependent but B must be ≥ 13 . Therefore $MAX_MWAIT_COUNT \geq 2^{13} - 1$ (8191) nanoseconds. If an MWAIT instruction requests a pause of more than MAX_MWAIT_COUNT nanoseconds, the virtual processor must saturate the monitor-wait count to its maximum value, that is, write a value of MAX_MWAIT_COUNT to the MWAIT register.

Implementation Note | An implementation should implement enough bits in the MWAIT register to satisfy the equation:

$$\frac{(C \times 2^B) - (S \times 4)}{C \times 2^B} \geq 0.99$$

where:

B is the number of bits that must be implemented in the MWAIT register

C is the minimum number of clock cycles per nanosecond (under any conditions except *Parked*), and

S is the number of strands that share an instruction pipeline, and 4 represents the number of slots in the pipeline needed to execute the four instructions in a minimal busy-wait loop (load, compare, branch, MWAIT).

Solving this equation for B yields:

$$B \geq \text{ceiling}(\log_2(100 \times S + C)) + 2$$

Implementation Note | While MWAIT ≠ 0, the strand should suspend forward progress and should release all shared hardware resources to other strands.

When the value of MWAIT reaches 0, a paused strand resumes forward progress and begins using shared processor resources again.

If any of the following are true:

- there is no previously executed (in program order) Load-Monitor instruction (see *Load-Monitor ASIs* on page 437), or
- the memory location(s) read by the most recently executed (in program order) Load-Monitor instruction are modified by other than the strand executing the MWAIT after the Load-Monitor has read the location(s), or
- there is a previously executed (in program order) MWAIT or synchronous or asynchronous trap later (in program order) than the most recently executed (in program order) Load-Monitor instruction,

then:

- the contents of the MWAIT register are forced to zero, thus terminating any MWAIT pause, and
- any later (in program order) MWAIT or WRMWAIT instructions are executed as NOPs unless executed after (in program order) a later (in program order) Load Monitor instruction.

The MWAIT register will be zero after the number of nanoseconds specified in the MWAIT instruction have elapsed (that is, MWAIT has auto-decremented down to 0) or when any of the following events occur, all of which cause MWAIT to be zeroed:

- an **unmasked** (PSTATE.ie = 1) disrupting trap request
- a deferred trap request

A **masked** trap request that occurs while a strand is paused has no effect on the suspension of the strand or on the contents of the MWAIT register. In particular, a masked trap request does **not** zero MWAIT.

Implementation	The MWAIT instruction is implemented as a WRAsr to ASR 28.
Note	See <i>MWait</i> on page 292 and <i>Write Ancillary State Register</i> on page 373 for details.

5.6 Register-Window PR State Registers

The state of the register windows is determined by the contents of a set of privileged registers. These state registers can be read/written by privileged software using the RDPR/WRPR instructions. An attempt by nonprivileged software to execute a RDPR or WRPR instruction causes a *privileged_opcode* exception. In addition, these registers are modified by instructions related to register windows and are used to generate traps that allow supervisor software to spill, fill, and clean register windows.

Privileged registers CWP, CANSERVE, CANRESTORE, OTHERWIN, and CLEANWIN contain values in the range 0 .. $N_REG_WINDOWS - 1$. An attempt to write a value greater than $N_REG_WINDOWS - 2$ to CANSERVE, CANRESTORE, or OTHERWIN violates the register window state definition in *Register Window State Definition* on page 66. All five of these registers should have the same width.

Programming	Hardware is not required to prevent a value greater than
Note	$N_REG_WINDOWS - 2$ from being written to CANSERVE, CANRESTORE, or OTHERWIN; it is up to system software to keep the window state consistent.

IMPL. DEP. #126-V9-Ms10Cs40: An attempt to write a value greater than $N_REG_WINDOWS - 1$ to CWP, CANSERVE, CANRESTORE, OTHERWIN, or CLEANWIN causes an implementation-dependent value in the range 0 .. $N_REG_WINDOWS - 1$ to be written to the register.

Although the architectural width of each of these five registers is 5 bits, their actual implemented

width is implementation dependent and shall be between $\lceil \log_2(N_REG_WINDOWS) \rceil$ and 5 bits, inclusive. If fewer than 5 bits are implemented, the unimplemented upper bits shall read as 0 and writes to them shall have no effect.

Implementation Note | A write to any privileged register, including PR state registers, may drain the virtual processor pipeline.

Programming Note | Privileged software should not assume that $N_REG_WINDOWS$ is a power of 2, because an implementation is free to implement a non-power-of-2 value for $N_REG_WINDOWS$.

Programming Note | Privileged software should not assume that $N_REG_WINDOWS$ is a fixed value, in order to support live migration to a virtual processor with a different number of implemented register windows.

For details of how the window-management registers are used, see *Register Window Management Instructions* on page 90.

5.6.1 Current Window Pointer (CWP^P) Register (PR 9) (A1)

The privileged CWP register, shown in FIGURE 5-23, is a counter that identifies the current window into the array of integer registers. See *Register Window Management Instructions* on page 90 and Chapter 12, *Traps*, for information regarding how hardware manipulates the CWP register.




FIGURE 5-23 Current Window Pointer Register

Implementation Note | For Oracle SPARC Architecture 2015 processors, $N_REG_WINDOWS = 8$. Therefore, the CWP register is implemented with 3 bits and the maximum value for CWP is 7.

5.6.2 Savable Windows (CANSAVE^P) Register (PR 10) (A1)

The privileged CANSAVE register, shown in FIGURE 5-24, contains the number of register windows following CWP that are not in use and are, hence, available to be allocated by a SAVE instruction without generating a window spill exception.



FIGURE 5-24 CANSAVE Register, Figure 5-24, page 88

IMPL. DEP. #600-Ms50(a): An attempt to use a WRPR instruction to write a value greater than $N_REG_WINDOWS - 1$ to the CANSAVE register causes an implementation-dependent value between 0 and $N_REG_WINDOWS - 1$ (inclusive) to be written to the register. Although the architectural width of each of this register is 5 bits, its actual implemented width is implementation dependent and shall be between $\lceil \log_2(N_REG_WINDOWS) \rceil$ and 5 bits, inclusive. If fewer than 5 bits are implemented, the unimplemented upper bits shall read as 0 and writes to them shall have no effect.

Programming Note Per *Register Window State Definition* on page 66, the maximum valid value of CANSAVE is $N_REG_WINDOWS - 2$. However, implementations are not required to prevent out-of-range values (notably, $N_REG_WINDOWS - 1$) from being written to CANSAVE, so software should take care to never write such values to it.

5.6.3 Restorable Windows (CANRESTORE^P) Register (PR 11) (A1)

The privileged CANRESTORE register, shown in FIGURE 5-25, contains the number of register windows preceding CWP that are in use by the current program and can be restored (by the RESTORE instruction) without generating a window fill exception.




FIGURE 5-25 CANRESTORE Register

IMPL. DEP. #600-Ms50(b): An attempt to use a WRPR instruction to write a value greater than $N_REG_WINDOWS - 1$ to the CANRESTORE register causes an implementation-dependent value between 0 and $N_REG_WINDOWS - 1$ (inclusive) to be written to the register. Although the architectural width of each of this register is 5 bits, its actual implemented width is implementation dependent and shall be between $\lceil \log_2(N_REG_WINDOWS) \rceil$ and 5 bits, inclusive. If fewer than 5 bits are implemented, the unimplemented upper bits shall read as 0 and writes to them shall have no effect.

Programming Note Per *Register Window State Definition* on page 66, the maximum valid value of CANRESTORE is $N_REG_WINDOWS - 2$. However, implementations are not required to prevent out-of-range values (notably, $N_REG_WINDOWS - 1$) from being written to CANRESTORE, so software should take care to never write such values to it.

5.6.4 Clean Windows (CLEANWIN^P) Register (PR 12) (A1)

The privileged CLEANWIN register, shown in FIGURE 5-26, contains the number of windows that can be used by the SAVE instruction without causing a *clean_window* exception.




FIGURE 5-26 CLEANWIN Register

The CLEANWIN register counts the number of register windows that are “clean” with respect to the current program; that is, register windows that contain only zeroes, valid addresses, or valid data from that program. Registers in these windows need not be cleaned before they can be used. The count includes the register windows that can be restored (the value in the CANRESTORE register) and the register windows following CWP that can be used without cleaning. When a clean window is requested (by a SAVE instruction) and none is available, a *clean_window* exception occurs to cause the next window to be cleaned.

IMPL. DEP. #600-Ms50(c): An attempt to use a WRPR instruction to write a value greater than $N_REG_WINDOWS - 1$ to the CLEANWIN register causes an implementation-dependent value between 0 and $N_REG_WINDOWS - 1$ (inclusive) to be written to the register. Although the architectural width of each of this register is 5 bits, its actual implemented width is implementation dependent and shall be between $\lceil \log_2(N_REG_WINDOWS) \rceil$ and 5 bits, inclusive. If fewer than 5 bits are implemented, the unimplemented upper bits shall read as 0 and writes to them shall have no effect.

Implementation Note For Oracle SPARC Architecture 2015 processors, $N_REG_WINDOWS = 8$. Therefore, the CLEANWIN register is implemented with 3 bits and the maximum value for CLEANWIN is 7. When this register is written by the WRPR instruction, bits 63:3 of the data written are ignored.

5.6.5 Other Windows (OTHERWIN^P) Register (PR 13) (A1)

The privileged OTHERWIN register, shown in FIGURE 5-27, contains the count of register windows that will be spilled/filled by a separate set of trap vectors based on the contents of WSTATE.other. If OTHERWIN is zero, register windows are spilled/filled by use of trap vectors based on the contents of WSTATE.normal.

The OTHERWIN register can be used to split the register windows among different address spaces and handle spill/fill traps efficiently by use of separate spill/fill vectors.




FIGURE 5-27 OTHERWIN Register

IMPL. DEP. #600-Ms50(d): An attempt to use a WRPR instruction to write a value greater than $N_REG_WINDOWS - 1$ to the OTHERWIN register causes an implementation-dependent value between 0 and $N_REG_WINDOWS - 1$ (inclusive) to be written to the register. Although the architectural width of each of this register is 5 bits, its actual implemented width is implementation dependent and shall be between $\lceil \log_2(N_REG_WINDOWS) \rceil$ and 5 bits, inclusive. If fewer than 5 bits are implemented, the unimplemented upper bits shall read as 0 and writes to them shall have no effect.

Programming Note Per *Register Window State Definition* on page 66, the maximum valid value of OTHERWIN is $N_REG_WINDOWS - 2$. However, implementations are not required to prevent out-of-range values (notably, $N_REG_WINDOWS - 1$) from being written to OTHERWIN, so software should take care to never write such values to it.

5.6.6 Window State (WSTATE^P) Register (PR 14) (A1)

The privileged WSTATE register, shown in FIGURE 5-28, specifies bits that are inserted into TT[TL]{4:2} on traps caused by window spill and fill exceptions. These bits are used to select one of eight different window spill and fill handlers. If OTHERWIN = 0 at the time a trap is taken because of a window spill or window fill exception, then the WSTATE.normal bits are inserted into TT[TL]. Otherwise, the WSTATE.other bits are inserted into TT[TL]. See *Register Window State Definition*, below, for details of the semantics of OTHERWIN.




FIGURE 5-28 WSTATE Register

5.6.7 Register Window Management

The state of the register windows is determined by the contents of the set of privileged registers described in *Register-Window PR State Registers* on page 62. Those registers are affected by the instructions described in *Register Window Management Instructions* on page 90. Privileged software can read/write these state registers directly by using RDPR/WRPR instructions.

5.6.7.1 Register Window State Definition

For the state of the register windows to be consistent, the following must always be true:

$$\text{CANSAVE} + \text{CANRESTORE} + \text{OTHERWIN} = N_REG_WINDOWS - 2$$

FIGURE 5-3 on page 37 shows how the register windows are partitioned to obtain the above equation. The partitions are as follows:

- The current window plus the window that must not be used because it overlaps two other valid windows. In FIGURE 5-3, these are windows 0 and 5, respectively. They are always present and account for the “2” subtracted from $N_REG_WINDOWS$ in the right-hand side of the above equation.
- Windows that do not have valid contents and that can be used (through a SAVE instruction) without causing a spill trap. These windows (windows 1–4 in FIGURE 5-3) are counted in CANSAVE.
- Windows that have valid contents for the current address space and that can be used (through the RESTORE instruction) without causing a fill trap. These windows (window 7 in FIGURE 5-3) are counted in CANRESTORE.
- Windows that have valid contents for an address space other than the current address space. An attempt to use these windows through a SAVE (RESTORE) instruction results in a spill (fill) trap to a separate set of trap vectors, as discussed in the following subsection. These windows (window 6 in FIGURE 5-3) are counted in OTHERWIN.

In addition,

$$\text{CLEANWIN} \geq \text{CANRESTORE}$$

since CLEANWIN is the sum of CANRESTORE and the number of clean windows following CWP.

For the window-management features of the architecture described in this section to be used, the state of the register windows must be kept consistent at all times, except within the trap handlers for window spilling, filling, and cleaning. While window traps are being handled, the state may be inconsistent. Window spill/fill trap handlers should be written so that a nested trap can be taken without destroying state.

Programming Note System software is responsible for keeping the state of the register windows consistent at all times. Failure to do so will cause undefined behavior. For example, CANSAVE, CANRESTORE, and OTHERWIN must never be greater than or equal to $N_REG_WINDOWS - 1$.

5.6.7.2 Register Window Traps

Window traps are used to manage overflow and underflow conditions in the register windows, support clean windows, and implement the FLUSHW instruction.

See *Register Window Traps* on page 465 for a detailed description of how fill, spill, and *clean_window* traps support register windowing.

5.7 Non-Register-Window PR State Registers

The registers described in this section are visible only to software running in privileged mode (that is, when $PSTATE.priv = 1$), and may be accessed with the WRPR and RDPR instructions. (An attempt to execute a WRPR or RDPR instruction in nonprivileged mode causes a *privileged_opcode* exception.)

Each virtual processor provides a full set of these state registers.

Implementation | A write to any privileged register, including PR state registers,
Note | may drain the CPU pipeline.

5.7.1 Trap Program Counter (TPC^P) Register (PR 0) (A1)

The privileged Trap Program Counter register (TPC; FIGURE 5-29) contains the program counter (PC) from the previous trap level. There are *MAXPTL* instances of the TPC, but only one is accessible at any time. The current value in the TL register determines which instance of the TPC[TL] register is accessible. An attempt to read or write the TPC register when TL = 0 causes an *illegal_instruction* exception.

	RW	R
TPC ₁ ^P	pc_high62 (PC{63:2} from trap while TL = 0)	00
TPC ₂ ^P	pc_high62 (PC{63:2} from trap while TL = 1)	00
TPC ₃ ^P	pc_high62 (PC{63:2} from trap while TL = 2)	00
:	:	:
TPC _{MAXPTL} ^P	pc_high62 (PC{63:2} from trap while TL = MAXPTL - 1)	00

63 2 1 0

FIGURE 5-29 Trap Program Counter Register Stack

During normal operation, the value of TPC[*n*], where *n* is greater than the current trap level (*n* > TL), is undefined.

TABLE 5-16 lists the events that cause TPC to be read or written.

TABLE 5-16 Events that involve TPC, when executing with TL = *n*.

Event	Effect
Trap	TPC[<i>n</i> + 1] ← PC
RETRY instruction	PC ← TPC[<i>n</i>]
RDPR (TPC)	R[rd] ← TPC[<i>n</i>]
WRPR (TPC)	TPC[<i>n</i>] ← <i>value</i>

5.7.2 Trap Next PC (TNPC^P) Register (PR 1) (A1)

The privileged Trap Next Program Counter register (TNPC; FIGURE 5-30) is the next program counter (NPC) from the previous trap level. There are *MAXPTL* instances of the TNPC, but only one is accessible at any time. The current value in the TL register determines which instance of the TNPC register is accessible. An attempt to read or write the TNPC register when TL = 0 causes an *illegal_instruction* exception.

	RW	R
TNPC ₁ ^P	npc_high62 (NPC{63:2} from trap while TL = 0)	00
TNPC ₂ ^P	npc_high62 (NPC{63:2} from trap while TL = 1)	00
TNPC ₃ ^P	npc_high62 (NPC{63:2} from trap while TL = 2)	00
:	:	:
TNPC _{MAXPTL} ^P	npc_high62 (NPC{63:2} from trap while TL = MAXPTL - 1)	00

63 2 1 0

FIGURE 5-30 Trap Next Program Counter Register Stack

During normal operation, the value of $TNPC[n]$, where n is greater than the current trap level ($n > TL$), is undefined.

TABLE 5-17 lists the events that cause $TNPC$ to be read or written.

TABLE 5-17 Events that involve $TNPC$, when executing with $TL = n$.

Event	Effect
Trap	$TNPC[n + 1] \leftarrow NPC$
DONE instruction	$PC \leftarrow TNPC[n]$; $NPC \leftarrow TNPC[n] + 4$
RETRY instruction	$NPC \leftarrow TNPC[n]$
RDPR ($TNPC$)	$R[rd] \leftarrow TNPC[n]$
WRPR ($TNPC$)	$TNPC[n] \leftarrow value$

5.7.3 Trap State ($TSTATE^P$) Register (PR 2) (A1)

The privileged Trap State register ($TSTATE$; FIGURE 5-31) contains the state from the previous trap level, comprising the contents of the GL , CCR , ASI , $PSTATE$, and CWP registers from the previous trap level. There are $MAXPTL$ instances of the $TSTATE$ register, but only one is accessible at a time. The current value in the TL register determines which instance of $TSTATE$ is accessible. An attempt to read or write the $TSTATE$ register when $TL = 0$ causes an *illegal_instruction* exception.

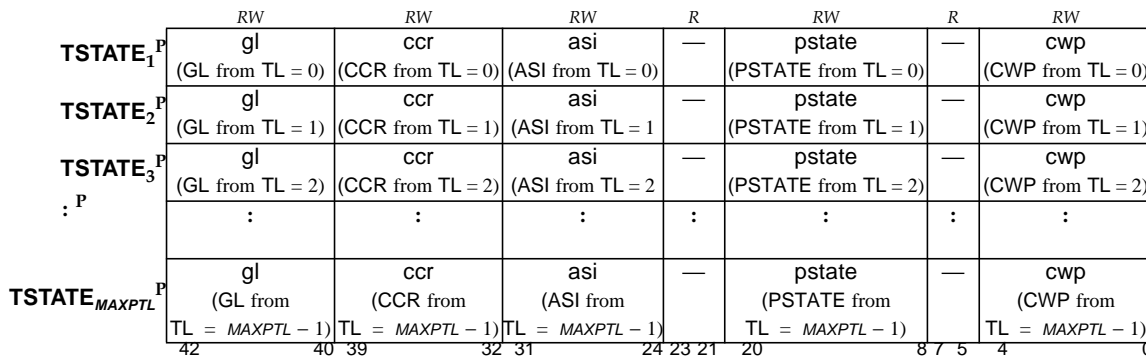


FIGURE 5-31 Trap State ($TSTATE$) Register Stack

During normal operation the value of $TSTATE[n]$, when n is greater than the current trap level ($n > TL$), is undefined.

V9 Compatibility Note | Because there are more bits in the Oracle SPARC Architecture's $PSTATE$ register than in a SPARC V9 $PSTATE$ register, a 13-bit $PSTATE$ value is stored in $TSTATE$ instead of the 10-bit value specified in the SPARC V9 architecture.

TABLE 5-18 lists the events that cause $TSTATE$ to be read or written.

TABLE 5-18 Events That Involve $TSTATE$, When Executing with $TL = n$

Event	Effect
Trap	$TSTATE[n + 1] \leftarrow (registers)$
DONE instruction	$(registers) \leftarrow TSTATE[n]$
RETRY instruction	$(registers) \leftarrow TSTATE[n]$
RDPR ($TSTATE$)	$R[rd] \leftarrow TSTATE[n]$
WRPR ($TSTATE$)	$TSTATE[n] \leftarrow value$

5.7.4 Trap Type (TT^P) Register (PR 3) (A1)

The privileged Trap Type register (TT; see FIGURE 5-32) contains the trap type of the trap that caused entry to the current trap level. There are *MAXPTL* instances of the TT register, but only one is accessible at a time. The current value in the TL register determines which instance of the TT register is accessible. An attempt to read or write the TT register when TL = 0 causes an *illegal_instruction* exception.




FIGURE 5-32 Trap Type Register Stack

During normal operation, the value of TT[n], where *n* is greater than the current trap level (*n* > TL), is undefined.

TABLE 5-19 lists the events that cause TT to be read or written.

TABLE 5-19 Events that involve TT, when executing with TL = *n*.

Event	Effect
Trap	TT[n + 1] ← (trap type)
RDPR (TT)	R[rd] ← TT[n]
WRPR (TT)	TT[n] ← <i>value</i>

5.7.5 Tick (TICK) Register (PR 4) (A1)

PR 4 and ASR 4 refer to the same physical TICK register. See *Tick (tick) Register (ASR 4)* on page 52 for a description of that register and how it can be accessed by the RDAsr, RDPR, and WRPR instructions.

5.7.6 Trap Base Address (TBA^P) Register (PR 5) (A1)

The privileged Trap Base Address register (TBA), shown in FIGURE 5-33, provides the upper 49 bits (bits 63:15) of the virtual address used to select the trap vector for a trap that is to be delivered to privileged mode. The lower 15 bits of the TBA always read as zero, and writes to them are ignored.




FIGURE 5-33 Trap Base Address Register

Details on how the full address for a trap vector is generated, using TBA and other state, are provided in *Trap-Table Entry Address to Privileged Mode* on page 450.

5.7.7 Processor State (PSTATE^P) Register (PR 6) (A1)

The privileged Processor State register (PSTATE), shown in , contains control fields for the current state of the virtual processor. There is only one instance of the PSTATE register per virtual processor.




FIGURE 5-34 PSTATE Fields

Writes to PSTATE are nondelayed; that is, new machine state written to PSTATE is visible to the next instruction executed. The privileged RDPR and WRPR instructions are used to read and write PSTATE, respectively.

The following subsections describe the fields of the PSTATE register.

Trap on Control Transfer (tct). PSTATE.tct enables the Trap-on-Control-Transfer feature. When PSTATE.tct = 1, the virtual processor monitors each control transfer instruction (CTI) to determine whether a *control_transfer_instruction* exception should be generated. If the virtual processor is executing a CTI, PSTATE.tct = 1, and a successful control transfer is going to occur as a result of execution of that CTI, the processor generates a *control_transfer_instruction* exception instead of completing execution of the control transfer instruction.

When the trap is taken, the address of the CTI (the value of PC when the CTI began execution) is saved in TPC[TL] and the value of NPC when the CTI began execution is saved in TNPC[TL].

During initial trap processing, before trap handler code is executed, the virtual processor sets PSTATE.tct to 0 (so that control transfers within the trap handler don't cause additional traps).

Programming Note Trap handler software for a *control_transfer_instruction* trap should take care when returning to the software that caused the trap. Execution of DONE or RETRY causes PSTATE.tct to be restored from TSTATE, normally setting PSTATE.tct back to 1. If trap handler software intends for *control_transfer_instruction* exceptions to be reenabled, then it must emulate the trapped control transfer instruction.

IMPL. DEP. #450-S20: Availability of the *control_transfer_instruction* exception feature is implementation dependent. If not implemented, trap type 074₁₆ is unused, PSTATE.tct always reads as zero, and writes to PSTATE.tct are ignored.

For the purposes of the *control_transfer_instruction* exception, a discontinuity in instruction-fetch addresses caused by a WRPR to PSTATE that changes the value of PSTATE.am (and thus, potentially the more-significant 32 bits of the address of the next instruction; see page 73) is *not* considered a control transfer. Only explicit CTIs can generate a *control_transfer_instruction* exception.

Current Little Endian (cle). This bit affects the endianness of data accesses performed using an implicit ASI. When PSTATE.cle = 1, all data accesses using an implicit ASI are performed in little-endian byte order. When PSTATE.cle = 0, all data accesses using an implicit ASI are performed in big-endian byte order. Specific ASIs used are shown in TABLE 6-3 on page 83. Note that the endianness of a data access may be further affected by TTE.ie used by the MMU.

Instruction accesses are unaffected by PSTATE.cle and are always performed in big-endian byte order.

Trap Little Endian (tle). When a trap is taken, the current PSTATE register is pushed onto the trap stack. During a virtual processor trap to privileged mode, the PSTATE.tle bit is copied into PSTATE.cle in the new PSTATE register. This behavior allows system software to have a different implicit byte ordering than the current process. Thus, if PSTATE.tle is set to 1, data accesses using an implicit ASI in the trap handler are little-endian.

The original state of `PSTATE.cle` is restored when the original `PSTATE` register is restored from the trap stack.

Memory Model (mm). This 2-bit field determines the memory model in use by the virtual processor. The defined values for an Oracle SPARC Architecture virtual processor are listed in TABLE 5-20.

TABLE 5-20 `PSTATE.mm` Encodings

mm Value	Selected Memory Model
00	Total Store Order (TSO)
01	<i>Reserved</i>
10	<i>Implementation dependent</i> (impl. dep. #113-V9-Ms10)
11	<i>Implementation dependent</i> (impl. dep. #113-V9-Ms10)

The current memory model is determined by the value of `PSTATE.mm`. Software should refrain from writing the values 01_2 , 10_2 , or 11_2 to `PSTATE.mm` because they are implementation-dependent or reserved for future extensions to the architecture, and in any case not currently portable across implementations.

- **Total Store Order (TSO)** — Loads are ordered with respect to earlier loads. Stores are ordered with respect to earlier loads and stores. Thus, loads can bypass earlier stores but cannot bypass earlier loads; stores cannot bypass earlier loads or stores.

IMPL. DEP. #113-V9-Ms10: Whether memory models represented by `PSTATE.mm = 102` or `112` are supported in an Oracle SPARC Architecture processor is implementation dependent. If the 10_2 model is supported, then when `PSTATE.mm = 102` the implementation must correctly execute software that adheres to the RMO model described in *The SPARC Architecture Manual-Version 9*. If the 11_2 model is supported, its definition is implementation dependent.

IMPL. DEP. #119-Ms10: The effect of writing an unimplemented memory model designation into `PSTATE.mm` is implementation dependent.

SPARC V9 Compatibility Notes	<p>The PSO memory model described in SPARC V8 and SPARC V9 architecture specifications was never implemented in a SPARC V9 implementation and is not included in the Oracle SPARC Architecture specification.</p> <p>The RMO memory model described in the SPARC V9 specification was implemented in some non-Sun SPARC V9 implementations, but is not directly supported in Oracle SPARC Architecture 2015 implementations. All software written to run correctly under RMO will run correctly under TSO on an Oracle SPARC Architecture 2015 implementation.</p>
---	--

Enable FPU (pef). When set to 1, the `PSTATE.pef` bit enables the floating-point unit. This allows privileged software to manage the FPU. For the FPU to be usable, both `PSTATE.pef` and `FPRS.fef` must be set to 1. Otherwise, any floating-point instruction that tries to reference the FPU causes an *fp_disabled* trap.

If an implementation does not contain a hardware FPU, `PSTATE.pef` always reads as 0 and writes to it are ignored.

Address Mask (am). The `PSTATE.am` bit is provided to allow 32-bit SPARC software to run correctly on a 64-bit SPARC processor. When `PSTATE.am = 1`, bits 63:32 of virtual addresses are masked out (treated as 0).

When `PSTATE.am = 0`:

- bits 63:0 of each virtual address are passed along in the virtual address to the memory system.

When `PSTATE.am = 1`:

- bits 63:32 of each virtual address are masked out (treated as 0) in the virtual address passed to the memory system, as described below
- bits 31:0 of each virtual addresses are passed along as-is in the virtual address passed to the memory system

`PSTATE.am` only affects virtual addresses (including those referenced using `ASI_AS_IF_USER*` ASIs); it does not affect real addresses.

When `PSTATE.am = 0`, the full 64 bits of all instruction and data addresses are *preserved* at all points in the virtual processor.

When an MMU is disabled, `PSTATE.am` has no effect on (does not cause masking of) addresses.

Programming Note	It is the responsibility of privileged software to manage the setting of the <code>PSTATE.am</code> bit, since hardware masks virtual addresses when <code>PSTATE.am = 1</code> . Misuse of the <code>PSTATE.am</code> bit can result in undesirable behavior. In particular, <code>PSTATE.am</code> should <i>not</i> be set to 1 in privileged mode. The <code>PSTATE.am</code> bit should always be set to 1 when 32-bit nonprivileged software is executed.
-------------------------	---

Instances in which the more-significant 32 bits of a virtual address **are masked** when `PSTATE.am = 1` include:

- Before any data virtual address is sent out of the virtual processor (notably, to the memory system, which includes MMU, internal caches, and external caches).
- Before any instruction virtual address is sent out of the virtual processor (notably, to the memory system, which includes MMU, internal caches, and external caches)
- When the value of `PC` is stored to a general-purpose register by a `CALL`, `JMPL`, or `RDPC` instruction (closed impl.dep. #125-V9-Cs10)
- When the values of `PC` and `NPC` are written to `TPC[TL]` and `TNPC[TL]` (respectively) during a trap (closed impl.dep. #125-V9-Cs10)
- Before any address is sent to a watchpoint comparator

Programming Note	A 64-bit comparison is always used when performing a masked watchpoint address comparison with the Instruction or Data VA watchpoint register. When <code>PSTATE.am = 1</code> , the more significant 32 bits of the VA watchpoint register must be zero for a match (and resulting trap) to occur.
-------------------------	---

When `PSTATE.am = 1`, the more-significant 32 bits of a virtual address **are explicitly preserved and not masked** out in the following cases:

- When a target address is written to `NPC` by a control transfer instruction
- When `NPC` is incremented to `NPC + 4` during execution of an instruction that is not a taken control transfer
- When the address is a nontranslating address (an address that is never translated by an MMU, such as the address of an internal register accessed using an ASI)
- When a `WRPR` instruction writes to `TPC[TL]` or `TNPC[TL]`

Programming Note Since writes to PSTATE are nondelayed (see page 70), a change to PSTATE.am can affect which instruction is executed immediately after the write to PSTATE.am. Specifically, if a WRPR to the PSTATE register changes the value of PSTATE.am from '0' to '1', and NPC{63:32} when the WRPR began execution was nonzero, then the next instruction executed after the WRPR will be from the address indicated in NPC{31:0} (with the more-significant 32 address bits set to zero).

- When a RDPR instruction reads from TPC[TL] or TNPC[TL]

If (1) TSTATE[TL].pstate.am = 1 and (2) a DONE or RETRY instruction is executed¹, it is implementation dependent whether the DONE or RETRY instruction masks (zeroes) the more-significant 32 bits of the values it places into PC and NPC (impl. dep. #417-S10).

Programming Note Because of implementation dependency #417-S10, great care must be taken in trap handler software if TSTATE[TL].pstate.am = 1 and the trap handler wishes to write a nonzero value to the more-significant 32 bits of TPC[TL] or TNPC[TL].

Programming Note PSTATE.am affects the operation of the edge-handling instructions, EDGE<8 | 16 | 32>[L]*. See *Edge Handling Instructions* on page 149 and *Edge Handling Instructions (no CC)* on page 150.

Privileged Mode (priv). When PSTATE.priv = 1, the virtual processor is operating in privileged mode.

When PSTATE.priv = 0, the processor is operating in nonprivileged mode

PSTATE_interrupt_enable (ie). PSTATE.ie controls when the virtual processor can take traps due to disrupting exceptions (such as interrupts or errors unrelated to instruction processing).

Outstanding disrupting exceptions that are destined for privileged mode can only cause a trap when the virtual processor is in nonprivileged or privileged mode and PSTATE.ie = 1. At all other times, they are held pending. For more details, see *Conditioning of Disrupting Traps* on page 448.

SPARC V9 Compatibility Note Since the Oracle SPARC Architecture provides a more general “alternate globals” facility (through use of the GL register) than does SPARC V9, an Oracle SPARC Architecture processor does not implement the SPARC V9 PSTATE.ag bit.

5.7.8 Trap Level Register (TL^P) (PR 7) (A1)

The privileged Trap Level register (TL; FIGURE 5-35) specifies the current trap level. TL = 0 is the normal (nontrap) level of operation. TL > 0 implies that one or more traps are being processed.




FIGURE 5-35 Trap Level Register

¹ which sets PSTATE.am to '1', by restoring the value from TSTATE[TL].pstate.am to PSTATE.am

The maximum valid value that the TL register may contain is $MAXPTL$, which is always equal to the number of supported trap levels beyond level 0.

IMPL. DEP. #101-V9-CS10: The architectural parameter $MAXPTL$ is a constant for each implementation; its legal values are from 2 to 6 (supporting from 2 to 6 levels of saved trap state). In a typical implementation $MAXPTL = MAXPGL$ (see impl. dep. #401-S10). Architecturally, $MAXPTL$ must be ≥ 2 .

In an Oracle SPARC Architecture 2015 implementation, $MAXPTL = 2$. See Chapter 12, *Traps*, for more details regarding the TL register.

The effect of writing to TL with a WRPR instruction is summarized in TABLE 5-21.

TABLE 5-21 Effect of WRPR of Value x to Register TL

Value x Written with WRPR	Privilege Level when Executing WRPR	
	Nonprivileged	Privileged
$x \leq MAXPTL$	<i>privileged_opcode</i> exception	TL $\leftarrow x$
$x > MAXPTL$		TL $\leftarrow MAXPTL$ (no exception generated)

Writing the TL register with a WRPR instruction does not alter any other machine state; that is, it is *not* equivalent to taking a trap or returning from a trap.

Programming Note An Oracle SPARC Architecture implementation only needs to implement sufficient bits in the TL register to encode the maximum trap level value. In an implementation where $MAXPTL \leq 3$, bits 63:2 of data written to the TL register using the WRPR instruction are ignored; only the least-significant two bits (bits 1:0) of TL are actually written. For example, if $MAXPTL = 2$, writing a value of 05_{16} to the TL register causes a value of 1_{16} to actually be stored in TL.

Implementation Note $MAXPTL = 2$ for all Oracle SPARC Architecture 2015 processors. Writing a value between 3 and 7 to the TL register in privileged mode causes a 2 to be stored in TL.

Programming Note Although it is possible for privileged software to set $TL > 0$ for nonprivileged software[†], an Oracle SPARC Architecture virtual processor's behavior when executing with $TL > 0$ in nonprivileged mode is undefined.

[†] by executing a WRPR to TSTATE followed by DONE instruction or RETRY instruction.

5.7.9 Processor Interrupt Level (PIL^P) Register (PR 8) A1

The privileged Processor Interrupt Level register (PIL; see FIGURE 5-36) specifies the interrupt level above which the virtual processor will accept an *interrupt_level_n* interrupt. Interrupt priorities are mapped so that interrupt level 2 has greater priority than interrupt level 1, and so on. See TABLE 12-4 on page 453 for a list of exception and interrupt priorities.




FIGURE 5-36 Processor Interrupt Level Register

V9 Compatibility Note | On SPARC V8 processors, the level 15 interrupt is considered to be nonmaskable, so it has different semantics from other interrupt levels. SPARC V9 processors do not treat a level 15 interrupt differently from other interrupt levels.

5.7.10 Global Level Register (GL^P) (PR 16) A1

The privileged Global Level (GL) register selects which set of global registers is visible at any given time.

FIGURE 5-37 illustrates the Global Level register.




FIGURE 5-37 Global Level Register, GL

When a trap occurs, GL is stored in TSTATE[TL].gl, GL is incremented, and a new set of global registers (R[1] through R[7]) becomes visible. A DONE or RETRY instruction restores the value of GL from TSTATE[TL].

The valid range of values that the GL register may contain is 0 to *MAXPGL*, where *MAXPGL* is one fewer than the number of global register sets available to the virtual processor.

IMPL. DEP. #401-S10: The architectural parameter *MAXPGL* is a constant for each implementation; its legal values are from 2 to 7 (supporting from 3 to 8 sets of global registers). In a typical implementation, *MAXPGL* = *MAXPTL* (see impl. dep. #101-V9-CS10). Architecturally, *MAXPGL* must be ≥ 2 .

In all Oracle SPARC Architecture 2015 implementations, *MAXPGL* = 2 (impl. dep. #401-S10).

IMPL. DEP. #400-S10: Although GL is defined as a 3-bit register, an implementation may implement any subset of those bits sufficient to encode the values from 0 to *MAXPGL* for that implementation. If any bits of GL are not implemented, they read as zero and writes to them are ignored.

GL operates similarly to TL, in that it increments during entry to a trap, but the values of GL and TL are independent. That is, $TL = n$ does not imply that $GL = n$, and $GL = n$ does not imply that $TL = n$. Furthermore, there may be a different total number of global levels (register sets) than there are trap levels; that is, *MAXPTL* and *MAXPGL* are not necessarily equal.

The GL register can be accessed directly with the RDPR and WRPR instructions (as privileged register number 16). Writing the GL register directly with WRPR will change the set of global registers visible to all instructions subsequent to the WRPR.

In privileged mode, attempting to write a value greater than *MAXPGL* to the GL register causes *MAXPGL* to be written to GL.

The effect of writing to GL with a WRPR instruction is summarized in TABLE 5-22.

TABLE 5-22 Effect of WRPR to Register GL

Value <i>x</i> Written with WRPR	Privilege Level when WRPR Is Executed		
	Nonprivileged	Privileged	
$x \leq \text{MAXPGL}$		$GL \leftarrow x$	
$x > \text{MAXPGL}$	<i>privileged_opcode</i> exception	$GL \leftarrow \text{MAXPGL}$	(no exception generated)

Since TSTATE itself is software-accessible, it is possible that when a DONE or RETRY is executed to return from a trap handler, the value of GL restored from TSTATE[TL] will be different from that which was saved into TSTATE[TL] when the trap occurred.

Instruction Set Overview

Instructions are fetched by the virtual processor from memory and are executed, annulled, or trapped. Instructions are encoded in 4 major formats and partitioned into 11 general categories. Instructions are described in the following sections:

- **Instruction Execution** on page 77.
- **Instruction Formats** on page 78.
- **Instruction Categories** on page 78.

6.1 Instruction Execution

The instruction at the memory location specified by the program counter is fetched and then executed. Instruction execution may change program-visible virtual processor and/or memory state. As a side effect of its execution, new values are assigned to the program counter (PC) and the next program counter (NPC).

An instruction may generate an exception if it encounters some condition that makes it impossible to complete normal execution. Such an exception may in turn generate a precise trap. Other events may also cause traps: an exception caused by a previous instruction (a deferred trap), an interrupt or asynchronous error (a disrupting trap), or a reset request (a reset trap). If a trap occurs, control is vectored into a trap table. See Chapter 12, *Traps*, for a detailed description of exception and trap processing.

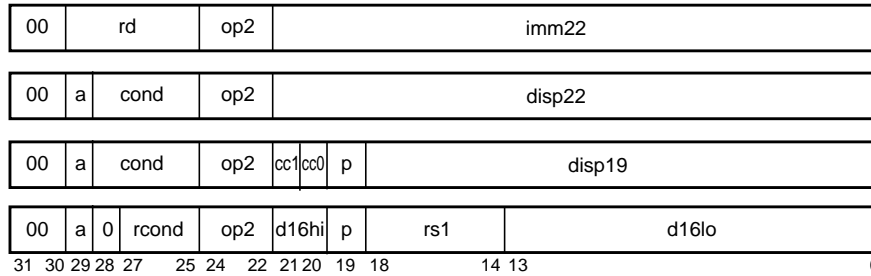
If a trap does not occur and the instruction is not a control transfer, the next program counter is copied into the PC, and the NPC is incremented by 4 (ignoring arithmetic overflow if any). There are two types of control-transfer instructions (CTIs): delayed and immediate. For a delayed CTI, at the end of the execution of the instruction, NPC is copied into the PC and the target address is copied into NPC. For an immediate CTI, at the end of execution, the target is copied to PC and target + 4 is copied to NPC. In the SPARC instruction set, many CTIs do not transfer control until after a delay of one instruction, hence the term “delayed CTI” (DCTI). Thus, the two program counters provide for a delayed-branch execution model.

For each instruction access and each normal data access, an 8-bit address space identifier (ASI) is appended to the 64-bit memory address. Load/store alternate instructions (see *Address Space Identifiers (ASIs)* on page 83) can provide an arbitrary ASI with their data addresses or can use the ASI value currently contained in the ASI register.


6.2 Instruction Formats

Every instruction is encoded in a single 32-bit word. The most typical 32-bit formats are shown in FIGURE 6-1. For detailed formats for specific instructions, see individual instruction descriptions in the *Instructions* chapter.

$op = 00_2$: *SETHI, Branches, and ILLTRAP*



$op = 01_2$: *CALL*



$op = 10_2$ or 11_2 : *Arithmetic, Logical, Moves, Tcc, Loads, Stores, Prefetch, and Misc*



FIGURE 6-1 Summary of Instruction Formats

6.3 Instruction Categories

Oracle SPARC Architecture instructions can be grouped into the following categories:

- Memory access
- Memory synchronization
- Integer arithmetic
- Control transfer (CTI)
- Conditional moves
- Register window management
- State register access
- Privileged register access
- Floating-point operate
- Implementation dependent
- Reserved

These categories are described in the following subsections.

6.3.1 Memory Access Instructions

Load, store, load-store, and PREFETCH instructions are the only instructions that access memory. All of the memory access instructions except CASA, CASXA, and Partial Store use either two R registers or an R register and simm13 to calculate a 64-bit byte memory address. For example, Compare and Swap uses a single R register to specify a 64-bit byte memory address. To this 64-bit address, an ASI is appended that encodes address space information.

The destination field of a memory reference instruction specifies the R or F register(s) that supply the data for a store or that receive the data from a load or LDSTUB. For SWAP, the destination register identifies the R register to be exchanged atomically with the calculated memory location. For Compare and Swap, an R register is specified, the value of which is compared with the value in memory at the computed address. If the values are equal, then the destination field specifies the R register that is to be exchanged atomically with the addressed memory location. If the values are unequal, then the destination field specifies the R register that is to receive the value at the addressed memory location; in this case, the addressed memory location remains unchanged. LDFSR/LDXEFSR/LDXFSR and STFSR/STXFSR are special load and store instructions that load or store the floating-point status register, FSR, instead of acting on an R or F register.

The destination field of a PREFETCH instruction (fcn) is used to encode the type of the prefetch.

Memory is byte (8-bit) addressable. Integer load and store instructions support byte, halfword (2 bytes), word (4 bytes), and doubleword/extended-word (8 bytes) accesses. Floating-point load and store instructions support word, doubleword, and quadword memory accesses. LDSTUB accesses bytes, SWAP accesses words, CASA accesses words, and CASXA accesses doublewords. The LDTXA (load twin-extended-word) instruction accesses a quadword (16 bytes) in memory. Block loads and stores access 64-byte aligned data. PREFETCH accesses at least 64 bytes.

Programming Note	For some instructions, by use of simm13, any location in the lowest or highest 4 Kbytes of an address space can be accessed without the use of a register to hold part of the address.
-------------------------	--

6.3.1.1 Memory Alignment Restrictions

A halfword access must be aligned on a 2-byte boundary, a word access (including an instruction fetch) must be aligned on a 4-byte boundary, an extended-word (LDX, LDXA, STX, STXA) or integer twin word (LDTW, LDTWA, STTW, STTWA) access must be aligned on an 8-byte boundary, an integer twin-extended-word (LDTXA) access must be aligned on a 16-byte boundary, and a Block Load (LDBLOCKF^D) or Store (STBLOCKF^D) access must be aligned on a 64-byte boundary.

A floating-point doubleword access (LDDF, LDDFA, STDF, STDFA) should be aligned on an 8-byte boundary, but is only required to be aligned on a word (4-byte) boundary. A floating-point doubleword access to an address that is 4-byte aligned but not 8-byte aligned may result in less efficient and nonatomic access (causes a trap and is emulated in software (impl. dep. #109-V9-Cs10)), so 8-byte alignment is recommended.

A floating-point quadword access (LDQF, LDQFA, STQF, STQFA) should be aligned on a 16-byte boundary, but is only required to be aligned on a word (4-byte) boundary. A floating-point quadword access to an address that is 4-byte or 8-byte aligned but not 16-byte aligned may result in less efficient and nonatomic access (causes a trap and is emulated in software (impl. dep. #111-V9-Cs10)), so 16-byte alignment is recommended.

An improperly aligned address in a load, store, or load-store instruction causes a *mem_address_not_aligned* exception to occur, with these exceptions:

- An LDDF or LDDFA instruction accessing an address that is word aligned but not doubleword aligned may cause an *LDDF_mem_address_not_aligned* exception (impl. dep. #109-V9-Cs10).
- An STDF or STDFA instruction accessing an address that is word aligned but not doubleword aligned may cause an *STDF_mem_address_not_aligned* exception (impl. dep. #110-V9-Cs10).

- An LDQF or LDQFA instruction accessing an address that is word aligned but not quadword aligned may cause an *LDQF_mem_address_not_aligned* exception (impl. dep. #111-V9-Cs10a).

Implementation | Although the architecture provides for the
Note | *LDQF_mem_address_not_aligned* exception, Oracle SPARC
 Architecture 2015 implementations do not currently generate it.

- An STQF or STQFA instruction accessing an address that is word aligned but not quadword aligned may cause an *STQF_mem_address_not_aligned* exception (impl. dep. #112-V9-Cs10a).

Implementation | Although the architecture provides for the
Note | *STQF_mem_address_not_aligned* exception, Oracle SPARC
 Architecture 2015 implementations do not currently generate it.

6.3.1.2 Addressing Conventions

An Oracle SPARC Architecture virtual processor uses big-endian byte order for all instruction accesses and, by default, for data accesses. It is possible to access data in little-endian format by use of selected ASIs. It is also possible to change the default byte order for implicit data accesses. See *Processor State (pstateP) Register (PR 6)* on page 69 for more information.¹

Big-endian Addressing Convention. Within a multiple-byte integer, the byte with the smallest address is the most significant; a byte’s significance decreases as its address increases. The big-endian addressing conventions are described in TABLE 6-1 and illustrated in FIGURE 6-2.

TABLE 6-1 Big-endian Addressing Conventions

Term	Definition
byte	A load/store byte instruction accesses the addressed byte in both big- and little-endian modes.
halfword	For a load/store halfword instruction, two bytes are accessed. The most significant byte (bits 15–8) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 1.
word	For a load/store word instruction, four bytes are accessed. The most significant byte (bits 31–24) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 3.
doubleword or extended word	For a load/store extended or floating-point load/store double instruction, eight bytes are accessed. The most significant byte (bits 63:56) is accessed at the address specified in the instruction; the least significant byte (bits 7:0) is accessed at the address + 7. For the deprecated integer load/store twin word instructions (LDTW, LDTWA [†] , STTW, STTWA), two big-endian words are accessed. The word at the address specified in the instruction corresponds to the even register specified in the instruction; the word at address + 4 corresponds to the following odd-numbered register. [†] Note that the LDTXA instruction, which is not an LDTWA operation but does share LDTWA’s opcode, is <i>not</i> deprecated.
quadword	For a load/store quadword instruction, 16 bytes are accessed. The most significant byte (bits 127–120) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 15.

¹ Readers interested in more background information on big- vs. little-endian can also refer to Cohen, D., “On Holy Wars and a Plea for Peace,” *Computer* 14:10 (October 1981), pp. 48-54.




FIGURE 6-2 Big-endian Addressing Conventions

Little-endian Addressing Convention. Within a multiple-byte integer, the byte with the smallest address is the least significant; a byte’s significance increases as its address increases. The little-endian addressing conventions are defined in TABLE 6-2 and illustrated in FIGURE 6-3.

TABLE 6-2 Little-endian Addressing Convention

Term	Definition
byte	A load/store byte instruction accesses the addressed byte in both big- and little-endian modes.
halfword	For a load/store halfword instruction, two bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 15–8) is accessed at the address + 1.
word	For a load/store word instruction, four bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 31–24) is accessed at the address + 3.
doubleword or extended word	<p>For a load/store extended or floating-point load/store double instruction, eight bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 63–56) is accessed at the address + 7.</p> <p>For the deprecated integer load/store twin word instructions (LDTW, LDTWA[†], STTW, STTWA), two little-endian words are accessed. The word at the address specified in the instruction corresponds to the even register in the instruction; the word at the address specified in the instruction +4 corresponds to the following odd-numbered register. With respect to little-endian memory, an LDTW/LDTWA (STTW/STTWA) instruction behaves as if it is composed of two 32-bit loads (stores), each of which is byte-swapped independently before being written into each destination register (memory word).</p> <p>[†]Note that the LDTXA instruction, which is not an LDTWA operation but does share LDTWA’s opcode, is <i>not</i> deprecated.</p>
quadword	For a load/store quadword instruction, 16 bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 127–120) is accessed at the address + 15.

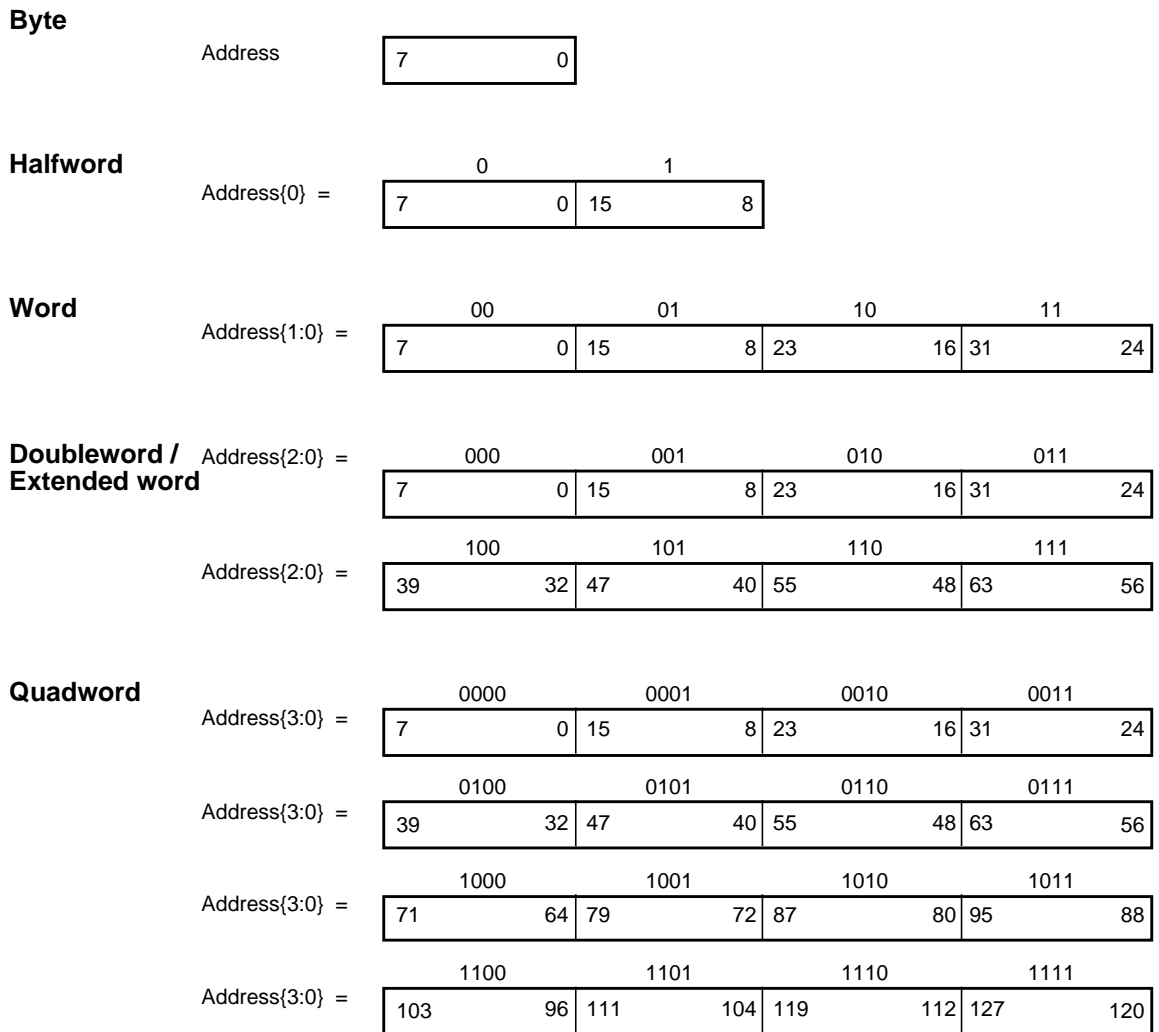


FIGURE 6-3 Little-endian Addressing Conventions

6.3.1.3 Address Space Identifiers (ASIs)

Alternate-space load, store, and load-store instructions specify an *explicit* ASI to use for their data access; when $i = 0$, the explicit ASI is provided in the instruction's `imm_asi` field, and when $i = 1$, it is provided in the ASI register.

Non-alternate-space load, store, and load-store instructions use an *implicit* ASI value that depends on the current trap level (TL) and the value of `PSTATE.cle`. Instruction fetches use an implicit ASI that depends only on the current trap level. The cases are enumerated in TABLE 6-3.

TABLE 6-3 ASIs Used for Data Accesses and Instruction Fetches

Access Type	TL	PSTATE.cle	ASI Used
Instruction Fetch	= 0	any	ASI_PRIMARY
	> 0	any	ASI_NUCLEUS*

TABLE 6-3 ASIs Used for Data Accesses and Instruction Fetches

Access Type	TL	PSTATE.cle	ASI Used
Non-alternate-space Load, Store, or Load-Store (implicit ASI)	= 0	0	ASI_PRIMARY
		1	ASI_PRIMARY_LITTLE
	> 0	0	ASI_NUCLEUS*
		1	ASI_NUCLEUS_LITTLE**
Alternate-space Load, Store, or Load-Store	any	any	ASI explicitly specified in the instruction (subject to privilege-level restrictions)

* On some early SPARC V9 implementations, ASI_PRIMARY may have been used for this case.

** On some early SPARC V9 implementations, ASI_PRIMARY_LITTLE may have been used for this case.

OSA 2015 See also *Memory Addressing and Alternate Address Spaces* on page 409.

ASIs 00₁₆-7F₁₆ are restricted; only software with sufficient privilege is allowed to access them. An attempt to access a restricted ASI by insufficiently-privileged software results in a *privileged_action* exception (impl. dep #103-V9-Ms10(6)). ASIs 80₁₆ through FF₁₆ are unrestricted; software is allowed to access them regardless of the virtual processor's privilege mode, as summarized in TABLE 6-4.

TABLE 6-4 Allowed Accesses to ASIs

Value	Access Type	Processor Mode (PSTATE.priv)	Result of ASI Access
00 ₁₆ -7F ₁₆	Restricted	Nonprivileged (0)	<i>privileged_action</i> exception
		Privileged (1)	Valid access
80 ₁₆ -FF ₁₆	Unrestricted	Nonprivileged (0)	Valid access
		Privileged (1)	Valid access

IMPL. DEP. #29-V8: Some Oracle SPARC Architecture 2015 ASIs are implementation dependent. See TABLE 10-1 on page 425 for details.

V9 Compatibility Note | In SPARC V9, many ASIs were defined to be implementation dependent.

An Oracle SPARC Architecture implementation decodes all 8 bits of ASI specifiers (impl. dep. #30-V8-Cu3).

V9 Compatibility Note | In SPARC V9, an implementation could choose to decode only a subset of the 8-bit ASI specifier.

6.3.1.4 Separate Instruction Memory

A SPARC V9 implementation may choose to access instruction and data through the same address space and use hardware to keep data and instruction memory consistent at all times. It may also choose to overload independent address spaces for data and instructions and allow them to become inconsistent when data writes are made to addresses shared with the instruction space.

Programming Note | A SPARC V9 program containing self-modifying code should use FLUSH instruction(s) after executing stores to modify instruction memory and before executing the modified instruction(s), to ensure the consistency of program execution.

6.3.2 Memory Synchronization Instructions

Two forms of memory barrier (MEMBAR) instructions allow programs to manage the order and completion of memory references. Ordering MEMBARs induce a partial ordering between sets of loads and stores and future loads and stores. Sequencing MEMBARs exert explicit control over completion of loads and stores (or other instructions). Both barrier forms are encoded in a single instruction, with subfunctions bit-encoded in cmask and mmask fields.

6.3.3 Integer Arithmetic and Logical Instructions

The integer arithmetic and logical instructions generally compute a result that is a function of two source operands and either write the result in a third (destination) register R[rd] or discard it. The first source operand is R[rs1]. The second source operand depends on the i bit in the instruction; if i = 0, then the second operand is R[rs2]; if i = 1, then the second operand is the constant simm10, simm11, or simm13 from the instruction itself, sign-extended to 64 bits.

Note | The value of R[0] always reads as zero, and writes to it are ignored.

6.3.3.1 Setting Condition Codes

Most integer arithmetic instructions have two versions: one sets the integer condition codes (icc and xcc) as a side effect; the other does not affect the condition codes. A special comparison instruction for integer values is not needed since it is easily synthesized with the “subtract and set condition codes” (SUBcc) instruction. See *Synthetic Instructions* on page 516 for details.

6.3.3.2 Shift Instructions

Shift instructions shift an R register left or right by a constant or variable amount. None of the shift instructions change the condition codes.

6.3.3.3 Set High 22 Bits of Low Word

The “set high 22 bits of low word of an R register” instruction (SETHI) writes a 22-bit constant from the instruction into bits 31 through 10 of the destination register. It clears the low-order 10 bits and high-order 32 bits, and it does not affect the condition codes. Its primary use is to construct constants in registers.

6.3.3.4 Integer Multiply/Divide

The integer multiply instruction performs a $64 \times 64 \rightarrow 64$ -bit operation; the integer divide instructions perform $64 \div 64 \rightarrow 64$ -bit operations. For compatibility with SPARC V8 processors, $32 \times 32 \rightarrow 64$ -bit multiply instructions, $64 \div 32 \rightarrow 32$ -bit divide instructions, and the Multiply Step instruction are provided. Division by zero causes a *division_by_zero* exception.

6.3.3.5 Tagged Add/Subtract

The tagged add/subtract instructions assume tagged-format data, in which the tag is the two low-order bits of each operand. If either of the two operands has a nonzero tag or if 32-bit arithmetic overflow occurs, tag overflow is detected. If tag overflow occurs, then TADDcc and TSUBcc set the CCR.icc.v bit; if 64-bit arithmetic overflow occurs, then they set the CCR.xcc.v bit.

The trapping versions (TADDccTV, TSUBccTV) of these instructions are deprecated. See *Tagged Add* on page 362 and *Tagged Subtract* on page 367 for details.

6.3.4 Control-Transfer Instructions (CTIs)

The basic control-transfer instruction types are as follows:

- Conditional branch (Bicc, BPcc, BPr, FBfcc, FBPfcc)
- Compare and Branch (C*Bcond)
- Unconditional branch
- Call and link (CALL)
- Jump and link (JEMPL, RETURN)
- Return from trap (DONE, RETRY)
- Trap (Tcc)
-

A control-transfer instruction functions by changing the value of the next program counter (NPC) or by changing the value of both the program counter (PC) and the next program counter (NPC). When only NPC is changed, the effect of the transfer of control is delayed by one instruction. Most control transfers are of the delayed variety. The instruction following a delayed control-transfer instruction is said to be in the *delay slot* of the control-transfer instruction.

Some control transfer instructions (branches) can optionally annul, that is, not execute, the instruction in the delay slot, based on the setting of an *annul bit* in the instruction. The effect of the annul bit depends upon whether the transfer is taken or not taken and whether the branch is conditional or unconditional. Annulled delay instructions neither affect the program-visible state, nor can they cause a trap.

Programming Note	<p>The annul bit increases the likelihood that a compiler can find a useful instruction to fill the delay slot after a branch, thereby reducing the number of instructions executed by a program. For example, the annul bit can be used to move an instruction from within a loop to fill the delay slot of the branch that closes the loop.</p> <p>Likewise, the annul bit can be used to move an instruction from either the “else” or “then” branch of an “if-then-else” program block to the delay slot of the branch that selects between them. Since a full set of conditions is provided, a compiler can arrange the code (possibly reversing the sense of the condition) so that an instruction from either the “else” branch or the “then” branch can be moved to the delay slot. Use of annulled branches provided some benefit in older, single-issue SPARC implementations. On an Oracle SPARC Architecture implementation, the only benefit of annulled branches might be a slight reduction in code size. Therefore, the use of annulled branch instructions is no longer encouraged.</p>
-------------------------	--

TABLE 6-5 defines the value of the program counter and the value of the next program counter after execution of each instruction. Conditional branches have two forms: branches that test a condition (including branch-on-register), represented in the table by Bcc, and branches that are unconditional, that is, always or never taken, represented in the table by BA and BN, respectively. The effect of an annulled branch is shown in the table through explicit transfers of control, rather than by fetching and annulling the instruction.

TABLE 6-5 Control-Transfer Characteristics (1 of 2)

Instruction Group	Address Form	Delayed?	Taken?	Annul Bit?	New PC	New NPC
Non-CTIs	—	—	—	—	NPC	NPC + 4
Bcc	PC-relative	Yes	Yes	0	NPC	EA
Bcc	PC-relative	Yes	No	0	NPC	NPC + 4
Bcc	PC-relative	Yes	Yes	1	NPC	EA

TABLE 6-5 Control-Transfer Characteristics (Continued) (2 of 2)

Instruction Group	Address Form	Delayed?	Taken?	Annul Bit?	New PC	New NPC
Bcc	PC-relative	Yes	No	1	NPC + 4	NPC + 8
BA	PC-relative	Yes	Yes	0	NPC	EA
BA	PC-relative	No	Yes	1	EA	EA + 4
BN	PC-relative	Yes	No	0	NPC	NPC + 4
BN	PC-relative	Yes	No	1	NPC + 4	NPC + 8
CALL	PC-relative	Yes	—	—	NPC	EA
C*Bcond	PC-relative	No	Yes	—	EA	EA + 4
C*Bcond	PC-relative	No	No	—	NPC	NPC + 4
JMPL, RETURN	Register-indirect	Yes	—	—	NPC	EA
DONE	Trap state	No	—	—	TNPC[TL]	TNPC[TL] + 4
RETRY	Trap state	No	—	—	TPC[TL]	TNPC[TL]
Tcc	Trap vector	No	Yes	—	EA	EA + 4
Tcc	Trap vector	No	No	—	NPC	NPC + 4

The effective address, “EA” in TABLE 6-5, specifies the target of the control-transfer instruction. The effective address is computed in different ways, depending on the particular instruction.

- **PC-relative effective address** — A PC-relative effective address is computed by sign extending the instruction’s immediate field to 64-bits, left-shifting the word displacement by 2 bits to create a byte displacement, and adding the result to the contents of the PC.
- **Register-indirect effective address** — If $i = 0$, a register-indirect effective target address is $R[rs1] + R[rs2]$. If $i = 1$, a register-indirect effective target address is $R[rs1] + \text{sign_ext}(imm13)$.
- **Trap vector effective address** — A trap vector effective address first computes the software trap number as the least significant 7 or 8 bits of $R[rs1] + R[rs2]$ if $i = 0$, or as the least significant 7 or 8 bits of $R[rs1] + imm_trap\#$ if $i = 1$. Whether 7 or 8 bits are used depends on the privilege level — 7 bits are used in nonprivileged mode and 8 bits are used in privileged mode. The trap level, TL, is incremented. The hardware trap type is computed as $256 +$ the software trap number and stored in $TT[TL]$. The effective address is generated by combining the contents of the TBA register with the trap type and other data; see *Trap Processing* on page 458 for details.
- **Trap state effective address** — A trap state effective address is not computed but is taken directly from either $TPC[TL]$ or $TNPC[TL]$.

SPARC V8 Compatibility Note | The SPARC V8 architecture specified that the delay instruction was always fetched, even if annulled, and that an annulled instruction could not cause any traps. The SPARC V9 architecture does not require the delay instruction to be fetched if it is annulled.

6.3.4.1 Conditional Branches

A conditional branch transfers control if the specified condition is TRUE. If the annul bit is 0, the instruction in the delay slot is always executed. If the annul bit is 1, the instruction in the delay slot is executed only when the conditional branch is taken.

Note | The annulling behavior of a taken conditional branch is different from that of an unconditional branch.

6.3.4.2 Unconditional Branches

An unconditional branch transfers control unconditionally if its specified condition is “always”; it never transfers control if its specified condition is “never.” If the annul bit is 0, then the instruction in the delay slot is always executed. If the annul bit is 1, then the instruction in the delay slot is *never* executed.

Note | The annul behavior of an unconditional branch is different from that of a taken conditional branch.

6.3.4.3 CALL and JMPL Instructions

The CALL instruction writes the contents of the PC, which points to the CALL instruction itself, into R[15] (*out* register 7) and then causes a delayed transfer of control to a PC-relative effective address. The value written into R[15] is visible to the instruction in the delay slot.

The JMPL instruction writes the contents of the PC, which points to the JMPL instruction itself, into R[rd] and then causes a register-indirect delayed transfer of control to the address given by “R[rs1] + R[rs2]” or “R[rs1] + a signed immediate value.” The value written into R[rd] is visible to the instruction in the delay slot.

When PSTATE.am = 1, the value of the high-order 32 bits transmitted to R[15] by the CALL instruction or to R[rd] by the JMPL instruction is zero.

6.3.4.4 RETURN Instruction

The RETURN instruction is used to return from a trap handler executing in nonprivileged mode. RETURN combines the control-transfer characteristics of a JMPL instruction with R[0] specified as the destination register and the register-window semantics of a RESTORE instruction.

6.3.4.5 DONE and RETRY Instructions

The DONE and RETRY instructions are used by privileged software to return from a trap. These instructions restore the machine state to values saved in the TSTATE register stack.

RETRY returns to the instruction that caused the trap in order to reexecute it. DONE returns to the instruction pointed to by the value of NPC associated with the instruction that caused the trap, that is, the next logical instruction in the program. DONE presumes that the trap handler did whatever was requested by the program and that execution should continue.

6.3.4.6 Trap Instruction (Tcc)

The Tcc instruction initiates a trap if the condition specified by its cond field matches the current state of the condition code specified in its cc field; otherwise, it executes as a NOP. If the trap is taken, it increments the TL register, computes a trap type that is stored in TT[TL], and transfers to a computed address in a trap table pointed to by a trap base address register.

A Tcc instruction can specify one of 256 software trap types (128 when in nonprivileged mode). When a Tcc is taken, 256 plus the 7 (in nonprivileged mode) or 8 (in privileged mode) least significant bits of the Tcc’s second source operand are written to TT[TL]. The only visible difference between a software trap generated by a Tcc instruction and a hardware trap is the trap number in the TT register. See Chapter 12, *Traps*, for more information.

Programming Note | Tcc can be used to implement breakpointing, tracing, and calls to privileged or hyperprivileged software. Tcc can also be used for runtime checks, such as out-of-range array index checks or integer overflow checks.

6.3.4.7 DCTI Couples (E2)

A delayed control transfer instruction (DCTI) in the delay slot of another DCTI is referred to as a “DCTI couple”. The use of DCTI couples is deprecated in the Oracle SPARC Architecture; no new software should place a DCTI in the delay slot of another DCTI, because on future Oracle SPARC Architecture implementations DCTI couples may execute either slowly or differently than the programmer assumes it will.

SPARC V8 and SPARC V9 Compatibility Note	The SPARC V8 architecture left behavior undefined for a DCTI couple. The SPARC V9 architecture defined behavior in that case, but as of UltraSPARC Architecture 2005, <i>use of DCTI couples was deprecated</i> . Software should not expect high performance from DCTI couples, and performance of DCTI couples should be expected to decline further in future processors.
---	--

Programming Note	As noted in TABLE 6-5 on page 86, an annulled branch-always (branch-always with <code>a = 1</code>) instruction is not architecturally a DCTI. However, since not all implementations make that distinction, for optimal performance, a DCTI should not be placed in the instruction word immediately following an annulled branch-always instruction (BA,A or BPA,A).
-----------------------------	---

6.3.5 Conditional Move Instructions

This subsection describes two groups of instructions that copy or move the contents of any integer or floating-point register.

MOVcc and FMOVcc Instructions. The MOVcc and FMOVcc instructions copy the contents of any integer or floating-point register to a destination integer or floating-point register if a condition is satisfied. The condition to test is specified in the instruction and can be any of the conditions allowed in conditional delayed control-transfer instructions. This condition is tested against one of the six sets of condition codes (`icc`, `xcc`, `fcc0`, `fcc1`, `fcc2`, and `fcc3`), as specified by the instruction. For example:

```
fmovdgd          %fcc2, %f20, %f22
```

moves the contents of the double-precision floating-point register `%f20` to register `%f22` if floating-point condition code number 2 (`fcc2`) indicates a greater-than relation (`FSR.fcc2 = 2`). If `fcc2` does not indicate a greater-than relation (`FSR.fcc2 ≠ 2`), then the move is not performed.

The MOVcc and FMOVcc instructions can be used to eliminate some branches in programs. In most implementations, branches will be more expensive than the MOVcc or FMOVcc instructions. For example, the C statement:

```
if (A > B) X = 1; else X = 0;
```

can be coded as

```
cmp          %i0, %i2      ! (A > B)
or           %g0, 0, %i3   ! set X = 0
movg        %xcc, 1, %i3  ! overwrite X with 1 if A > B
```

to eliminate the need for a branch.

MOVr and FMOVr Instructions. The MOVr and FMOVr instructions allow the contents of any integer or floating-point register to be moved to a destination integer or floating-point register if the contents of a register satisfy a specified condition. The conditions to test are enumerated in TABLE 6-6.

TABLE 6-6 MOVr and FMOVr Test Conditions

Condition	Description
NZ	Nonzero
Z	Zero
GEZ	Greater than or equal to zero
LZ	Less than zero
LEZ	Less than or equal to zero
GZ	Greater than zero

Any of the integer registers (treated as a signed value) may be tested for one of the conditions, and the result used to control the move. For example,

```
movrnz    %i2, %i4, %i6
```

moves integer register %i4 to integer register %i6 if integer register %i2 contains a nonzero value.

MOVr and FMOVr can be used to eliminate some branches in programs or can emulate multiple unsigned condition codes by using an integer register to hold the result of a comparison.

6.3.6 Register Window Management Instructions

This subsection describes the instructions that manage register windows in the Oracle SPARC Architecture. The privileged registers affected by these instructions are described in *Register-Window PR State Registers* on page 62.

6.3.6.1 SAVE Instruction

The SAVE instruction allocates a new register window and saves the caller's register window by incrementing the CWP register.

If CANSAVE = 0, then execution of a SAVE instruction causes a window spill exception, that is, one of the *spill_n_<normal|other>* exceptions.

If CANSAVE ≠ 0 but the number of clean windows is zero, that is, (CLEANWIN – CANRESTORE) = 0, then SAVE causes a *clean_window* exception.

If SAVE does not cause an exception, it performs an ADD operation, decrements CANSAVE, and increments CANRESTORE. The source registers for the ADD operation are from the old window (the one to which CWP pointed before the SAVE), while the result is written into a register in the new window (the one to which the incremented CWP points).

6.3.6.2 RESTORE Instruction

The RESTORE instruction restores the previous register window by decrementing the CWP register.

If CANRESTORE = 0, execution of a RESTORE instruction causes a window fill exception, that is, one of the *fill_n_<normal|other>* exceptions.

If RESTORE does not cause an exception, it performs an ADD operation, decrements CANRESTORE, and increments CANSAVE. The source registers for the ADD are from the old window (the one to which CWP pointed before the RESTORE), and the result is written into a register in the new window (the one to which the decremented CWP points).

Programming Note	<p>This note describes a common convention for use of register windows, SAVE, RESTORE, CALL, and JMPL instructions.</p> <p>A procedure is invoked by execution of a CALL (or a JMPL) instruction. If the procedure requires a register window, it executes a SAVE instruction in its prologue code. A routine that does not allocate a register window of its own (possibly a leaf procedure) should not modify any windowed registers except <i>out</i> registers 0 through 6. This optimization, called “Leaf-Procedure Optimization”, is routinely performed by SPARC compilers.</p> <p>A procedure that uses a register window returns by executing both a RESTORE and a JMPL instruction. A procedure that has not allocated a register window returns by executing a JMPL only. The target address for the JMPL instruction is normally 8 plus the address saved by the calling instruction, that is, the instruction after the instruction in the delay slot of the calling instruction.</p> <p>The SAVE and RESTORE instructions can be used to atomically establish a new memory stack pointer in an R register and switch to a new or previous register window.</p>
-------------------------	---

6.3.6.3 SAVED Instruction

SAVED is a privileged instruction used by a spill trap handler to indicate that a window spill has completed successfully. It increments CANSAVE and decrements either OTHERWIN or CANRESTORE, depending on the conditions at the time SAVED is executed.

See *SAVED* on page 324 for details.

6.3.6.4 RESTORED Instruction

RESTORED is a privileged instruction, used by a fill trap handler to indicate that a window has been filled successfully. It increments CANRESTORE and decrements either OTHERWIN or CANSAVE, depending on the conditions at the time RESTORED is executed. RESTORED also manipulates CLEANWIN, which is used to ensure that no address space’s data become visible to another address space through windowed registers.

See *RESTORED* on page 317 for details.

6.3.6.5 Flush Windows Instruction

The FLUSHW instruction flushes all of the register windows, except the current window, by performing repetitive spill traps. The FLUSHW instruction causes a spill trap if any register window (other than the current window) has valid contents. The number of windows with valid contents is computed as:

$$N_REG_WINDOWS - 2 - CANSAVE$$

If this number is nonzero, the FLUSHW instruction causes a spill trap. Otherwise, FLUSHW has no effect. If the spill trap handler exits with a RETRY instruction, the FLUSHW instruction continues causing spill traps until all the register windows except the current window have been flushed.

6.3.7 Ancillary State Register (ASR) Access

The read/write state register instructions access program-visible state and status registers. These instructions read/write the state registers into/from R registers. A read/write Ancillary State register instruction is privileged only if the accessed register is privileged.

The supported RDasr and WRasr instructions are described in *Ancillary State Registers* on page 48.

6.3.8 Privileged Register Access

The read/write privileged register instructions access state and status registers that are visible only to privileged software. These instructions read/write privileged registers into/from R registers. The read/write privileged register instructions are privileged.

6.3.9 Floating-Point Operate (FPop) Instructions

Floating-point operate instructions (FPops) compute a result that is a function of one, two, or three source operands and place the result in one or more destination F registers, with one exception: floating-point compare operations do not write to an F register but instead update one of the *fccn* fields of the FSR.

The term “FPop” refers to instructions in the FPop1, FMAf, and FPop2 opcode spaces. FPop instructions do not include FBfcc instructions, loads and stores between memory and the F registers, or non-floating-point operations that read or write F registers.

The FMOVcc instructions function for the floating-point registers as the MOVcc instructions do for the integer registers. See *MOVcc and FMOVcc Instructions* on page 89.

The FMOVr instructions function for the floating-point registers as the MOVr instructions do for the integer registers. See *MOVr and FMOVr Instructions* on page 90.

If no floating-point unit is present or if PSTATE.pef = 0 or FPRS.fef = 0, then any instruction, including an FPop instruction, that attempts to access an FPU register generates an *fp_disabled* exception.

All FPop instructions clear the *ftt* field and set the *cexc* field unless they generate an exception. Floating-point compare instructions also write one of the *fccn* fields. All FPop instructions that can generate IEEE exceptions set the *cexc* and *aexc* fields unless they generate an exception. FABS<sd|q>, FMOV<sd|q>, FMOVcc<sd|q>, FMOVr<sd|q>, and FNEG<sd|q> cannot generate IEEE exceptions, so they clear *cexc* and leave *aexc* unchanged.

IMPL. DEP. #3-V8: An implementation may indicate that a floating-point instruction did not produce a correct IEEE Std 754-1985 result by generating an *fp_exception_other* exception with FSR.ftt = unfinished_FPop. In this case, software running in a mode with greater privileges must emulate any functionality not present in the hardware.

See *ftt = 2 (unfinished_FPop)* on page 45 to see which instructions can produce an *fp_exception_other* exception (with FSR.ftt = unfinished_FPop).

6.3.10 Implementation-Dependent Instructions

The SPARC V9 architecture provided two instruction spaces that are entirely implementation dependent: IMPDEP1 and IMPDEP2 .

In the Oracle SPARC Architecture, the IMPDEP1 opcode space is used by many VIS instructions. The IMPDEP2B opcode space is primarily used for implementation of floating-point multiply-add/multiply-subtract instructions. The remaining opcodes in IMPDEP1 and IMPDEP2 are now marked as reserved opcodes.

6.3.11 Reserved Opcodes and Instruction Fields

If a conforming Oracle SPARC Architecture 2015 implementation attempts to execute an instruction bit pattern that is not specifically defined in this specification, it behaves as follows:

- If the instruction bit pattern encodes an implementation-specific extension to the instruction set, that extension is executed.
- If the instruction does not encode an extension to the instruction set, then the instruction bit pattern is invalid and causes an *illegal_instruction* exception.

See Appendix A, *Opcode Maps*, for an enumeration of the reserved instruction bit patterns (opcodes).

Programming Note	For software portability, software (such as assemblers, static compilers, and dynamic compilers) that generates SPARC instructions must always generate zeroes in instruction fields marked “reserved” (“—”).
-------------------------	---

Instructions

Oracle SPARC Architecture 2015 extends the standard SPARC V9 instruction set with additional classes of instructions:

- Enhanced functionality:
 - Instructions for alignment (*Align Address* on page 117)
 - Array handling (*Three-Dimensional Array Addressing* on page 120)
 - Byte-permutation instructions (*Byte Mask and Shuffle* on page 125 and *CMASK* on page 139)
 - Edge handling (*Edge Handling Instructions* on pages 149 and 150)
 - Logical operations on floating-point registers (*f Register Logical Operate (1 operand)* on page 225)
 - Partitioned arithmetic (*Partitioned Add* on page 201, *Partitioned Add with Saturation* on page 204, or *pstate.pef = 0Partitioned Subtract* on page 219, and *Partitioned Subtract with Saturation* on page 222)
 - Pixel manipulation (*FEXPAND* on page 165, *FPACK* on page 197, and *FPMERGE* on page 216)
- Efficient memory access
 - Partial store (*DAE_invalid_asi eDAE_invalid_asi eStore Partial Floating-Point* on page 347)
 - Short floating-point loads and stores (*Store Short Floating-Point* on page 350)
 - Block load and store (*Block Load* on page 243 and *Block Store* on page 337)
- Efficient interval arithmetic: SIAM (*Set Interval Arithmetic Mode* on page 330) and all instructions that reference GSR.im
- Floating-point Multiply-Add and Multiply-Subtract (FMA) instructions (*Floating-Point Multiply-Add and Multiply-Subtract (fused)* on page 175)
- Direct moves between integer and floating-point registers: *Move Floating-Point Register to Integer Register* on page 284 and *Move Integer Register to Floating-Point Register* on page 285
- Efficient fused compare-and-branch instructions: *Compare and Branch* on page 136
- Cryptographic algorithm instructions, such as the AES instructions starting on page 112, the Camellia instructions starting on page 131, the DES instructions starting on page 143, Montgomery Multiplication and Squaring instructions on pages 272 and 276, and Multiple-Precision Multiply on page 286

TABLE 7-5 provides a quick index of instructions, alphabetically by architectural instruction name.

TABLE 7-4 summarizes the instruction set, listed within functional categories.

Within these tables and throughout the rest of this chapter, and in Appendix A, *Opcode Maps*, certain opcodes are marked with mnemonic superscripts. The superscripts and their meanings are defined in TABLE 7-1.

TABLE 7-1 Instruction Superscripts

Superscript	Meaning
D	Deprecated instruction (do not use in new software)
H _{dis}	Privileged action if in nonprivileged or privileged mode and access is disabled
N	Nonportable instruction
P	Privileged instruction
P _{ASI}	Privileged action if bit 7 of the referenced ASI is 0
P _{ASR}	Privileged instruction if the referenced ASR register is privileged
P _{dis}	Privileged action if in nonprivileged mode (PSTATE.priv = 0) and nonprivileged access is disabled

TABLE 7-2 Oracle SPARC Architecture 2015 Instruction Set - Alphabetical (1 of 3)

Page	Instruction	Page	Instruction	Page	Instruction
110	ADD (ADDcc)	152	FABS<s d q>	227	FORNOT<1 2><s d>
110	ADDC (ADDCcc)	153	FADD<s d q>	227	FOR<s d>
111	ADDXC (ADDXCcc) ^N	154	FALIGNDATAg	197	FPACK<16 32 FIX>
112	AES_DROUND<01 23>[_LAST] ^N	155	FALIGNDATAi	201	FPADD<8 16 32 64>
112	AES_EROUND<01 23>[_LAST] ^N	227	FANDNOT<1 2><s d>	201	FPADD<16,32>s
115	AES_KEXPAND0 ^N	227	FAND<s d>	204	FPADD[U]S<16,32>[s]
112	AES_KEXPAND1 ^N	157	FBfcc ^D	207	FPCMP
115	AES_KEXPAND2 ^N	159	FBPfcc		
117	ALIGNADDRESS[_LITTLE]	161	FCHKSM16		
118	ALLCLEAN	162	FCMP<s d q>	210	FPCMPU
		162	FCMPE<s d q>		
		164	FDIV<s d q>	213	FPMADDX[HI]
		190	FdMULq	214	FPMAX
		165	FEXPAND	214	FPMAXU
118	AND (ANDcc)	166	FHADD<s d>	216	FPMERGE
118	ANDN (ANDNcc)	167	FHSUB<s d>	217	FPMIN
120	ARRAY<8 16 32>	168	FiTO<s d q>	217	FPMINU
123	Bicc	169	FLCMP<s d>		
125	BMASK	171	FLUSH		
126	BPcc	174	FLUSHW		
128	BPr	175	FMADD<s d>		
125	BSHUFFLE	177	FMEAN16		
130	CALL	179	FMOV<s d q>	219	FPSUB
131	CAMELLIA_F ^N	180	FMOV<s d q>cc	222	FPSUBS
133	CAMELLIA_FL ^N	184	FMOV<s d q>R		
133	CAMELLIA_FLI ^N	175	FMSUB<s d>	228	FS<LL RL RA><16 32>
134	CASA ^{PASI}	190	FMUL<s d q>	190	FsMULd
134	CASXA ^{PASI}	186	FMUL8[SU UL]x16	230	FSQRT<s d q>
136	CBcond	186	FMUL8x16	226	FSRC<1 2><s d>
139	CMASK<8 16 32> ^N	186	FMUL8x16[AU AL]	234	FSUB<s d q>
141	CRC32C ^N	186	FMULD8[SU UL]x16	227	FXNOR<s d>
145	DES_IP ^N	191	FNADD	227	FXOR<s d>
145	DES_IIP ^N	227	FNAND<s d>	235	FxTO<s d q>
145	DES_KEXPAND ^N	193	FNEG<s d q>	225	FZERO<s d>
143	DES_ROUND ^N	194	FNHADD<s d>		
		175	FNHADD<s d>	236	ILLTRAP
147	DONE ^P	175	FNMSUB<s d>	237	INVALW
149	EDGE<8 16 32>[L]cc	195	FNMUL<s d>	238	JMPL
150	EDGE<8 16 32>[L]N	227	FNOR<s d>		
232	F<s d q>TO<s d q>	226	FNOT<1 2><s d>	243	LDBLOCKF ^D
231	F<s d q>TOi	195	FNsMULd	246	LDDF
231	F<s d q>TOx	225	FONE<s d>	248	LDDFA ^{PASI}

TABLE 7-2 Oracle SPARC Architecture 2015 Instruction Set - Alphabetical (2 of 3)

Page	Instruction	Page	Instruction	Page	Instruction
246	LDF				
248	LDFA ^{PASI}			330	SIAM
251	LDFSR ^D				
246	LDQF			331	SLL
248	LDQFA ^{PASI}	296	OR (ORcc)	331	SLLX
		296	ORN (ORNcc)	333	SMUL ^D (SMULcc ^D)
		297	OTHERW	331	SRA
		298	PAUSE	331	SRAX
		299	PDIST ^D	331	SRL
239	LDSB	300	PDISTN	331	SRLX
240	LDSBA ^{PASI}	301	POPC	334	STB
239	LDSH	303	PREFETCH	335	STBA ^{PASI}
240	LDSHA ^{PASI}	303	PREFETCHA ^{PASI}	479	STBAR ^D
253	LDSHORTF	310	RDASI	337	STBLOCKF
255	LDSTUB	310	RDAsr ^{PASR}	340	STDF
256	LDSTUBA ^{PASI}	310	RDCCR	342	STDFA ^{PASI}
239	LDSW	310	RDCFR	340	STF
240	LDSWA ^{PASI}			342	STFA ^{PASI}
262	LDTXA ^N	310	RDFPRS	345	STFSR ^D
257	LDTW ^D	310	RDGSR	334	STH
259	LDTWA ^{D, PASI}			335	STHA ^{PASI}
239	LDUB	310	RDPC		
240	LDUBA ^{PASI}				
239	LDUH				
240	LDUHA ^{PASI}	313	RDPR ^P		
239	LDUW	310	RDSOFTINT ^P	347	STPARTIALF
240	LDUWA ^{PASI}	310	RDSTICK ^{Pdis,Hdis}	340	STQF
239	LDX	310	RDSTICK_CMPR ^P	350	STSHORTF
240	LDXA ^{PASI}	310	RDTICK ^{Pdis,Hdis}	342	STQFA ^{PASI}
264	LDXEFSR	315	RESTORE	352	STTW ^D
264	LDXFSR	317	RESTORED ^P	354	STTWA ^{D, PASI}
266	LZCNT	318	RETRY ^P	334	STW
268	MD5 ^N	320	RETURN	335	STWA ^{PASI}
269	MEMBAR			334	STX
272	MONTMUL ^N			335	STXA ^{PASI}
276	MONTSQR ^N			356	STXFSR
280	MOVcc			357	SUB (SUBcc)
284	MOVfTOi	324	SAVED ^P	357	SUBC (SUBCcc)
285	MOViTOf	322	SAVE	358	SUBX (SUBXcc)
283	MOVr	325	SDIV ^D (SDIVcc ^D)	360	SWAPA ^{D, PASI}
286	MPMUL ^N	291	SDIVX	359	SWAP ^D
406	MULScc ^D	327	SETHI	362	TADDcc
291	MULX	328	SHA1 ^N	363	TADDccTV ^D
292	MWAIT	328	SHA256 ^N	364	Tcc
294	NOP			367	TSUBcc
295	NORMALW	328	SHA512 ^N	368	TSUBccTV ^D
369	UDIV ^D (UDIVcc ^D)	373	WRGSR	373	WRSTICK_CMPR ^P

TABLE 7-2 Oracle SPARC Architecture 2015 Instruction Set - Alphabetical (3 of 3)

Page	Instruction	Page	Instruction	Page	Instruction
291	UDIVX			373	WRSTICK ^P
371	UMUL ^D (UMULcc ^D)	373	WRPAUSE	373	WRY ^D
372	UMULXHI			378	XMONTMUL ^N
373	WRASI			381	XMONTSQR ^N
373	WRasr ^{PASR}	536	WRPR ^P	384	XMPMUL ^N
373	WRCCR	373	WRSOFTINT_CLR ^P	389	XMULX[HI]N
		373	WRSOFTINT_SET ^P	390	XNOR (XNORcc)
373	WRFPRS	373	WRSOFTINT ^P	390	XOR (XORcc)

Oracle SPARC Architecture instructions are grouped into “feature sets”. Each set comprises a collection of architectural features that were introduced at the same time. Those feature sets are listed in the following table:

TABLE 7-3 Oracle SPARC Architecture Feature Sets

Architectural Feature Set Name	Sun Studio Compiler Option that enables generation of instructions in Feature Set (-xarch=)	SPARC Processor in which Feature Set was first Implemented (year)	Oracle SPARC Architecture Specification in which Feature Set first appeared
SPARC V9	(enabled by default)	UltraSPARC I (1995)	<i>The SPARC Architecture Manual-Version 9</i>
VIS 1	<code>sparcvis</code>	UltraSPARC I (1995)	UltraSPARC Architecture 2005
VIS 2	<code>sparcvis2</code>	UltraSPARC III (2001)	UltraSPARC Architecture 2005
FMAf	<code>sparcfmaf</code>	SPARC T3 (2010)	Oracle SPARC Architecture 2011
VIS 3, VIS 3B	<code>sparcvis3</code>	SPARC T3 (2010)	Oracle SPARC Architecture 2011
IMA	<code>sparcima</code>	SPARC T4 (2011)	Oracle SPARC Architecture 2011
SPARC4	<code>sparc4</code>	SPARC T4 (Core S3, 2011)	Oracle SPARC Architecture 2011
SPARC5, VIS 4	<code>sparc5</code>	SPARC M7/T7 (Core S4, 2015)	Oracle SPARC Architecture 2015

TABLE 7-4 Instruction Set - by Functional Category (1 of 7)

Instruction	Category and Function	Page	Added in
Data Movement Operations, Between R Registers			
MOVcc	Move integer register if condition is satisfied	280	
MOVr	Move integer register on contents of integer register	283	
Data Movement Operations, Between F Registers			
FMOV<sd q>	Floating-point move	179	
FMOV<sd q>cc	Move floating-point register if condition is satisfied	180	
FMOV<sd q>R	Move f-p reg. if integer reg. contents satisfy condition	184	
FSRC<1 2><sd>	Copy source to destination	226	VIS 1 (1995)
Data Movement Operations, Between R and F Registers			
MOVfToi	Move floating-point register to integer register	284	VIS 3 (2010)
MOViTof	Move integer register to floating-point register	285	VIS 3 (2010)
Data Conversion Instructions			
FiTO<sd q>	Convert 32-bit integer to floating-point	168	
F<sd q>TOi	Convert floating point to integer	231	
F<sd q>TOx	Convert floating point to 64-bit integer	231	
F<sd q>TO<sd q>	Convert between floating-point formats	232	
FxTO<sd q>	Convert 64-bit integer to floating-point	235	
Logical Operations on R Registers			
AND (ANDcc)	Logical and (and modify condition codes)	118	
OR (ORcc)	Inclusive- or (and modify condition codes)	296	
ORN (ORNcc)	Inclusive- or not (and modify condition codes)	296	
XNOR (XNORcc)	Exclusive- nor (and modify condition codes)	390	
XOR (XORcc)	Exclusive- or (and modify condition codes)	390	
Logical Operations on F Registers			
FAND<sd>	Logical and operation	227	VIS 1 (1995)
FANDNOT<1 2><sd>	Logical and operation with one inverted source	227	VIS 1 (1995)
FNAND<sd>	Logical nand operation	227	VIS 1 (1995)
FNOR<sd>	Logical nor operation	227	VIS 1 (1995)
FNOT<1 2><sd>	Copy negated source	226	VIS 1 (1995)
FONE<sd>	One fill	225	VIS 1 (1995)
FOR<sd>	Logical or operation	227	VIS 1 (1995)
FORNOT<1 2><sd>	Logical or operation with one inverted source	227	VIS 1 (1995)
FXNOR<sd>	Logical xnor operation	227	VIS 1 (1995)
FXOR<sd>	Logical xor operation	227	VIS 1 (1995)
FZERO<sd>	Zero fill	225	VIS 1 (1995)
Shift Operations on R Registers			
SLL	Shift left logical	331	
SLLX	Shift left logical, extended	331	
SRA	Shift right arithmetic	331	
SRAX	Shift right arithmetic, extended	331	
SRL	Shift right logical	331	
SRLX	Shift right logical, extended	331	

TABLE 7-4 Instruction Set - by Functional Category (2 of 7)

Instruction	Category and Function	Page	Added in
Special Addressing and Data Alignment Operations			
ALIGNADDRESS[_LITTLE]	Calculate address for misaligned data	117	VIS 1 (1995)
ARRAY<8 16 32>	3-D array addressing instructions	120	VIS 1 (1995)
FALIGNDATAg	Perform data alignment for misaligned data (using GSR.align)	154	VIS 1 (1995)
FALIGNDATAi	Perform data alignment for misaligned data (using integer register)	154	OSA 2015 (VIS 4)
Control Transfers			
Bicc	Branch on integer condition codes	123	
BPcc	Branch on integer condition codes with prediction	126	
BPr	Branch on contents of integer register with prediction	128	
CALL	Call and link	130	
CBcond	Compare and Branch	136	OSA 2011
DONE ^P	Return from trap	147	
FBfcc ^D	Branch on floating-point condition codes	157	
FBPfcc	Branch on floating-point condition codes with prediction	159	
ILLTRAP	Illegal instruction	236	
JMPL	Jump and link	238	
RETRY ^P	Return from trap and retry	318	
RETURN	Return	320	
Tcc	Trap on integer condition codes	364	
Byte Permutation			
BMASK	Set the GSR.mask field	125	VIS 2 (2001)
BSHUFFLE	Permute bytes as specified by GSR.mask	125	VIS 2 (2001)
CMASK<8 16 32> ^N	Create GSR.mask from SIMD comparison result	139	
Data Formatting Operations on F Registers			
FEXPAND	Pixel expansion	165	VIS 1 (1995)
FPACK<16 32 FIX>	Pixel packing	197	VIS 1 (1995)
FPMERGE	Pixel merge	216	VIS 1 (1995)
Memory Operations to/from F Registers			
LDBLOCKF ^D	Block loads	243	VIS 1 (1995)
STBLOCKF	Block stores	337	VIS 1 (1995)
LDDF	Load double floating-point	246	
LDDFA ^{PASI}	Load double floating-point from alternate space	248	
LDF	Load floating-point	246	
L DFA ^{PASI}	Load floating-point from alternate space	248	
LDQF	Load quad floating-point	246	
LDQFA ^{PASI}	Load quad floating-point from alternate space	248	
LDSHORTF	Short floating-point loads	253	VIS 1 (1995)
STDF	Store double floating-point	340	
STDFA ^{PASI}	Store double floating-point into alternate space	342	
STF	Store floating-point	340	
STFA ^{PASI}	Store floating-point into alternate space	342	
STPARTIALF	Partial Store instructions	347	VIS 1 (1995)
STQF	Store quad floating point	340	
STQFA ^{PASI}	Store quad floating-point into alternate space	342	

TABLE 7-4 Instruction Set - by Functional Category (3 of 7)

Instruction	Category and Function	Page	Added in
STSHORTF	Short floating-point stores	350	VIS 1 (1995)
<i>Atomic (Load-Store) Memory Operations to/from R Registers</i>			
CASA ^{PASI}	Compare and swap word in alternate space	134	
CASXA ^{PASI}	Compare and swap doubleword in alternate space	134	
LDSTUB	Load-store unsigned byte	255	
LDSTUBA ^{PASI}	Load-store unsigned byte in alternate space	256	
SWAP ^D	Swap integer register with memory	359	
SWAPA ^{D, PASI}	Swap integer register with memory in alternate space	360	
<i>Memory Operations to/from R Registers</i>			
LDSB	Load signed byte	239	
LDSBA ^{PASI}	Load signed byte from alternate space	240	
LDSH	Load signed halfword	239	
LDSHA ^{PASI}	Load signed halfword from alternate space	240	
LDSW	Load signed word	239	
LDSWA ^{PASI}	Load signed word from alternate space	240	
LDTXA ^N	Load integer twin extended word from alternate space	262	VIS 2 (2001)+
LDTW ^{D, PASI}	Load integer twin word	257	
LDTWA ^{D, PASI}	Load integer twin word from alternate space	259	
LDUB	Load unsigned byte	239	
LDUBA ^{PASI}	Load unsigned byte from alternate space	240	
LDUH	Load unsigned halfword	239	
LDUHA ^{PASI}	Load unsigned halfword from alternate space	240	
LDUW	Load unsigned word	239	
LDUWA ^{PASI}	Load unsigned word from alternate space	240	
LDX	Load extended	239	
LDXA ^{PASI}	Load extended from alternate space	240	
STB	Store byte	334	
STBA ^{PASI}	Store byte into alternate space	335	
STTW ^D	Store twin word	352	
STTWA ^{D, PASI}	Store twin word into alternate space	354	
STH	Store halfword	334	
STHA ^{PASI}	Store halfword into alternate space	335	
STW	Store word	334	
STWA ^{PASI}	Store word into alternate space	335	
STX	Store extended	334	
STXA ^{PASI}	Store extended into alternate space	335	
<i>Memory Operations — Miscellaneous</i>			
LDFSR ^D	Load floating-point state register (lower)	251	
LDXEFSR	Load Entire floating-point state register	264	VIS 3 (2010), VIS 3B (2010)
LDXFSR	Load floating-point state register	264	
MWAIT	Monitor Wait	292	OSA 2015
MEMBAR	Memory barrier	269	
PREFETCH	Prefetch data	303	
PREFETCHA ^{PASI}	Prefetch data from alternate space	303	

TABLE 7-4 Instruction Set - by Functional Category (4 of 7)

Instruction	Category and Function	Page	Added in
STFSR ^D	Store floating-point state register (lower)	345	
STXFSR	Store floating-point state register	356	
<i>Floating-Point Arithmetic Operations</i>			
FABS<s d q>	Floating-point absolute value	152	
FADD<s d q>	Floating-point add	153	
FDIV<s d q>	Floating-point divide	164	
FdMULq	Floating-point multiply double to quad	190	
FHADD<s d>	Floating-point add and halve	166	VIS 3 (2010)
FHSUB<s d>	Floating-point subtract and halve	167	VIS 3 (2010)
FMADD(s,d)	Floating-point multiply-add single/double (fused)	175	
FMSUB(s,d)	Floating-point multiply-subtract single/double (fused)	175	
FMUL<s d q>	Floating-point multiply	190	
FNADD<s d>	Floating-point add and negate	191	VIS 3 (2010)
FNHADD<s d>	Floating-point add, halve, and negate	194	VIS 3 (2010)
FNMADD(s,d)	Floating-point negative multiply-add single/double (fused)	175	
FNMUL<s d>	Floating-point multiply and negate	195	VIS 3 (2010)
FNEG<s d q>	Floating-point negate	193	
FNMSUB(s,d)	Floating-point negative multiply-subtract single/double (fused)	175	
FNsMULd	Floating-point multiply single to double, and negate	195	VIS 3 (2010)
FsMULd	Floating-point multiply single to double	190	
FSQRT<s d q>	Floating-point square root	230	
FSUB<s d q>	Floating-point subtract	234	
<i>Floating-Point Comparison Operations</i>			
FCMP<s d q>	Floating-point compare	162	
FCMPE<s d q>	Floating-point compare (exception if unordered)	162	
FLCMP{s,d}	Lexicographic compare	169	VIS 3 (2010)
<i>Register-Window Control Operations</i>			
ALLCLEAN ^P	Mark all register window sets as “clean”	118	
INVALW ^P	Mark all register window sets as “invalid”	237	
FLUSHW	Flush register windows	174	
NORMALW ^P	“Other” register windows become “normal” register windows	295	
OTHERW ^P	“Normal” register windows become “other” register windows	297	
RESTORE	Restore caller’s window	315	
RESTORED ^P	Window has been restored	317	
SAVE	Save caller’s window	322	
SAVED ^P	Window has been saved	324	
<i>Miscellaneous Operations</i>			
FLUSH	Flush instruction memory	171	
NOP	No operation	294	
<i>Integer Arithmetic Operations on R Registers</i>			
ADD (ADDcc)	Add (and modify condition codes)	110	
ADDC (ADDCcc)	Add with carry (and modify condition codes)	110	
ADDXC (ADDXCcc)	Add with extended carry (and modify condition codes)	111	VIS 3 (2010)
MULX	Multiply 64-bit integers	291	

TABLE 7-4 Instruction Set - by Functional Category (5 of 7)

Instruction	Category and Function	Page	Added in
SDIV ^D (SDIVcc ^D)	32-bit signed integer divide (and modify condition codes)	325	
SDIVX	64-bit signed integer divide	291	
SMUL ^D (SMULcc ^D)	Signed integer multiply (and modify condition codes)	333	
SUB (SUBBcc)	Subtract (and modify condition codes)	357	
SUBC (SUBCcc)	Subtract with carry (and modify condition codes)	357	
SUBXC (SUBXCcc)	Subtract with extended carry (and modify condition codes)	358	OSA 2015
TADDcc	Tagged add and modify condition codes (trap on overflow)	362	
TADDccTV ^D	Tagged add and modify condition codes (trap on overflow)	363	
TSUBcc	Tagged subtract and modify condition codes (trap on overflow)	367	
TSUBccTV ^D	Tagged subtract and modify condition codes (trap on overflow)	368	
UDIV ^D (UDIVcc ^D)	Unsigned integer divide (and modify condition codes)	369	
UDIVX	64-bit unsigned integer divide	291	
UMUL ^D (UMULcc ^D)	Unsigned integer multiply (and modify condition codes)	371	
UMULXHI	64 × 64 multiply yielding upper 64 bits of product	372	VIS 3 (2010)
XMULX[HI]	XOR Multiply	389	VIS 3 (2010)
Integer SIMD (Partitioned) Operations on F Registers			
FMEAN16	16-bit partitioned average	177	VIS 3 (2010)
FPADD	Partitioned integer add	201	VIS 1 (1995),3 , OSA 2015 (VIS 4)
FPADDS	Partitioned integer add with saturation	204	VIS 3 (2010), OSA 2015 (VIS 4)
FPCMP	Partitioned Compare signed integer values	207	VIS 1 (1995) , OSA 2015 (VIS 4)
FPCMPU	Partitioned Compare unsigned integer values	210	VIS 3 (2010), VIS 3B (2010), OSA 2015 (VIS 4)
FPMAX	Partitioned Integer Maximum	214	OSA 2015 (VIS 4)
FPMAXU	Partitioned Integer Maximum, Unsigned	214	OSA 2015 (VIS 4)
FPMIN	Partitioned Integer Minimum	217	OSA 2015 (VIS 4)
FPMINU	Partitioned Integer Minimum, Unsigned	217	OSA 2015 (VIS 4)
FPSUB<16,32>[S]	Partitioned Integer Subtract	219	VIS 1 (1995) , OSA 2015 (VIS 4)
FPSUBS	Partitioned Integer Subtract with Saturation	222	VIS 3 (2010), OSA 2015 (VIS 4)
FSLAS<16,32>	Partitioned Shift Left Arithmetic, with Saturation	228	OSA 2011 (VIS 3)
FSL<16,32>	Partitioned Shift Left Logical, 16- or 32-bit elements	228	OSA 2011 (VIS 3)
FSRA<16,32>	Partitioned Shift Right Arithmetic, 16- or 32-bit elements	228	OSA 2011 (VIS 3)
FSRL<16,32>	Partitioned Shift Right Logical, 16- or 32-bit elements	228	OSA 2011 (VIS 3)
FS<LL RL RA><16 32>	16- or 32-bit partitioned shift, left or right	228	VIS 3 (2010)
Integer Arithmetic and Logical Operations on F Registers			
FPADD64	Integer add, 64-bit F registers	201	VIS 3B (2010)
FMUL8x16	8x16 partitioned product	186	VIS 1 (1995)
FMUL8x16[AU AL]	8x16 upper/lower α partitioned product	186	VIS 1 (1995)
FMUL8[SU UL]x16	8x16 upper/lower partitioned product	186	VIS 1 (1995)
FMULD8[SU UL]x16	8x16 upper/lower partitioned product	186	VIS 1 (1995)
FPMADDX[HI]	64-bit Integer multiply-add (low and high 64 bit results)	213	OSA 2011

TABLE 7-4 Instruction Set - by Functional Category (6 of 7)

Instruction	Category and Function	Page	Added in
FPSUB64	Integer Subtract, 64-bit F registers	201	VIS 3B (2010)
<i>Miscellaneous Operations on R Registers</i>			
LZCNT	Leading zeroes count, on 64-bit integer register	266	VIS 3 (2010)
POPC	Population count	301	
SETHI	Set high 22 bits of low word of integer register	327	
<i>Miscellaneous Operations on F Registers</i>			
EDGE<8 16 32>[L]cc	Edge handling instructions (and modify condition codes)	149	VIS 1 (1995)
EDGE<8 16 32>[L]N	Edge handling instructions	150	VIS 2 (2001)
FCHKSM16	16-bit partitioned checksum	161	VIS 3 (2010)
PDIST ^D	Pixel component distance	299	VIS 1 (1995)
PDISTN	Distance between eight 8-bit components with no accumulation	300	VIS 3 (2010)
<i>Cryptographic and Secure Hash Operations</i>			
AES_DROUND<01 23>	AES Decryption, Columns 0&1 (or 2&3)	112	Crypto
AES_DROUND<01 23>_LAST	AES Decryption, Columns 0&1 (or 2&3), Last Round	112	Crypto
AES_EROUND<01 23>	AES Encryption, Columns 0&1 (or 2&3)	112	Crypto
AES_EROUND<01 23>_LAST	AES Encryption, Columns 0&1 (or 2&3), Last Round	112	Crypto
AES_KEXPAND0	AES Key Expansion, without RCON	115	Crypto
AES_KEXPAND1	AES Key Expansion, with RCON	112	Crypto
AES_KEXPAND2	AES Key Expansion, without SBOX	115	Crypto
CAMELLIA_F ^N	Camellia "F" operation	131	Crypto
CAMELLIA_FL ^N	Camellia "FL" operation	133	Crypto
CAMELLIA_FLI ^N	Camellia "FLI" operation	133	Crypto
CRC32C ^N	Cyclic Redundancy Check	141	Crypto
DES_IP ^N	DES Initial Permutation	145	Crypto
DES_IIP ^N	DES Inverse Initial Permutation	145	Crypto
DES_KEXPAND ^N	DES Key Expand	145	Crypto
DES_ROUND ^N	DES Round Operations	143	Crypto
MD5 ^N	MD5 Hash Operation	268	Crypto
MPMUL ^N	Multiple-Precision Multiply	286	Crypto
MONTMUL ^N	Montgomery Multiplication	272	Crypto
MONTMUL ^N	Montgomery Squaring	276	Crypto
SHA1 ^N	SHA1 Secure Hash operation	328	Crypto
SHA256 ^N	SHA256 Secure Hash operation	328	Crypto
SHA512 ^N	SHA512 Secure Hash operation	328	Crypto
XMONTMUL ^N	Montgomery XOR Multiplication	381	Crypto
XMONTMUL ^N	Montgomery XOR Squaring	384	Crypto
XMPMUL ^N	Multiple-Precision XOR Multiply	378	Crypto
<i>Control and Status Register Access</i>			
PAUSE	Pause Virtual Processor	298	OSA 2011
RDASI	Read ASI register	310	
RDAsr ^{PASR}	Read ancillary state register	310	
RDCCR	Read Condition Codes register (CCR)	310	
RDFPRS	Read Floating-Point Registers State register (FPRS)	310	
RDGSR	Read General Status register (GSR)	310	

TABLE 7-4 Instruction Set - by Functional Category (7 of 7)

Instruction	Category and Function	Page	Added in
RDPC	Read Program Counter register (PC)	310	
RDPR ^P	Read privileged register	313	
RDSOFTINT ^P	Read per-virtual processor Soft Interrupt register (SOFTINT)	310	
RDSTICK ^{P_{dis},H_{dis}}	Read System Tick register (STICK)	310	
RDSTICK_CMPR ^P	Read System Tick Compare register (STICK_CMPR)	310	
RDTICK ^{P_{dis},H_{dis}}	Read Tick register (TICK)	310	
RDY ^D	Read Y register	310	
SIAM	Set interval arithmetic mode	330	VIS 2 (2001)
WRASI	Write ASI register	373	
WRAsr ^{P_{ASR}}	Write ancillary state register	373	
WRCCR	Write Condition Codes register (CCR)	373	
WRFPSR	Write Floating-Point Registers State register (FPSR)	373	
WRGSR	Write General Status register (GSR)	373	
WRPR ^P	Write privileged register	536	
WRSOFTINT ^P	Write per-virtual processor Soft Interrupt register (SOFTINT)	373	
WRSOFTINT_CLR ^P	Clear bits of per-virtual processor Soft Interrupt register (SOFTINT)	373	
WRSOFTINT_SET ^P	Set bits of per-virtual processor Soft Interrupt register (SOFTINT)	373	
WRSTICK ^P	Write System Tick register (STICK)	373	
WRSTICK_CMPR ^P	Write System Tick Compare register (STICK_CMPR)	373	
WRY ^D	Write Y register	373	

In the remainder of this chapter, related instructions are grouped into subsections. Each subsection contains the following sets of information:

(1) Instruction Table. This section of an instruction page lists the instructions that are defined in the subsection, including the values of the field(s) that uniquely identify the instruction(s) and its assembly language syntax. In the rightmost column, Software (alphabetic) and Implementation (numeric) classifications for the instructions are provided. The meaning of the alphabetic Software Classifications is as follows:

Software Usage Class	How this feature may be used	Attributes
A “Use Freely”	Use freely.	<ul style="list-style-type: none"> ■ Compilers always free to use (no option to disable use). ■ Executes well across all implementations.
B “Use Carefully”	Use with care/forethought in portable software	<ul style="list-style-type: none"> ■ Usage is being <i>phased in</i>. ■ A compiler option exists to enable/disable references to this feature; by default, use is <i>enabled</i>.
C “New Feature”	Use only in platform-specific software (privileged code, DLLs, and non-portable applications)	<ul style="list-style-type: none"> ■ New feature; usage is being <i>phased in</i>. ■ A compiler option† exists to enable/disable references to this feature; by default, use is <i>disabled</i>. ■ An assembler option† exists to enable/disable references to this feature; by default, use is <i>disabled</i>. If use is enabled, reference to feature triggers a warning; if disabled, reference triggers an error message.
D “Deprecated”	Use in portable software is strongly discouraged.	<ul style="list-style-type: none"> ■ Usage is being <i>phased out</i> and this feature may not perform as well in future implementations. ■ A compiler option† exists to enable/disable use of this feature; by default, use is <i>disabled</i>. ■ An assembler option† exists to enable/disable references to this feature; by default, use is <i>disabled</i>. If use is enabled, reference to feature triggers a warning; if disabled, reference triggers an error message. ■
N “Non-portable (platform-specific)”	Only use in platform-specific software (privileged code, hyperprivileged code, DLLs, JIT code, and [if absolutely necessary] non-portable applications)	<ul style="list-style-type: none"> ■ A compiler option† exists to enable/disable references to this feature; by default, use is <i>disabled</i>. ■ An assembler option† exists to enable/disable references to this feature; by default, use is <i>disabled</i>. If use is enabled, reference to feature triggers a warning; if disabled, reference triggers an error message. ■

(2) Illustration of Instruction Format(s). These illustrations show how the instruction is encoded in a 32-bit word in memory. In them, a dash (—) indicates that the field is *reserved* for future versions of the architecture and must be 0 in any instance of the instruction. If a conforming Oracle SPARC Architecture implementation encounters nonzero values in these fields, its behavior is as defined in *Reserved Opcodes and Instruction Fields* on page 93.

(3) Description. This subsection describes the operation of the instruction, its features, restrictions, and exception-causing conditions.

(4) Exceptions. The exceptions that can occur as a consequence of attempting to execute the instruction(s). Exceptions due to an *IAE_**, and interrupts are not listed because they can occur on any instruction. An instruction not implemented in hardware generates an *illegal_instruction* exception and therefore will not generate any of the other exceptions listed. Exceptions are listed in order of trap priority (see *Trap Priorities* on page 458), from highest to lowest priority.


(5) See Also. A list of related instructions (on selected pages).

Note | This specification does not contain any timing information (in either cycles or elapsed time), since timing is always implementation dependent.

ADD

7.1 Add

Instruction	op3	Operation	Assembly Language Syntax	Class
ADD	00 0000	Add	add <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1
ADDcc	01 0000	Add and modify cc's	addcc <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1
ADDC	00 1000	Add with 32-bit Carry	addc <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1
ADDCcc	01 1000	Add with 32-bit Carry and modify cc's	addccc <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1



Description If $i = 0$, ADD and ADDcc compute “ $R[rs1] + R[rs2]$ ”. If $i = 1$, they compute “ $R[rs1] + \text{sign_ext}(\text{simm13})$ ”. In either case, the sum is written to $R[rd]$.

ADDC and ADDCcc (“ADD with carry”) also add the CCR register’s 32-bit carry (icc.c) bit. That is, if $i = 0$, they compute “ $R[rs1] + R[rs2] + \text{icc.c}$ ” and if $i = 1$, they compute “ $R[rs1] + \text{sign_ext}(\text{simm13}) + \text{icc.c}$ ”. In either case, the sum is written to $R[rd]$.

ADDcc and ADDCcc modify the integer condition codes (CCR.icc and CCR.xcc). Overflow occurs on addition if both operands have the same sign and the sign of the sum is different from that of the operands.

Programming Note | ADDC and ADDCcc read the 32-bit condition codes’ carry bit (CCR.icc.c), not the 64-bit condition codes’ carry bit (CCR.xcc.c).

SPARC V8 Compatibility Note | ADDC and ADDCcc were previously named ADDX and ADDXcc, respectively, in SPARC V8.

An attempt to execute an ADD, ADDcc, ADDC or ADDCcc instruction when $i = 0$ and reserved instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*


■ **See Also** Add Extended with 64-bit Carry on page 111

ADDXC

7.2 Add Extended with 64-bit Carry VIS 3

The ADDXC instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
ADDXC ^N	0 0001 0001	Add Extended with 64-bit Carry	addxc <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2011
ADDXCcc ^N	0 0001 0011	Add Extended with 64-bit Carry and modify cc's	addxcc <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2011



Description ADDXC and ADDXCcc (“ADD extended with carry”) both compute “R[rs1] + R[rs2] + xcc.c” and write the sum into R[rd].

In addition, ADDXCcc modifies the integer condition codes (CCR.icc and CCR.xcc). Overflow occurs on addition if both operands have the same sign and the sign of the sum is different from that of the operands.

Programming Note | ADDXC and ADDXCcc can be used in conjunction with MULX and UMULXHI to speed up large multiword integer multiplication computations.

Exceptions None

See Also Add on page 110
Subtract Extended with 64-bit Carry on page 358


AES Crypto (4-operand)

7.3

AES Cryptographic Operations (4 operand) Crypto

The AES instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	op5	Operation	Assembly Language Syntax		Class
AES_EROUND01	0000	AES Encrypt columns 0&1	aes_eround01	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_EROUND23	0001	AES Encrypt columns 2&3	aes_eround23	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_DROUND01	0010	AES Decrypt columns 0&1	aes_dround01	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_DROUND23	0011	AES Decrypt columns 2&3	aes_dround23	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_EROUND01_LAST	0100	AES Encrypt columns 0&1 last round	aes_eround01_1	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_EROUND23_LAST	0101	AES Encrypt columns 2&3 last round	aes_eround23_1	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_DROUND01_LAST	0110	AES Decrypt columns 0&1 last round	aes_dround01_1	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_DROUND23_LAST	0111	AES Decrypt columns 2&3 last round	aes_dround23_1	$freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}$	N1
AES_KEXPAND1	1000	AES Key expansion with RCON	aes_kexpand1	$freq_{rs1}, freq_{rs2}, imm5, freq_{rd}$	N1



Description

The AES instructions support encryption, decryption and key expansion for the Advanced Encryption Standard. This standard is available as FIPS-197 on <http://nist.gov>. Encryption is performed by looping on a set of primitive functions including SubBytes, ShiftRows, MixColumns, and AddRoundKey. Decryption loops the inverse set of functions which include InvShiftRows, InvSubBytes, AddRoundKey, and InvMixColumns. Key expansion utilizes SubBytes and RotWord functions as well as the XOR operation. A number of temporary variables are used in the functional descriptions below. For example, `data_sb` is the result of applying the SubBytes function to the RS2 and RS3 operand.

AES_EROUND01:

```

data_sb{127:0} ← SubBytes( FD[rs2]{63:0} :: FD[rs3]{63:0} )
data_sr{127:0} ← ShiftRows( data_sb{127:0} );
data_mc{127:0} ← MixColumns( data_sr{127:0} );
FD[rd]{63:0} ← AddRoundKey( data_mc{63:0}, FD[rs1]{63:0} );
    
```

AES_EROUND23:

```

data_sb{127:0} ← SubBytes( FD[rs2]{63:0} :: FD[rs3]{63:0} );
data_sr{127:0} ← ShiftRows( data_sb{127:0} );
data_mc{127:0} ← MixColumns( data_sr{127:0} );
FD[rd]{63:0} ← AddRoundKey( data_mc{127:64}, FD[rs1]{63:0} );
    
```

AES_DROUND01:

```

data_sr{127:0} ← InvShiftRows( FD[rs2]{63:0} :: FD[rs3]{63:0} );
data_sb{127:0} ← InvSubBytes( data_sr{127:0} );
data_ark{63:0} ← AddRoundKey( data_sb{63:0}, FD[rs1]{63:0} );
FD[rd]{63:0} ← InvMixColumn( data_ark{63:0} );
    
```


AES Crypto (4-operand)

AES_DROUND23:

$data_sr\{127:0\} \leftarrow \text{InvShiftRows}(F_D[rs2]\{63:0\} :: F_D[rs3]\{63:0\});$
 $data_sb\{127:0\} \leftarrow \text{InvSubBytes}(data_sr\{127:0\});$
 $data_ark\{63:0\} \leftarrow \text{AddRoundKey}(data_sb\{127:64\} , F_D[rs1]\{63:0\});$
 $F_D[rd]\{63:0\} \leftarrow \text{InvMixColumn}(data_ark\{63:0\});$

AES_EROUND01_LAST:

$data_sb\{127:0\} \leftarrow \text{SubBytes}(F_D[rs2]\{63:0\} :: F_D[rs3]\{63:0\});$
 $data_sr\{127:0\} \leftarrow \text{ShiftRows}(data_sb\{127:0\});$
 $F_D[rd]\{63:0\} \leftarrow \text{AddRoundKey}(data_sr\{63:0\} , F_D[rs1]\{63:0\});$

AES_EROUND23_LAST:

$data_sb\{127:0\} \leftarrow \text{SubBytes}(F_D[rs2]\{63:0\} :: F_D[rs3]\{63:0\});$
 $data_sr\{127:0\} \leftarrow \text{ShiftRows}(data_sb\{127:0\});$
 $F_D[rd]\{63:0\} \leftarrow \text{AddRoundKey}(data_sr\{127:64\} , F_D[rs1]\{63:0\});$

AES_DROUND01_LAST:

$data_sr\{127:0\} \leftarrow \text{InvShiftRows}(F_D[rs2]\{63:0\} :: F_D[rs3]\{63:0\});$
 $data_sb\{127:0\} \leftarrow \text{InvSubBytes}(data_sr\{127:0\});$
 $F_D[rd]\{63:0\} \leftarrow \text{AddRoundKey}(data_sb\{63:0\} , F_D[rs1]\{63:0\});$

AES_DROUND23_LAST:

$data_sr\{127:0\} \leftarrow \text{InvShiftRows}(F_D[rs2]\{63:0\} :: F_D[rs3]\{63:0\});$
 $data_sb\{127:0\} \leftarrow \text{InvSubBytes}(data_sr\{127:0\});$
 $F_D[rd]\{63:0\} \leftarrow \text{AddRoundKey}(data_sb\{127:64\} , F_D[rs1]\{63:0\});$

AES_KEXPAND1:

$data_rw\{31:0\} \leftarrow \text{RotWord}(F_D[rs2]\{31:0\});$
 $data_sb\{31:0\} \leftarrow \text{SubBytes}(data_rw\{31:0\});$
 $F_D[rd]\{63:32\} \leftarrow data_sb\{31:0\} \text{ xor } F_D[rs1]\{63:32\} \text{ xor } rcon\{31:0\};$
 $F_D[rd]\{31:0\} \leftarrow data_sb\{31:0\} \text{ xor } F_D[rs1]\{63:32\} \text{ xor } rcon\{31:0\} \text{ xor } F_D[rs1]\{31:0\};$

where rcon is a function of imm5, as shown below:

imm5	rcon{31:0}
00000 ₂	0100 0000 ₁₆
00001 ₂	0200 0000 ₁₆
00010 ₂	0400 0000 ₁₆
00011 ₂	0800 0000 ₁₆
00100 ₂	1000 0000 ₁₆
00101 ₂	2000 0000 ₁₆
00110 ₂	4000 0000 ₁₆
00111 ₂	8000 0000 ₁₆
01000 ₂	1b00 0000 ₁₆
01001 ₂	3600 0000 ₁₆
01010 ₂	0000 0000 ₁₆
11111 ₂	

AES Crypto (4-operand)

Programming Note The AES instructions are components of the overall AES algorithm. To perform an encryption or decryption, the key must first be expanded. Key expansion is done only once per session key. The expanded keys are then applied to all blocks for that session. In the following example, expanded keys are stored in F0 thru F42, plain text is stored in F52 and F54, F56 and F58 are scratch registers, F60 and F62 hold the cipher text. For each block, the following instruction sequence can be applied for an AES 128 ECB encryption:

```
fxor          %f00, %f52,      %f52    //initial ARK
fxor          %f02, %f54,      %f54
aes_eround01  %f04, %f52, %f54, %f56    // Round 1
aes_eround23  %f06, %f52, %f54, %f58
aes_eround01  %f08, %f56, %f58, %f52    // Round 2
aes_eround23  %f10, %f56, %f58, %f54
aes_eround01  %f12, %f52, %f54, %f56    // Round 3
aes_eround23  %f14, %f52, %f54, %f58
aes_eround01  %f16, %f56, %f58, %f52    // Round 4
aes_eround23  %f18, %f56, %f58, %f54
aes_eround01  %f20, %f52, %f54, %f56    // Round 5
aes_eround23  %f22, %f52, %f54, %f58
aes_eround01  %f24, %f56, %f58, %f52    // Round 6
aes_eround23  %f26, %f56, %f58, %f54
aes_eround01  %f28, %f52, %f54, %f56    // Round 7
aes_eround23  %f30, %f52, %f54, %f58
aes_eround01  %f32, %f56, %f58, %f52    // Round 8
aes_eround23  %f34, %f56, %f58, %f54
aes_eround01  %f36, %f52, %f54, %f56    // Round 9
aes_eround23  %f38, %f52, %f54, %f58
aes_eround01_1 %f40, %f56, %f58, %f60    // Round 10
aes_eround23_1 %f42, %f56, %f58, %f62
```

If `CFR.aes = 0`, an attempt to execute any AES instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that `CFR.aes = 1` before executing any of these AES instructions. If `CFR.aes = 0`, then software should assume that an attempt to execute one of the AES instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.

Therefore, if `CFR.aes = 0`, software should perform the corresponding AES operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*

See Also *AES Cryptographic Operations (3 operand)* on page 115

AES Crypto (3-operand)

7.4 AES Cryptographic Operations (3 operand) Crypto

The AES instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class
AES_KEXPAND0	1 0011 0000	AES Key expansion without RCON	<code>aes_kexpand0 <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	N1
AES_KEXPAND2	1 0011 0001	AES Key expansion without SBOX	<code>aes_kexpand2 <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	N1



Description AES_KEXPAND0:

```

data_sb{31:0} ← SubBytes( FD[rs2]{31:0} );
FD[rd]{63:32} ← data_sb{31:0} xor FD[rs1]{63:32};
FD[rd]{31:0}  ← data_sb{31:0} xor FD[rs1]{63:32} xor FD[rs1]{31:0};
    
```

AES_KEXPAND2:

```

FD[rd]{63:32} ← FD[rs2]{31:0} xor FD[rs1]{63:32};
FD[rd]{31:0}  ← FD[rs2]{31:0} xor FD[rs1]{63:32} xor FD[rs1]{31:0};
    
```

Programming Note The AES instructions are components of the overall AES algorithm. To perform an encryption or decryption, the key must first be expanded. Key expansion is done only once per session key; the expanded keys are then applied to all blocks for that session. The following is an example of key expansion for AES 128. The original keys are loaded into F0 and F2. The expanded keys are stored in F4 thru F42:

```

aes_kexpand1 %f0, %f2, 0x0, %f4    !# w[4] , w[5]
aes_kexpand2 %f2, %f4, %f6        !# w[6] , w[7]
aes_kexpand1 %f4, %f6, 0x1, %f8    !# w[8] , w[9]
aes_kexpand2 %f6, %f8, %f10       !# w[10], w[11]
aes_kexpand1 %f8, %f10, 0x2, %f12  !# w[12], w[13]
aes_kexpand2 %f10, %f12, %f14     !# w[14], w[15]
aes_kexpand1 %f12, %f14, 0x3, %f16 !# w[16], w[17]
aes_kexpand2 %f14, %f16, %f18     !# w[18], w[19]
aes_kexpand1 %f16, %f18, 0x4, %f20 !# w[20], w[21]
aes_kexpand2 %f18, %f20, %f22     !# w[22], w[23]
aes_kexpand1 %f20, %f22, 0x5, %f24 !# w[24], w[25]
aes_kexpand2 %f22, %f24, %f26     !# w[26], w[27]
aes_kexpand1 %f24, %f26, 0x6, %f28 !# w[28], w[29]
aes_kexpand2 %f26, %f28, %f30     !# w[30], w[31]
aes_kexpand1 %f28, %f30, 0x7, %f32 !# w[32], w[33]
aes_kexpand2 %f30, %f32, %f34     !# w[34], w[35]
aes_kexpand1 %f32, %f34, 0x8, %f36 !# w[36], w[37]
aes_kexpand2 %f34, %f36, %f38     !# w[38], w[39]
aes_kexpand1 %f36, %f38, 0x9, %f40 !# w[40], w[41]
aes_kexpand2 %f38, %f40, %f42     !# w[42], w[43]
    
```

If CFR.aes = 0, an attempt to execute any AES instruction causes a *compatibility_feature* exception.

AES Crypto (3-operand)

Programming Note Software *must* check that `CFR.aes = 1` before executing any of these AES instructions. If `CFR.aes = 0`, then software should assume that an attempt to execute one of the AES instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.

Therefore, if `CFR.aes = 0`, software should perform the corresponding AES operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*

■ *See Also* *AES Cryptographic Operations (4 operand)* on page 112

ALIGNADDRESS

7.5 Align Address VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
ALIGNADDRESS	0 0001 1000	Calculate address for misaligned data access	<code>alignaddr <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	A1	UA 2005
ALIGNADDRESS_LITTLE	0 0001 1010	Calculate address for misaligned data access, little-endian	<code>alignaddr1 <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	A1	UA 2005



Description ALIGNADDRESS adds two integer values, R[rs1] and R[rs2], and stores the result (with the least significant 3 bits forced to 0) in the integer register R[rd]. The least significant 3 bits of the result are stored in the GSR.align field.

ALIGNADDRESS_LITTLE is the same as ALIGNADDRESS except that the two's complement of the least significant 3 bits of the result is stored in GSR.align.

Note ALIGNADDRESS_LITTLE generates the opposite-endian byte ordering for a subsequent FALIGNDATAg operation.

A byte-aligned 64-bit load can be performed as shown below.

```
alignaddr  Address, Offset, Address !set GSR.align
ldd       [Address], %d0
ldd       [Address + 8], %d2
faligndata %d0, %d2, %d4          !use GSR.align to select bytes
```

If the floating-point unit is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an ALIGNADDRESS or ALIGNADDRESS_LITTLE instruction causes an *fp_disabled* exception.


Exceptions *fp_disabled*

See Also Align Data (using gsr.align) on page 154

ALLCLEAN

7.6 Mark All Register Window Sets “Clean”

Instruction	Operation	Assembly Language Syntax	Class	Added
ALLCLEAN ^P	Mark all register window sets as “clean”	allclean	A1	UA 2005



Description The ALLCLEAN instruction marks all register window sets as “clean”; specifically, it performs the following operation:

$$\text{CLEANWIN} \leftarrow (N_REG_WINDOWS - 1)$$

Programming Note ALLCLEAN is used to indicate that all register windows are “clean”; that is, do not contain data belonging to other address spaces. It is needed because the value of *N_REG_WINDOWS* is not known to privileged software.

An attempt to execute an ALLCLEAN instruction when reserved instruction bits 18:0 are nonzero causes an *illegal_instruction* exception.

An attempt to execute an ALLCLEAN instruction in nonprivileged mode (*PSTATE.priv* = 0) causes a *privileged_opcode* exception.

Exceptions

illegal_instruction
privileged_opcode

See Also

INVALW on page 237
NORMALW on page 295
OTHERW on page 297
RESTORED on page 317
SAVED on page 324

Exceptions..),Exceptions..),Exceptions..),Exceptions.. 7.7),AND Logical Operation

Instruction	op3	Operation	Assembly Language Syntax	Class
AND	00 0001	and	and <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1
ANDcc	01 0001	and and modify cc’s	andcc <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1
ANDN	00 0101	and not	andn <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1
ANDNcc	01 0101	and not and modify cc’s	andncc <i>reg_{rs1}, reg_or_imm, reg_{rd}</i>	A1



Description These instructions implement bitwise logical **and** operations. They compute “R[rs1] **op** R[rs2]” if *i* = 0, or “R[rs1] **op** sign_ext(simm13)” if *i* = 1, and write the result into R[rd].

AND, ANDN

ANDcc and ANDNcc modify the integer condition codes (*icc* and *xcc*). They set the condition codes as follows:

- *icc.v*, *icc.c*, *xcc.v*, and *xcc.c* are set to 0
- *icc.n* is copied from bit 31 of the result
- *xcc.n* is copied from bit 63 of the result
- *icc.z* is set to 1 if bits 31:0 of the result are zero (otherwise to 0)
- *xcc.z* is set to 1 if all 64 bits of the result are zero (otherwise to 0)


ANDN and ANDNcc logically negate their second operand before applying the main (**and**) operation.

An attempt to execute an AND, ANDcc, ANDN or ANDNcc instruction when *i* = 0 and reserved instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*

7.8 Three-Dimensional Array Addressing vis 1

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
ARRAY8	0 0001 0000	Convert 8-bit 3D address to blocked byte address	array8 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1	UA 2005
ARRAY16	0 0001 0010	Convert 16-bit 3D address to blocked byte address	array16 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1	UA 2005
ARRAY32	0 0001 0100	Convert 32-bit 3D address to blocked byte address	array32 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1	UA 2005



Description These instructions convert three-dimensional (3D) fixed-point addresses contained in R[rs1] to a blocked-byte address; they store the result in R[rd]. Fixed-point addresses typically are used for address interpolation for planar reformatting operations. Blocking is performed at the 64-byte level to maximize external cache block reuse, and at the 64-Kbyte level to maximize TLB entry reuse, regardless of the orientation of the address interpolation. These instructions specify an element size of 8 bits (ARRAY8), 16 bits (ARRAY16), or 32 bits (ARRAY32).

The second operand, R[rs2], specifies the power-of-2 size of the X and Y dimensions of a 3D image array. The legal values for R[rs2] and their meanings are shown in TABLE 7-5. Illegal values produce undefined results in the destination register, R[rd].

TABLE 7-5 3D R[rs2] Array X and Y Dimensions

R[rs2] Value (<i>n</i>)	Number of Elements
0	64
1	128
2	256
3	512
4	1024
5	2048

Implementation Note Architecturally, an illegal R[rs2] value (>5) causes the array instructions to produce undefined results. For historic reference, past implementations of these instructions have ignored R[rs2]{63:3} and have treated R[rs2] values of 6 and 7 as if they were 5.

The array instructions facilitate 3D texture mapping and volume rendering by computing a memory address for data lookup based on fixed-point x, y, and z coordinates. The data are laid out in a blocked fashion, so that points which are near one another have their data stored in nearby memory locations.

If the texture data were laid out in the obvious fashion (the z = 0 plane, followed by the z = 1 plane, etc.), then even small changes in z would result in references to distant pages in memory. The resulting lack of locality would tend to result in TLB misses and poor performance. The three versions of the array instruction, ARRAY8, ARRAY16, and ARRAY32, differ only in the scaling of the computed memory offsets. ARRAY16 shifts its result left by one position and ARRAY32 shifts left by two in order to handle 16- and 32-bit texture data.

When using the array instructions, a “blocked-byte” data formatting structure is imposed. The $N \times N \times M$ volume, where $N = 2^n \times 64$, $M = m \times 32$, $0 \leq n \leq 5$, $1 \leq m \leq 16$ should be composed of $64 \times 64 \times 32$ smaller volumes, which in turn should be composed of $4 \times 4 \times 2$ volumes. This data structure is optimal for 16-bit data. For 16-bit data, the $4 \times 4 \times 2$ volume has 64 bytes of data, which is ideal for reducing cache-line misses; the $64 \times 64 \times 32$ volume will have 256 Kbytes of data, which is good for improving the TLB hit rate. FIGURE 7-1 illustrates how the data has to be organized, where the origin

ARRAY<8|16|32>

(0,0,0) is assumed to be at the lower-left front corner and the x coordinate varies faster than y than z. That is, when traversing the volume from the origin to the upper right back, you go from left to right, front to back, bottom to top.




FIGURE 7-1 Blocked-Byte Data Formatting Structure

The array instructions have 2 inputs:

The (x,y,z) coordinates are input via a single 64-bit integer organized in R[rs1] as shown in FIGURE 7-2.




FIGURE 7-2 Three-Dimensional Array Fixed-Point Address Format

Note that z has only 9 integer bits, as opposed to 11 for x and y. Also note that since (x,y,z) are all contained in one 64-bit register, they can be incremented or decremented simultaneously with a single add or subtract instruction (ADD or SUB).

So for a 512 x 512 x 32 or a 512 x 512 x 256 volume, the size value is 3. Note that the x and y size of the volume must be the same. The z size of the volume is a multiple of 32, ranging between 32 and 512.

The array instructions generate an integer memory offset, that when added to the base address of the volume, gives the address of the volume element (voxel) and can be used by a load instruction. The offset is correct only if the data has been reformatted as specified above.

The integer parts of x, y, and z are converted to the following blocked-address formats as shown in FIGURE 7-3 for ARRAY8, FIGURE 7-4 for ARRAY16, and FIGURE 7-5 for ARRAY32.




FIGURE 7-3 Three-Dimensional Array Blocked-Address Format (ARRAY8)

ARRAY<8|16|32>




FIGURE 7-4 Three-Dimensional Array Blocked-Address Format (ARRAY16)




FIGURE 7-5 Three Dimensional Array Blocked-Address Format (ARRAY32)

The bits above Z upper are set to 0. The number of zeroes in the least significant bits is determined by the element size. An element size of 8 bits has no zeroes, an element size of 16 bits has one zero, and an element size of 32 bits has two zeroes. Bits in X and Y above the size specified by R[rs2] are ignored.

TABLE 7-6 ARRAY8 Description

Result (R[rd]) Bits	Source (R[rs1] Bits	Field Information
1:0	12:11	X_integer{1:0}
3:2	34:33	Y_integer{1:0}
4	55	Z_integer{0}
8:5	16:13	X_integer{5:2}
12:9	38:35	Y_integer{5:2}
16:13	59:56	Z_integer{4:1}
17+n-1:17	17+n-1:17	X_integer{6+n-1:6}
17+2n-1:17+n	39+n-1:39	Y_integer{6+n-1:6}
20+2n:17+2n	63:60	Z_integer{8:5}
63:20+2n+1	n/a	0

In the above description, if $n = 0$, there are 64 elements, so X_integer{6} and Y_integer{6} are not defined. That is, result{20:17} equals Z_integer{8:5}.

Note To maximize reuse of external cache and TLB data, software should block array references of a large image to the 64-Kbyte level. This means processing elements within a $32 \times 32 \times 64$ block.

The code fragment below shows assembly of components along an interpolated line at the rate of one component per clock.

```

add      Addr, DeltaAddr, Addr
array8  Addr, %g0, bAddr
ldda    [bAddr] #ASI_FL8_PRIMARY, data
faligndata data, accum, accum


```

Exceptions None

7.9 Branch on Integer Condition Codes (Bicc)

Opcode	cond	Operation	icc Test	Assembly Language Syntax	Class
BA	1000	Branch Always	1	ba{ , a} label	A1
BN	0000	Branch Never	0	bn{ , a} label	A1
BNE	1001	Branch on Not Equal	not Z	bne [†] { , a} label	A1
BE	0001	Branch on Equal	Z	be [‡] { , a} label	A1
BG	1010	Branch on Greater	not (Z or (N xor V))	bg{ , a} label	A1
BLE	0010	Branch on Less or Equal	Z or (N xor V)	ble{ , a} label	A1
BGE	1011	Branch on Greater or Equal	not (N xor V)	bge{ , a} label	A1
BL	0011	Branch on Less	N xor V	bl{ , a} label	A1
BGU	1100	Branch on Greater Unsigned	not (C or Z)	bgu{ , a} label	A1
BLEU	0100	Branch on Less or Equal Unsigned	C or Z	bleu{ , a} label	A1
BCC	1101	Branch on Carry Clear (Greater Than or Equal, Unsigned)	not C	bcc [◇] { , a} label	A1
BCS	0101	Branch on Carry Set (Less Than, Unsigned)	C	bcs [∇] { , a} label	A1
BPOS	1110	Branch on Positive	not N	bpos{ , a} label	A1
BNEG	0110	Branch on Negative	N	bneg{ , a} label	A1
BVC	1111	Branch on Overflow Clear	not V	bvc{ , a} label	A1
BVS	0111	Branch on Overflow Set	V	bvs{ , a} label	A1

[†] synonym: bnz [‡] synonym: bz [◇] synonym: bgeu [∇] synonym: blu



Programming Note To set the annul (a) bit for Bicc instructions, append “, a” to the opcode mnemonic. For example, use “bgu, a label”. In the preceding table, braces signify that the “, a” is optional.

Unconditional branches and icc-conditional branches are described below:

- **Unconditional branches** (BA, BN) — If its annul bit is 0 (a = 0), a BN (Branch Never) instruction is treated as a NOP. If its annul bit is 1 (a = 1), the following (delay) instruction is annulled (not executed). In neither case does a transfer of control take place.

BA (Branch Always) causes an unconditional PC-relative, delayed control transfer to the address “PC + (4 × sign_ext(disp22))”. If the annul (a) bit of the branch instruction is 1, the delay instruction is annulled (not executed). If the annul bit is 0 (a = 0), the delay instruction is executed.
- **icc-conditional branches** — Conditional Bicc instructions (all except BA and BN) evaluate the 32-bit integer condition codes (icc), according to the cond field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken, that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign_ext(disp22))”. If FALSE, the branch is not taken.

Bicc

If a conditional branch is taken, the delay instruction is always executed regardless of the value of the annul field. If a conditional branch is not taken and the annul bit is 1 ($a = 1$), the delay instruction is annulled (not executed).

Note | The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

Annulment, delay instructions, and delayed control transfers are described further in Chapter 6, *Instruction Set Overview*.

Programming Note | For optimal performance, a DCTI should not be placed in the instruction word immediately following an annulled branch-always instruction. For additional information, see *DCTI Couples* on page 89.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20), $PSTATE.tct = 1$, and the Bicc instruction will cause a transfer of control (BA or taken conditional branch), then Bicc generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the Bicc instruction) is stored in TPC[TL] and the value of NPC from before the Bicc was executed is stored in TNPC[TL].

Note that BN never causes a *control_transfer_instruction* exception.

Exceptions *control_transfer_instruction* (impl. dep. #450-S20)

BMASK / BSHUFFLE

7.10 Byte Mask and Shuffle VIS 2

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
BMASK	0 0001 1001	Set the GSR.mask field in preparation for a subsequent BSHUFFLE instruction	bmask <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1	UA 2007
BSHUFFLE	0 0100 1100	Permute 16 bytes as specified by GSR.mask	bshuffle <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	B1	UA 2007



Description BMASK adds two integer registers, R[rs1] and R[rs2], and stores the result in the integer register R[rd]. The least significant 32 bits of the result are stored in the GSR.mask field.

BSHUFFLE concatenates the two 64-bit floating-point registers F_D[rs1] (more significant half) and F_D[rs2] (less significant half) to form a 128-bit (16-byte) value. Bytes in the concatenated value are numbered from most significant to least significant, with the most significant byte being byte 0. BSHUFFLE extracts 8 of those 16 bytes and stores the result in the 64-bit floating-point register F_D[rd]. Bytes in F_D[rd] are also numbered from most to least significant, with the most significant being byte 0. The following table indicates which source byte is extracted from the concatenated value to generate each byte in the destination register, F_D[rd].

Destination Byte (in F[rd])	Source Byte
0 (most significant)	(F _D [rs1] :: F _D [rs2]) {GSR.mask{31:28}}
1	(F _D [rs1] :: F _D [rs2]) {GSR.mask{27:24}}
2	(F _D [rs1] :: F _D [rs2]) {GSR.mask{23:20}}
3	(F _D [rs1] :: F _D [rs2]) {GSR.mask{19:16}}
4	(F _D [rs1] :: F _D [rs2]) {GSR.mask{15:12}}
5	(F _D [rs1] :: F _D [rs2]) {GSR.mask{11:8}}
6	(F _D [rs1] :: F _D [rs2]) {GSR.mask{7:4}}
7 (least significant)	(F _D [rs1] :: F _D [rs2]) {GSR.mask{3:0}}

If the floating-point unit is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a BMASK or BSHUFFLE instruction causes an *fp_disabled* exception.

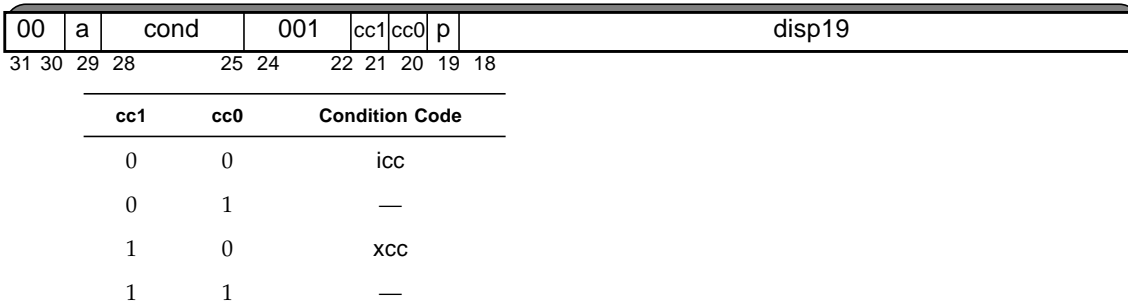
Exceptions *fp_disabled*

■ *See Also* CMASK on page 139

7.11 Branch on Integer Condition Codes with Prediction (BPcc)

Instruction	cond	Operation	cc Test	Assembly Language Syntax	Class
BPA	1000	Branch Always	1	ba{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPn	0000	Branch Never	0	bn{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPNE	1001	Branch on Not Equal	not Z	bnet{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPE	0001	Branch on Equal	Z	be‡{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPG	1010	Branch on Greater	not (Z or (N xor V))	bg{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPLe	0010	Branch on Less or Equal	Z or (N xor V)	ble{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPGE	1011	Branch on Greater or Equal	not (N xor V)	bge{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPL	0011	Branch on Less	N xor V	bl{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPGU	1100	Branch on Greater Unsigned	not (C or Z)	bgu{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPLeU	0100	Branch on Less or Equal Unsigned	C or Z	bleu{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPCC	1101	Branch on Carry Clear (Greater than or Equal, Unsigned)	not C	bcc◊{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPCS	0101	Branch on Carry Set (Less than, Unsigned)	C	bcs∇{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPPOS	1110	Branch on Positive	not N	bpos{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPNEG	0110	Branch on Negative	N	bneg{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPVC	1111	Branch on Overflow Clear	not V	bvc{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1
BPVS	0111	Branch on Overflow Set	V	bvs{,a}{,pt ,pn} <i>i_or_x_cc, label</i>	A1

† synonym: bnz ‡ synonym: bz ◊ synonym: bgeu ∇ synonym: blu



Programming Note To set the annul (a) bit for BPcc instructions, append “, a” to the opcode mnemonic. For example, use `bgu, a %icc, label`. Braces in the preceding table signify that the “, a” is optional. To set the branch prediction bit, append to an opcode mnemonic either “, pt” for predict taken or “, pn” for predict not taken. If neither “, pt” nor “, pn” is specified, the assembler defaults to “, pt”. To select the appropriate integer condition code, include “%icc” or “%xcc” before the label.

Description Unconditional branches and conditional branches are described below.

BPcc

- **Unconditional branches (BPA, BPN)** — A BPN (Branch Never with Prediction) instruction for this branch type ($op2 = 1$) may be used in the SPARC V9 architecture as an instruction prefetch; that is, the effective address ($PC + (4 \times \text{sign_ext}(\text{disp19}))$) specifies an address of an instruction that is expected to be executed soon. If the Branch Never's annul bit is 1 ($a = 1$), then the following (delay) instruction is annulled (not executed). If the annul bit is 0 ($a = 0$), then the following instruction is executed. Branch Always always causes a transfer of control, and In no case does Branch Never ever cause a transfer of control to take place.

BPA (Branch Always with Prediction) causes an unconditional PC-relative, delayed control transfer to the address “ $PC + (4 \times \text{sign_ext}(\text{disp19}))$ ”. If the annul bit of the branch instruction is 1 ($a = 1$), then the delay instruction is annulled (not executed). If the annul bit is 0 ($a = 0$), then the delay instruction is executed.

- **Conditional branches** — Conditional BPcc instructions (except BPA and BPN) evaluate one of the two integer condition codes (icc or xcc), as selected by $cc0$ and $cc1$, according to the $cond$ field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken; that is, the instruction causes a PC-relative, delayed control transfer to the address “ $PC + (4 \times \text{sign_ext}(\text{disp19}))$ ”. If FALSE, the branch is not taken.

If a conditional branch is taken, the delay instruction is always executed regardless of the value of the annul (a) bit. If a conditional branch is not taken and the annul bit is 1 ($a = 1$), the delay instruction is annulled (not executed).

Note | The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

The predict bit (p) is used to give the hardware a hint about whether the branch is expected to be taken. A 1 in the p bit indicates that the branch is expected to be taken; a 0 indicates that the branch is expected not to be taken.

Annulment, delay instructions, prediction, and delayed control transfers are described further in Chapter 6, *Instruction Set Overview*.

Programming Note | For optimal performance, a DCTI should not be placed in the instruction word immediately following an annulled branch-always instruction. For additional information, see *DCTI Couples* on page 89.

An attempt to execute a BPcc instruction with $cc0 = 1$ (a reserved value) causes an *illegal_instruction* exception.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20), $PSTATE.tct = 1$, and the BPcc instruction will cause a transfer of control (BPA or taken conditional branch), then BPcc generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the BPcc) is stored in $TPC[TL]$ and the value of NPC from before the BPcc was executed is stored in $TNPC[TL]$.

Note that BPN never causes a *control_transfer_instruction* exception.

Exceptions *illegal_instruction*
control_transfer_instruction (impl. dep. #450-S20)

■ *See Also* Branch on Integer Register with Prediction (BPr) on page 128

7.12 Branch on Integer Register with Prediction (BPr)

Instruction	rcond	Operation	Register Contents Test	Assembly Language Syntax	Class
—	000	<i>Reserved</i>	—	—	—
BRZ	001	Branch on Register Zero	$R[rs1] = 0$	<code>brz {,a}{,pt ,pn} <i>reg_{rs1}</i>, label</code>	A1
BRLEZ	010	Branch on Register Less Than or Equal to Zero	$R[rs1] \leq 0$	<code>brlez {,a}{,pt ,pn} <i>reg_{rs1}</i>, label</code>	A1
BRLZ	011	Branch on Register Less Than Zero	$R[rs1] < 0$	<code>brlz {,a}{,pt ,pn} <i>reg_{rs1}</i>, label</code>	A1
—	100	<i>Reserved</i>	—	—	—
BRNZ	101	Branch on Register Not Zero	$R[rs1] \neq 0$	<code>brnz {,a}{,pt ,pn} <i>reg_{rs1}</i>, label</code>	A1
BRGZ	110	Branch on Register Greater Than Zero	$R[rs1] > 0$	<code>brgz {,a}{,pt ,pn} <i>reg_{rs1}</i>, label</code>	A1
BRGEZ	111	Branch on Register Greater Than or Equal to Zero	$R[rs1] \geq 0$	<code>brgez {,a}{,pt ,pn} <i>reg_{rs1}</i>, label</code>	A1



† Some very early SPARC V9 implementations (circa 1995) ignored the value of bit 28 and executed a BPr instruction even if bit 28 = 1. Since then, all implementations have treated bits 31:30 = 00₂, bit 28 = 1, and bits 24:22 = 011₂ as an *illegal_instruction* or, in the case of implementations of Oracle SPARC Architecture 2011 and later, as a CBcond instruction (see page 136).

Programming Note To set the annul (a) bit for BPr instructions, append “, a” to the opcode mnemonic. For example, use “brz , a %i3, label.” In the preceding table, braces signify that the “, a” is optional. To set the branch prediction bit p, append either “, pt” for predict taken or “, pn” for predict not taken to the opcode mnemonic. If neither “, pt” nor “, pn” is specified, the assembler defaults to “, pt”.

Description These instructions branch based on the contents of R[rs1]. They treat the register contents as a signed integer value.

A BPr instruction examines all 64 bits of R[rs1] according to the rcond field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken; that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign_ext(d16hi :: d16lo))”. If FALSE, the branch is not taken.

If the branch is taken, the delay instruction is always executed, regardless of the value of the annul (a) bit. If the branch is not taken and the annul bit is 1 (a = 1), the delay instruction is annulled (not executed).

The predict bit (p) gives the hardware a hint about whether the branch is expected to be taken. If p = 1, the branch is expected to be taken; p = 0 indicates that the branch is expected not to be taken.

An attempt to execute a BPr instruction when instruction bit 28 = 1 or rcond is a reserved value (000₂ or 100₂) causes an *illegal_instruction* exception.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20), PSTATE.tct = 1, and the BPr instruction will cause a transfer of control (taken conditional branch), then BPr generates a *control_transfer_instruction* exception instead of causing a control transfer.

BPr

Annulment, delay instructions, prediction, and delayed control transfers are described further in Chapter 6, *Instruction Set Overview*.

Implementation Note If this instruction is implemented by tagging each register value with an N (negative) bit and Z (zero) bit, the table below can be used to determine if `rcond` is TRUE:

<u>Branch</u>	<u>Test</u>
BRNZ	not Z
BRZ	Z
BRGEZ	not N
BRLZ	N
BRLEZ	N or Z
BRGZ	not (N or Z)

Exceptions *illegal_instruction*
control_transfer_instruction (impl. dep. #450-S20)

■ *See Also* Branch on Integer Condition Codes with Prediction (BPcc) on page 126

CALL

7.13 Call and Link

Instruction	OP	Operation	Assembly Language Syntax	Class
CALL	01	Call and Link	<code>call label</code>	A1



Description The CALL instruction causes an unconditional, delayed, PC-relative control transfer to address $PC + (4 \times \text{sign_ext}(\text{disp30}))$. Since the word displacement (`disp30`) field is 30 bits wide, the target address lies within a range of -2^{31} to $+2^{31} - 4$ bytes. The PC-relative displacement is formed by sign-extending the 30-bit word displacement field to 62 bits and appending two low-order zeroes to obtain a 64-bit byte displacement.

The CALL instruction also writes the value of PC, which contains the address of the CALL, into R[15] (*out* register 7).

When `PSTATE.am = 1`, the more-significant 32 bits of the target instruction address are masked out (set to 0) before being sent to the memory system and in the address written into R[15]. (closed impl. dep. #125-V9-Cs10)

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20) and `PSTATE.tct = 1`, then CALL generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the CALL instruction) is stored in TPC[TL] and the value of NPC from before the CALL was executed is stored in TNPC[TL]. The full 64-bit (nonmasked) PC and NPC values are stored in TPC[TL] and TNPC[TL], regardless of the value of `PSTATE.am`.

Exceptions *control_transfer_instruction* (impl. dep. #450-S20)

■ *See Also* JMPL on page 238

CAMELLIA 4-operand Op

7.14 Camellia Operations (4 operand) Crypto

The Camellia instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	op5	Operation	Assembly Language Syntax	Class
CAMELLIA_F	1100	Camellia F operation	<code>camellia_f</code> <i>freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}</i>	N1



Description Camellia is a block cipher that produces a 128-bit output from a 128-bit input under the control of a 128, 192, or 256-bit key. The specifications for Camellia can be found at:
<http://info.isl.ntt.co.jp/crypt/eng/camellia/specifications.html>

The F function is applied 18 times for a 128-bit key and 24 times for a 192 or 256-bit key. Between groups of six F functions, the FL and FLI functions are applied to the upper and lower halves of the data respectively. XOR operations are applied at the beginning and end to complete the cipher. The temporary variable `data_f` is used in the functional description below and represents the result of applying the F function to the `rs3` data using the `rs1` key. The Camellia instructions operate on 64-bit floating-point registers.

CAMELLIA_F:

```

data_f{63:0} ← camellia_F( data=FD[rs3]{63:0} , key=FD[rs1]{63:0} );
FD[rd]{63:0} ← data_f{63:0} xor FD[rs2]{63:0};
    
```

CAMELLIA 4-operand Op

Programming Note The Camellia instructions are components of the overall Camellia algorithm. To perform an encryption or decryption, software must first expand the key. Key expansion is done only once per session key. The expanded keys are then applied to all blocks for that session. The following instruction sequence is an example of a 128-bit encrypt for one block. The expanded keys are loaded in F0 thru F50. The initial text is loaded in F_D[54] and F_D[52].

```
fxor      %f0, %f54, %f54    !# Pre-Whiten
fxor      %f2, %f52, %f52
camellia_f %f4, %f52, %f54, %f52    !# Round 1  F 1
camellia_f %f6, %f54, %f52, %f54    !# Round 1  F 2
camellia_f %f8, %f52, %f54, %f52    !# Round 1  F 3
camellia_f %f10, %f54, %f52, %f54   !# Round 1  F 4
camellia_f %f12, %f52, %f54, %f52   !# Round 1  F 5
camellia_f %f14, %f54, %f52, %f54   !# Round 1  F 6
camellia_fl %f16, %f54, %f54        !# FL
camellia_fli %f18, %f52, %f52       !# FLI
camellia_f %f20, %f52, %f54, %f52   !# Round 2  F 1
camellia_f %f22, %f54, %f52, %f54   !# Round 2  F 2
camellia_f %f24, %f52, %f54, %f52   !# Round 2  F 3
camellia_f %f26, %f54, %f52, %f54   !# Round 2  F 4
camellia_f %f28, %f52, %f54, %f52   !# Round 2  F 5
camellia_f %f30, %f54, %f52, %f54   !# Round 2  F 6
camellia_fl %f32, %f54, %f54        !# FL
camellia_fli %f34, %f52, %f52       !# FLI
camellia_f %f36, %f52, %f54, %f52   !# Round 3  F 1
camellia_f %f38, %f54, %f52, %f54   !# Round 3  F 2
camellia_f %f40, %f52, %f54, %f52   !# Round 3  F 3
camellia_f %f42, %f54, %f52, %f54   !# Round 3  F 4
camellia_f %f44, %f52, %f54, %f52   !# Round 3  F 5
camellia_f %f46, %f54, %f52, %f54   !# Round 3  F 6
fxor      %f48, %f52, %f52        !# Post-Whiten
fxor      %f50, %f54, %f54
```

If CFR.camellia = 0, an attempt to execute a CAMELLIA instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that CFR.camellia = 1 before executing any of these CAMELLIA instructions. If CFR.camellia = 0, then software should assume that an attempt to execute one of the CAMELLIA instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.

Therefore, if CFR.camellia = 0, software should perform the corresponding CAMELLIA operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*


■ *See Also* *Camellia Operations (3 Operand)* on page 133

Camellia 3-operand Ops

7.15 Camellia Operations (3 Operand) Crypto

The Camellia instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class
CAMELLIA_FL	1 0011 1100	Camellia FL operation	<code>camellia_fl</code> <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	N1
CAMELLIA_FLI	1 0011 1101	Camellia FLI operation	<code>camellia_fli</code> <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	N1



Description The CAMELLIA instructions operate on 64-bit floating-point registers.

CAMELLIA_FL: $F_D[rd]\{63:0\} \leftarrow \text{camellia FL}(\text{data}=F_D[rs2]\{63:0\}, \text{key}=F_D[rs1]\{63:0\});$

CAMELLIA_FLI: $F_D[rd]\{63:0\} \leftarrow \text{camellia FLI}(\text{data}=F_D[rs2]\{63:0\}, \text{key}=F_D[rs1]\{63:0\});$

If `CFR.camellia = 0`, an attempt to execute a CAMELLIA instruction causes a *compatibility_feature* exception.


Programming Note Software *must* check that `CFR.camellia = 1` before executing either of these CAMELLIA instructions. If `CFR.camellia = 0`, then software should assume that an attempt to execute the instruction instruction either
 (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
 (2) will execute, but perform some other operation.
 Therefore, if `CFR.camellia = 0`, software should perform the CAMELLIA operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions `fp_disabled`

See Also *Camellia Operations (4 operand)* on page 131

7.16 Compare and Swap

Instruction	op3	Operation	Assembly Language Syntax	Class
CASA ^{P_{ASI}}	11 1100	Compare and Swap Word from Alternate Space	<code>casa [reg_{rs1}] imm_asi, reg_{rs2}, reg_{rd}</code> <code>casa [reg_{rs1}] %asi, reg_{rs2}, reg_{rd}</code>	A1
CASXA ^{P_{ASI}}	11 1110	Compare and Swap Extended from Alternate Space	<code>casxa [reg_{rs1}] imm_asi, reg_{rs2}, reg_{rd}</code> <code>casxa [reg_{rs1}] %asi, reg_{rs2}, reg_{rd}</code>	A1



Description Concurrent processes use Compare-and-Swap instructions for synchronization and memory updates. Uses of compare-and-swap include spin-lock operations, updates of shared counters, and updates of linked-list pointers. The last two can use wait-free (nonlocking) protocols.

The CASXA instruction compares the value in register R[rs2] with the doubleword in memory pointed to by the doubleword address in R[rs1].

- If the values are equal, the value in R[rd] is swapped with the doubleword pointed to by the doubleword address in R[rs1].
- If the values are not equal, the contents of the doubleword pointed to by R[rs1] replaces the value in R[rd], but the memory location remains unchanged.

The CASA instruction compares the low-order 32 bits of register R[rs2] with a word in memory pointed to by the word address in R[rs1].

- If the values are equal, then the low-order 32 bits of register R[rd] are swapped with the contents of the memory word pointed to by the address in R[rs1] and the high-order 32 bits of register R[rd] are set to 0.
- If the values are not equal, the memory location remains unchanged, but the contents of the memory word pointed to by R[rs1] replace the low-order 32 bits of R[rd] and the high-order 32 bits of register R[rd] are set to 0.

A compare-and-swap instruction comprises three operations: a load, a compare, and a swap. The overall instruction is atomic; that is, no intervening interrupts or deferred traps are recognized by the virtual processor and no intervening update resulting from a compare-and-swap, swap, load, load-store unsigned byte, or store instruction to the doubleword containing the addressed location, or any portion of it, is performed by the memory system.

A compare-and-swap operation behaves as if it performs a store, either of a new value from R[rd] or of the previous value in memory. The addressed location must be writable, even if the values in memory and R[rs2] are not equal.

If $i = 0$, the address space of the memory location is specified in the `imm_asi` field; if $i = 1$, the address space is specified in the ASI register.

Exceptions. An attempt to execute a CASXA or CASA instruction when $i = 1$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

CASA / CASXA

Exceptions.. A *mem_address_not_aligned* exception is generated if the address in R[rs1] is not properly aligned.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, CASXA and CASA cause a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, CASXA and CASA cause a *privileged_action* exception.

Compatibility Note | An implementation might cause an exception because of an error during the store memory access, even though there was no error during the load memory access.

Programming Note | Compare and Swap (CAS) and Compare and Swap Extended (CASX) synthetic instructions are available for “big endian” memory accesses. Compare and Swap Little (CASL) and Compare and Swap Extended Little (CASXL) synthetic instructions are available for “little endian” memory accesses. See *Synthetic Instructions* on page 536 for the syntax of these synthetic instructions.

The compare-and-swap instructions do not affect the condition codes.

The compare-and-swap instructions can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with these instructions causes a *DAE_invalid_asi* exception.

ASIs valid for CASA and CASXA instructions	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE

Exceptions

- illegal_instruction*
- mem_address_not_aligned*
- privileged_action*
- VA_watchpoint*
- DAE_invalid_asi*
- DAE_privilege_violation*
- DAE_nc_page* (attempted access to noncacheable page)
- DAE_nfo_page* (attempted access to non-faulting-only page)

See Also

- CASA on page 134
- LDSTUB on page 255
- LDSTUBA on page 256
- SWAP on page 359
- SWAPA on page 360

CBcond (Compare and Branch)

7.17 Compare and Branch

Opcode	cc2	c_hi :: c_lo	Operation	Test	Assembly Language Syntax	Class
—	1000		<i>Reserved (illegal_instruction)</i>	—	—	—
—	0000		<i>Reserved (illegal_instruction)</i>	—	—	—



32-bit Compare and Branch Operations

CWBNE	0	1001	Compare and Branch if Not Equal	not Z	cwbne [†] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBE	0	0001	Compare and Branch if Equal	Z	cwbe [‡] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBG	0	1010	Compare and Branch if Greater	not (Z or (N xor V))	cwbg <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBLE	0	0010	Compare and Branch if Less or Equal	Z or (N xor V)	cwble <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBGE	0	1011	Compare and Branch if Greater or Equal	not (N xor V)	cwbge <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBL	0	0011	Compare and Branch if Less	N xor V	cwbl <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBGU	0	1100	Compare and Branch if Greater Unsigned	not (C or Z)	cwbg <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBLEU	0	0100	Compare and Branch if Less or Equal Unsigned	C or Z	cwbleu <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBCC	0	1101	Compare and Branch if Carry Clear (Greater Than or Equal, Unsigned)	not C	cwbcc [◊] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBCS	0	0101	Compare and Branch if Carry Set (Less Than, Unsigned)	C	cwbcs [∇] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBPOS	0	1110	Compare and Branch if Positive	not N	cwbpos <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBNEG	0	0110	Compare and Branch if Negative	N	cwbneg <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBVC	0	1111	Compare and Branch if Overflow Clear	not V	cwbvc <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CWBVS	0	0111	Compare and Branch if Overflow Set	V	cwbvs <i>reg_{rs1}, reg_or_imm5, label</i>	C1

64-bit Compare and Branch Operations

CXBNE	1	1001	Compare and Branch if Not Equal	not Z	cxbne [†] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBE	1	0001	Compare and Branch if Equal	Z	cxbe [‡] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBG	1	1010	Compare and Branch if Greater	not (Z or (N xor V))	cxbg <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBLE	1	0010	Compare and Branch if Less or Equal	Z or (N xor V)	cxble <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBGE	1	1011	Compare and Branch if Greater or Equal	not (N xor V)	cxbge <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBL	1	0011	Compare and Branch if Less	N xor V	cxbl <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBGU	1	1100	Compare and Branch if Greater Unsigned	not (C or Z)	cxbg <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBLEU	1	0100	Compare and Branch if Less or Equal Unsigned	C or Z	cxbleu <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBCC	1	1101	Compare and Branch if Carry Clear (Greater Than or Equal, Unsigned)	not C	cxbcc [◊] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBCS	1	0101	Compare and Branch if Carry Set (Less Than, Unsigned)	C	cxbcs [∇] <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBPOS	1	1110	Compare and Branch if Positive	not N	cxbpos <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBNEG	1	0110	Compare and Branch if Negative	N	cxbneg <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBVC	1	1111	Compare and Branch if Overflow Clear	not V	cxbvc <i>reg_{rs1}, reg_or_imm5, label</i>	C1
CXBVS	1	0111	Compare and Branch if Overflow Set	V	cxbvs <i>reg_{rs1}, reg_or_imm5, label</i>	C1

[†] *synonym: c<w|x>bnz* [‡] *synonym: c<w|x>bz* [◊] *synonym: c<w|x>bgeu* [∇] *synonym: c<w|x>b1u*



[†] Very early SPARC V9 implementations (circa 1995) ignored bit 28 and executed a BPr instruction when bits 31:30 = 00₂, bit 28 = 1, and bits 24:22 = 011₂. Later implementations treated that as an *illegal_instruction*. Starting with OSA 2011, it a CBcond instruction.

CBcond (Compare and Branch)

Description The compare-and-branch instruction compares two integer values, producing a TRUE or FALSE result.

If the comparison evaluates to FALSE, no control transfer (branch) occurs.

If the comparison evaluates to TRUE, the instruction causes a PC-relative, non-delayed transfer of control (branch) to address “PC + (4 × sign_ext(d10_hi :: d10_lo))”.

Since the word displacement available in this instruction is 10 bits, a CBcond instruction can target a PC-relative range of -512 to +511 words (-2048 to +2044 bytes).

The first comparison operand is always R[rs1]. If i = 0, the second comparison operand is R[rs2]; if i = 1, the second comparison operand is sign_ext(simm5).

If cc2 = 0 (the versions of these instructions with “W” in their mnemonic), the comparison is performed between the least-significant 32 bits of the two source operands, using 32-bit arithmetic semantics. If cc2 = 1 (the versions of these instructions with “X” in their mnemonic), the comparison is performed between all 64 bits of the two source operands, using 64-bit arithmetic semantics.

The CBcond instruction does not change architecturally-visible condition codes (the contents of the CCR register). The comparison (“SUBcc-like”) component of CBcond produces temporary condition codes, which are evaluated by its branch component, then discarded.

Programming Note The Compare and Branch (CBcond) instructions operate similarly to an integer comparison (SUBcc), followed by a conditional branch (Bicc), with a few differences:

- (1) CBcond uses only *temporary* condition codes (CCR is not updated)
- (2) CBcond implements a *non*-delayed control transfer (while Bicc is a dCTI),
- (3) CBcond has a shorter branch range (displacement) than Bicc, and
- (4) CBcond can perform 32- or 64-bit comparisons, while Bicc only tests 32-bit condition codes (BPcc can also test both 32- and 64-bit condition codes)

Given those caveats, then if i = 0, a CBcond instruction conceptually operates similarly to:

```
subcc reg_rs1, reg_rs2, %g0    !! subcc component; CCR not updated
bcond cc2, target              !! branch component; non-delayed transfer
```

And if i = 1, a CBcond instruction conceptually operates similarly to:

```
subcc reg_rs1, simm5, %g0    !! subcc component; CCR not updated
bcond cc2, target            !! branch component; non-delayed transfer
```

Programming Notes

- Because CBcond is a non-delayed control transfer instruction, it has no delay slot and no annul bit.
- CBcond neither needs, nor has, a static branch prediction bit.
- It is possible to construct CBcond to produce a non-delayed Branch Always instruction, using CBcond with two zero-value operands and a branch-if-equal condition. When i = 0, rs1 = rs2, CWBE and CXBE are effectively non-delayed “Branch Always” instructions.
- One can compare register contents with zero and branch based on that result by using a CBcond instruction with i = 1 and (imm5 = 0).

CBcond (Compare and Branch)

Programming Note For optimal performance, CBcond should not be the next sequential instruction executed after any other control transfer instruction (including another CBcond).

From the above, it follows that for optimal performance in switch statements (the code for which tends to contain a sequence of multiple compare-and-branch operations), any sequence of CBcond's should be separated by NOP(s), such as:

```
CBcond
nop
CBcond
nop
CBcond
...
```

Implementation Note • The concatenation `c_hi :: c_lo` encodes identical conditions to the `cond` field in other conditional branch instructions.

Exceptions. An attempt to execute a CBcond instruction when `c_lo = 0002` causes an *illegal_instruction* exception.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20), `PSTATE.tct = 1`, and the CBcond instruction will cause a transfer of control (taken conditional branch-and-compare), then CBcond generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the CBcond instruction) is stored in `TPC[TL]` and the value of NPC from before the CBcond was executed is stored in `TNPC[TL]`.

Exceptions *illegal_instruction*
control_transfer_instruction (impl. dep. #450-S20)

CMASK<8|16|32>

7.18 CMASK VIS 3

The CMASK instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
CMASK8	0 0001 1011	Set GSR.mask based on 8-bit mask	<code>cmask8 <i>reg_{rs2}</i></code>	C1	OSA 2011
CMASK16	0 0001 1101	Set GSR.mask based on 4-bit mask	<code>cmask16 <i>reg_{rs2}</i></code>	C1	OSA 2011
CMASK32	0 0001 1111	Set GSR.mask based on 2-bit mask	<code>cmask32 <i>reg_{rs2}</i></code>	C1	OSA 2011



Description The CMASK instructions compute a 32-bit set of byte-selection indexes and place the result in the mask field of the GSR register.

For CMASK8, CMASK16, and CMASK32, the input is contained in the least-significant 8, 4, and 2 bits, respectively, of R[rs2]; other bits of R[rs2] are ignored and should be set to 0 by software. The resultant GSR.mask values are listed in TABLE 7-7.

TABLE 7-7 Source Operand Effect on GSR.mask

R[rs2] bit	Value	CMASK8	CMASK16	CMASK32
R[rs2]{0}	0	GSR.mask{3:0} ← F ₁₆	GSR.mask{7:0} ← EF ₁₆	GSR.mask{15:0} ← CDEF ₁₆
	1	GSR.mask{3:0} ← 7 ₁₆	GSR.mask{7:0} ← 67 ₁₆	GSR.mask{15:0} ← 4567 ₁₆
R[rs2]{1}	0	GSR.mask{7:4} ← E ₁₆	GSR.mask{15:8} ← CD ₁₆	GSR.mask{31:16} ← 89AB ₁₆
	1	GSR.mask{7:4} ← 6 ₁₆	GSR.mask{15:8} ← 45 ₁₆	GSR.mask{31:16} ← 0123 ₁₆
R[rs2]{2}	0	GSR.mask{11:8} ← D ₁₆	GSR.mask{23:16} ← AB ₁₆	
	1	GSR.mask{11:8} ← 5 ₁₆	GSR.mask{23:16} ← 23 ₁₆	
R[rs2]{3}	0	GSR.mask{15:12} ← C ₁₆	GSR.mask{31:24} ← 89 ₁₆	
	1	GSR.mask{15:12} ← 4 ₁₆	GSR.mask{31:24} ← 01 ₁₆	
R[rs2]{4}	0	GSR.mask{19:16} ← B ₁₆		
	1	GSR.mask{19:16} ← 3 ₁₆		
R[rs2]{5}	0	GSR.mask{23:20} ← A ₁₆		
	1	GSR.mask{23:20} ← 2 ₁₆		
R[rs2]{6}	0	GSR.mask{27:24} ← 9 ₁₆		
	1	GSR.mask{27:24} ← 1 ₁₆		
R[rs2]{7}	0	GSR.mask{31:28} ← 8 ₁₆		
	1	GSR.mask{31:28} ← 0 ₁₆		

Note GSR.mask is generated such that if the corresponding input bit of CMASK in R[rs2] is 0, a subsequent BSHUFFLE instruction will select the corresponding byte from F_D[rs2]; if it is 1, a subsequent BSHUFFLE will select the corresponding byte from F_D[rs1].

CMASK<8|16|32>

Programming Note | It is envisioned that these instructions will consume the result of a prior SIMD compare instruction and generate the mask field necessary to allow the BSHUFFLE instruction to perform a conditional move. For example:

```
fcmpgt32 %f0, %f2, %i2
cmask32  %i2
bshuffle %f0, %f2, %f4
```

An attempt to execute a CMASK instruction when instruction bits 29:25 are nonzero or bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a CMASK instruction causes an *fp_disabled* exception.

Exceptions *illegal_instruction*
 fp_disabled

See Also Byte Mask and Shuffle on page 125
 Partitioned Signed Compare } on page 207
 FPCMPU on page 210

CRC32C

7.19 CRC32C Operation (3 operand) Crypto

The CR32C instructions is new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class
CRC32C ^N	1 0100 0111	two CRC32c operations	<i>crc32c</i> <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	N1



Description Accumulate a cyclic redundancy check (CRC) value using the polynomial 11EDC6F41₁₆ for two 32-bit input data blocks. The specification for this polynomial can be found at:

<http://www.faqs.org/rfcs/rfc3385.html>

The number of CRC32C instructions needed for a message length M is equal to “(M mod 8)” if evenly divisible by 8, or “(M mod 8) + 1” otherwise.

Software assistance is needed calculate the CRC32 value for a message whose length is not evenly divisible by 8. For cases where the remainder is 5, 6, or 7, the CRC32C instruction can be used if the remaining bytes are zeroed. For cases where the remainder is 1, 2, 3, or 4, software must process the last block, as CRC32C will not deliver a useful result.

The CRC32C algorithm calls for an inversion of the result on the final message block. This final inversion should be added after the final block is processed. See the Programming Note below for an example.

The CRC32C instruction operates on source values read from 64-bit floating-point registers F_D[rs1] and F_D[rs2], as follows:

Source Value	Location
IV	F _D [rs1]{31:0}
data1	F _D [rs2]{63:32}
data2	F _D [rs2]{31:0}

CRC32C: result_1{31:0} ← CRC32c(IV, data1);
 F_D[rd]{31:0} ← CRC32c(result_1, data2);
 F_D[rd]{63:32} ← 0000 0000₁₆;

CRC32C

Programming Note | The code example below performs the CRC32C algorithm across 32 bytes (four doublewords).

```
Let data = ( B31 :: B30 :: B29 :: ... :: B1 :: B0 )
           where Bn is byte n,
           B31 is the most significant byte, and
           B0 is the least significant byte

Let FD[2] = ( B31 :: B30 :: B29 :: B28 :: B27 :: B26 :: B25 :: B24 )
Let FD[4] = ( B23 :: B22 :: B21 :: B20 :: B19 :: B18 :: B17 :: B16 )
Let FD[6] = ( B15 :: B14 :: B13 :: B12 :: B11 :: B10 :: B9 :: B8 )
Let FD[8] = ( B7 :: B6 :: B5 :: B4 :: B3 :: B2 :: B1 :: B0 )

foned    %f0                !# initial IV = FFFF FFFF16
crc32c   %f0, %f2, %f0
crc32c   %f0, %f4, %f0
crc32c   %f0, %f6, %f0
crc32c   %f0, %f8, %f0
fnot1    %f0,                %f0 !# invert for final result
```

Upon completion of this code sequence, the final result will be in F[0]{31:0}.

If CFR.crc32c = 0, an attempt to execute a CRC32C instruction causes a *compatibility_feature* exception.

Exceptions *fp_disabled*

DES Crypto (4-operand)

7.20 DES Cryptographic Operations (4 operand) Crypto

The DES instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	op5	Operation	Assembly Language Syntax	Class
DES_ROUND	1001	two DES round operations	des_round <i>freq_{rs1}, freq_{rs2}, freq_{rs3}, freq_{rd}</i>	N1



Description The DES instructions support encryption and decryption for the Data Encryption Standard. This standard is available as FIPS-46 on “<http://nist.gov>”. Encryption and decryption start by applying Initial Permutation function (DES_IP) to the text. Encryption then loops on the DES round (DES_ROUND) applying the expanded keys in the forward direction. Decryption loops on the DES round applying the expanded keys in reverse order. Encryption and decryption end by applying the Inverse Initial Permutation function (DES_IIP). PC-2 for both keys is included in DES_ROUND. The DES instructions operate on 64-bit floating-point registers.

DES_ROUND:

- Perform two rounds of the DES algorithm.
- $F_D[rs1]$ is the expanded key for the first round.
- $F_D[rs2]$ is the expanded key for the second round.
- $F_D[rs3]$ is the text for the first round.

Programming Note

The DES instructions are components of the overall DES algorithm. To perform an encryption or decryption, software must first expand the key(s). Key expansion is done only once per session key. Prior to processing a DES block using these instructions, the DES key must go thru PC-1. After that, the PC-1 output is used to create 16 different keys using the shifts described by the DES algorithm. These 16 keys are then applied in pairs to the DES_ROUND instructions.

Note that these 16 keys are not equivalent to “k1” thru “k16” as described in the spec until PC-2 is applied as part of the DES_ROUND instruction. The expanded keys are then applied to all blocks for that session.

DES Crypto (4-operand)

For each block, the following instruction sequence can be used to encrypt or decrypt. Decryption applies the expanded keys in reverse order from encryption. For the following example, the expanded keys are stored in $F_D[0]$ thru $F_D[30]$ and text is in $F_D[32]$:

```
des_ip      %f32,          %f32
des_round   %f0,  %f2,  %f32, %f32
des_round   %f4,  %f6,  %f32, %f32
des_round   %f8,  %f10, %f32, %f32
des_round   %f12, %f14, %f32, %f32
des_round   %f16, %f18, %f32, %f32
des_round   %f20, %f22, %f32, %f32
des_round   %f24, %f26, %f32, %f32
des_round   %f28, %f30, %f32, %f32
des_iip     %f32,          %f32
```

If $CFR.des = 0$, an attempt to execute a DES instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that $CFR.des = 1$ before executing any of these DES instructions. If $CFR.des = 0$, then software should assume that an attempt to execute one of the DES instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.

Therefore, if $CFR.des = 0$, software should perform the corresponding DES operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*


■ *See Also* *DES Cryptographic Operations (2 operand)* on page 145

DES Crypto (2-operand)

7.21 DES Cryptographic Operations (2 operand) Crypto

The DES instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class
DES_IP	1 0011 0100	DES Initial Permutation	<code>des_ip</code> <i>freg_{rs1}, freg_{rd}</i>	N1
DES_IIP	1 0011 0101	DES Inverse Initial Permutation	<code>des_iip</code> <i>freg_{rs1}, freg_{rd}</i>	N1
DES_KEXPAND	1 0011 0110	DES Key Expand	<code>des_kexpand</code> <i>freg_{rs1}, imm5, freg_{rd}</i>	N1



Description The DES instructions operate on 64-bit floating-point registers.

DES_IP : Performs the Initial Permutation function on $F_D[rs1]$.
imm5 = 00000₂; other values of imm5 cause an *illegal_instruction* exception.

DES_IIP : Performs a 32-bit swap and then the Inverse Initial Permutation function on $F_D[rs1]$.
imm5 = 00000₂; other values of imm5 cause an *illegal_instruction* exception.

DES_KEXPAND:

```

if (imm5 = 000002)
then
    pc1[55:0] ← des_pc1(  $F_D[rs1]\{63:0\}$  );
     $F_D[rd]\{63:56\} \leftarrow 0000\ 0000_{16}$ ;
     $F_D[rd]\{55:0\} \leftarrow ( pc1\{54:28\} :: pc1\{55\} :: pc1\{26:0\} :: pc1\{27\} )$ 
else_if (imm5 = 000012)
then
     $F_D[rd]\{63:56\} \leftarrow 0000\ 0000_{16}$ ;
     $F_D[rd]\{55:0\} \leftarrow ( F_D[rs1]\{54:28\} :: F_D[rs1]\{55\} :: F_D[rs1]\{26:0\} :: F_D[rs1]\{27\} )$ 
else_if (imm5 = 000102)
then
     $F_D[rd]\{63:56\} \leftarrow 0000\ 0000_{16}$ ;
     $F_D[rd]\{55:0\} \leftarrow ( F_D[rs1]\{53:28\} :: F_D[rs1]\{55:54\} :: F_D[rs1]\{25:0\} :: F_D[rs1]\{27:26\} )$ 
else_if (imm5 = 000112)
then
     $F_D[rd]\{63:56\} \leftarrow 0000\ 0000_{16}$ ;
     $F_D[rd]\{55:0\} \leftarrow ( F_D[rs1]\{51:28\} :: F_D[rs1]\{55:52\} :: F_D[rs1]\{23:0\} :: F_D[rs1]\{27:24\} )$ 
else
    illegal_instruction exception
endif

```

Programming Note | The DES instructions are components of the overall DES algorithm. The DES_KEXPAND instruction can be used as an alternative to software key expansion. For each session, the following instruction sequence can be used to expand the key.

DES Crypto (2-operand)

In the following example, the original key is in $F_D[0]$ and the expanded keys are stored in $F_D[0]$ through $F_D[30]$.

```
des_kexpand    %f0 , 0, %f0
des_kexpand    %f0 , 1, %f2
des_kexpand    %f2 , 3, %f6
des_kexpand    %f2 , 2, %f4
des_kexpand    %f6 , 3, %f10
des_kexpand    %f6 , 2, %f8
des_kexpand    %f10, 3, %f14
des_kexpand    %f10, 2, %f12
des_kexpand    %f14, 1, %f16
des_kexpand    %f16, 3, %f20
des_kexpand    %f16, 2, %f18
des_kexpand    %f20, 3, %f24
des_kexpand    %f20, 2, %f22
des_kexpand    %f24, 3, %f28
des_kexpand    %f24, 2, %f26
des_kexpand    %f28, 1, %f30
```

If $CFR.des = 0$, an attempt to execute a DES instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that $CFR.des = 1$ before executing any of these DES instructions. If $CFR.des = 0$, then software should assume that an attempt to execute one of the DES instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.

Therefore, if $CFR.des = 0$, software should perform the corresponding DES operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*

■ *See Also* *DES Cryptographic Operations (4 operand)* on page 143

DONE

7.22 DONE

Instruction	op3	Operation	Assembly Language Syntax	Class
DONE ^P	11 1110	Return from Trap (skip trapped instruction)	done	A1



Description The DONE instruction restores the saved state from TSTATE[TL] (GL, CCR, ASI, PSTATE, and CWP), sets PC and NPC, and decrements TL. DONE sets $PC \leftarrow TNPC[TL]$ and $NPC \leftarrow TNPC[TL] + 4$ (normally, the value of NPC saved at the time of the original trap and address of the instruction immediately after the one referenced by the NPC).

Programming Notes The DONE and RETRY instructions are used to return from privileged trap handlers.
Unlike RETRY, DONE ignores the contents of TPC[TL].

If the saved TNPC[TL] was not altered by trap handler software, DONE causes execution to resume immediately *after* the instruction that originally caused the trap (as if that instruction was “done” executing).

Execution of a DONE instruction in the delay slot of a control-transfer instruction produces undefined results.

If software writes invalid or inconsistent state to TSTATE before executing DONE, virtual processor behavior during and after execution of the DONE instruction is undefined.

Note that since PSTATE.tct is automatically set to 0 during entry to a trap handler, execution of a DONE instruction at the end of a trap handler will not cause a *control_transfer_instruction* exception unless trap handler software has explicitly set PSTATE.tct to 1. During execution of the DONE instruction, the value of PSTATE.tct is restored from TSTATE.

Programming Notes If *control_transfer_instruction* traps are to be re-enabled (PSTATE.tct \leftarrow 1, restored from TSTATE[TL].pstate.tct) when trap handler software for the *control_transfer_instruction* trap returns, the trap handler must

- (1) emulate the trapped CTI, setting TPC[TL] and TNPC[TL] appropriately, remembering to compensate for annul bits) and
- (2) use a DONE (not RETRY) instruction to return.

If the CTI that caused the *control_transfer_instruction* trap was a DONE (RETRY) instruction, the trap handler must carefully emulate the trapped DONE (RETRY) (decrementing TL may suffice) before the trap handler returns using its own DONE (RETRY) instruction.

When PSTATE.am = 1, the more-significant 32 bits of the target instruction address are masked out (set to 0) before being sent to the memory system.

IMPL. DEP. #417-S10: If (1) TSTATE[TL].pstate.am = 1 and (2) a DONE instruction is executed (which sets PSTATE.am to '1' by restoring the value from TSTATE[TL].pstate.am to PSTATE.am), it is implementation dependent whether the DONE instruction masks (zeroes) the more-significant 32 bits of the values it places into PC and NPC.

DONE

Exceptions. In privileged mode (PSTATE.priv = 1), an attempt to execute DONE while TL = 0 causes an *illegal_instruction* exception. An attempt to execute DONE (in any mode) with instruction bits 18:0 nonzero causes an *illegal_instruction* exception.

In nonprivileged mode (PSTATE.priv = 0), an attempt to execute DONE causes a *privileged_opcode* exception.

Implementation | In nonprivileged mode, *illegal_instruction* exception due to TL = 0
Note | does not occur. The *privileged_opcode* exception occurs instead,
| regardless of the current trap level (TL).

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20) and PSTATE.tct = 1, then DONE generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the DONE instruction) is stored in TPC[TL] and the value of NPC from before the DONE was executed is stored in TNPC[TL]. The full 64-bit (nonmasked) PC and NPC values are stored in TPC[TL] and TNPC[TL], regardless of the value of PSTATE.am.

Exceptions *illegal_instruction*
 privileged_opcode
 control_transfer_instruction (impl. dep. #450-S20)

■ *See Also* RETRY on page 318

7.23 Edge Handling Instructions VIS 1

Instruction	opf	Operation	Assembly Language Syntax †		Class
EDGE8cc	0 0000 0000	Eight 8-bit edge boundary processing	edge8cc	$reg_{rs1}, reg_{rs2}, reg_{rd}$	B1
EDGE8Lcc	0 0000 0010	Eight 8-bit edge boundary processing, little-endian	edge8lcc	$reg_{rs1}, reg_{rs2}, reg_{rd}$	B1
EDGE16cc	0 0000 0100	Four 16-bit edge boundary processing	edge16cc	$reg_{rs1}, reg_{rs2}, reg_{rd}$	B1
EDGE16Lcc	0 0000 0110	Four 16-bit edge boundary processing, little-endian	edge16lcc	$reg_{rs1}, reg_{rs2}, reg_{rd}$	B1
EDGE32cc	0 0000 1000	Two 32-bit edge boundary processing	edge32cc	$reg_{rs1}, reg_{rs2}, reg_{rd}$	B1
EDGE32Lcc	0 0000 1010	Two 32-bit edge boundary processing, little-endian	edge32lcc	$reg_{rs1}, reg_{rs2}, reg_{rd}$	B1

† The original assembly language mnemonics for these instructions did not include the “cc” suffix, as appears in the names of all other instructions that set the integer condition codes. The old, non-“cc” mnemonics are deprecated. Over time, assemblers will support the new mnemonics for these instructions. In the meantime, some older assemblers may recognize only the mnemonics, without “cc”.



Description EDGE8[L]cc, EDGE16[L]cc, and EDGE32[L]cc operate identically to EDGE8[L]N, EDGE16[L]N, and EDGE32[L]N, respectively, but also set the integer condition codes (CCR.xcc and CCR.icc).

The integer condition codes are set by these instructions the same as a SUBcc instruction with the same operands (see *Subtract* on page 303).

Exceptions None

■ *See Also* EDGE<8|16|32>[L]N on page 150

7.24 Edge Handling Instructions (no CC) VIS 2

Instruction	opf	Operation	Assembly Language Syntax	Class
EDGE8N	0 0000 0001	Eight 8-bit edge boundary processing, no CC	edge8n <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1
EDGE8LN	0 0000 0011	Eight 8-bit edge boundary processing, little-endian, no CC	edge8ln <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1
EDGE16N	0 0000 0101	Four 16-bit edge boundary processing, no CC	edge16n <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1
EDGE16LN	0 0000 0111	Four 16-bit edge boundary processing, little-endian, no CC	edge16ln <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1
EDGE32N	0 0000 1001	Two 32-bit edge boundary processing, no CC	edge32n <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1
EDGE32LN	0 0000 1011	Two 32-bit edge boundary processing, little-endian, no CC	edge32ln <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	B1



Description These instructions handle the boundary conditions for parallel pixel scan line loops, where R[rs1] is the address of the next pixel to render and R[rs2] is the address of the last pixel in the scan line.

EDGE8LN, EDGE16LN, and EDGE32LN are little-endian versions of EDGE8N, EDGE16N, and EDGE32N, respectively. They produce an edge mask that is bit-reversed from their big-endian counterparts but are otherwise identical. This makes the mask consistent with the mask produced by the *Partitioned Unsigned Compare* on page 210 and consumed by the Partial Store instruction (see *DAE_invalid_asi eDAE_invalid_asi eStore Partial Floating-Point* on page 347) for little-endian data.

A 2-bit (EDGE32cc), 4-bit (EDGE16cc), or 8-bit (EDGE8cc) pixel mask is stored in the least significant bits of R[rd]. The mask is computed from left and right edge masks as follows:

1. The left edge mask is computed from the 3 least significant bits of R[rs1] (see TABLE 7-8) and the right edge mask is computed from the 3 least significant bits of R[rs2] (see TABLE 7-9).
2. If 32-bit address masking is **disabled** (PSTATE.am = 0) (that is, 64-bit addressing is in use) and the most significant 61 bits of R[rs1] are equal to the corresponding bits in R[rs2], R[rd] is set to the right edge mask **anded** with the left edge mask.
3. If 32-bit address masking is **enabled** (PSTATE.am = 1) (that is, 32-bit addressing is in use) and bits 31:3 of R[rs1] match bits 31:3 of R[rs2], R[rd] is set to the right edge mask **anded** with the left edge mask.
4. Otherwise, R[rd] is set to the left edge mask.

TABLE 7-8 and TABLE 7-9 list edge mask specifications.

TABLE 7-8 Left Edge Mask Specification

Edge Size	R[rs1] {2:0}	Left Edge	
		Big Endian	Little Endian
8	000	1111 1111	1111 1111
8	001	0111 1111	1111 1110
8	010	0011 1111	1111 1100
8	011	0001 1111	1111 1000
8	100	0000 1111	1111 0000
8	101	0000 0111	1110 0000
8	110	0000 0011	1100 0000

EDGE<8|16|32>{L}N

TABLE 7-8 Left Edge Mask Specification (Continued)

Edge Size	R[rs1] {2:0}	Left Edge	
		Big Endian	Little Endian
8	111	0000 0001	1000 0000
16	00x	1111	1111
16	01x	0111	1110
16	10x	0011	1100
16	11x	0001	1000
32	0xx	11	11
32	1xx	01	10

TABLE 7-9 Right Edge Mask Specification


Edge Size	R[rs2] {2:0}	Right Edge	
		Big Endian	Little Endian
8	000	1000 0000	0000 0001
8	001	1100 0000	0000 0011
8	010	1110 0000	0000 0111
8	011	1111 0000	0000 1111
8	100	1111 1000	0001 1111
8	101	1111 1100	0011 1111
8	110	1111 1110	0111 1111
8	111	1111 1111	1111 1111
16	00x	1000	0001
16	01x	1100	0011
16	10x	1110	0111
16	11x	1111	1111
32	0xx	10	01
32	1xx	11	11

Exceptions None

■ *See Also* EDGE<8,16,32>[L]cc on page 149

7.25 Floating-Point Absolute Value

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FABSs	11 0100	0 0000 1001	Absolute Value Single	<code>f abss</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FABSd	11 0100	0 0000 1010	Absolute Value Double	<code>f absd</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FABSq	11 0100	0 0000 1011	Absolute Value Quad	<code>f absq</code> <i>freq_{rs2}, freq_{rd}</i>	C3



Description FABS copies the source floating-point register(s) to the destination floating-point register(s), with the sign bit cleared (set to 0).

FABSs operates on single-precision (32-bit) floating-point registers, FABSd operates on double-precision (64-bit) floating-point register pairs, and FABSq operates on quad-precision (128-bit) floating-point register quadruples.

These instructions clear (set to 0) both `FSR.cexc` and `FSR.ftt`. They do not round, do not modify `FSR.aexc`, and do not treat floating-point NaN values differently from other floating-point values.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FABSq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FABS instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FABS instruction causes an *fp_disabled* exception.


An attempt to execute an FABSq instruction when `rs2{1} ≠ 0` or `rd{1} ≠ 0` causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (`FSR.ftt = invalid_fp_register` (FABSq only))

FADD

7.26 Floating-Point Add

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FADDs	11 0100	0 0100 0001	Add Single	<code>fadds</code> <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FADDd	11 0100	0 0100 0010	Add Double	<code>faddd</code> <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FADDq	11 0100	0 0100 0011	Add Quad	<code>faddq</code> <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C3



Description The floating-point add instructions add the floating-point register(s) specified by the `rs1` field and the floating-point register(s) specified by the `rs2` field. The instructions then write the sum into the floating-point register(s) specified by the `rd` field.

Rounding is performed as specified by `FSR.rd`.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute a `FADDq` instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an `FADD` instruction causes an *fp_disabled* exception.

An attempt to execute an `FADDq` instruction when (`rs1{1} ≠ 0`) or (`rs2{1} ≠ 0`) or (`rd{1:0} ≠ 0`) causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

Note An *fp_exception_other* with `FSR.ftt = unfinished_FPop` can occur if the operation detects unusual, implementation-specific conditions.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (`FSR.ftt = invalid_fp_register` (`FADDq` only))
fp_exception_other (`FSR.ftt = unfinished_FPop`)
fp_exception_ieee_754 (OF, UF, NX, NV)


See Also FMAf on page 175

FALIGNDATAg

7.27 Align Data (using GSR.align) VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
FALIGNDATAg †	0 0100 1000	Perform data alignment for misaligned data, using alignment indicated in GSR.align	<code>faligndata <i>freq_{rs1}</i>, <i>freq_{rs2}</i>, <i>freq_{rd}</i></code>	A1

† The original architectural name for this instruction was “FALIGNDATA”; the “g” was later appended for clarity.



Description FALIGNDATAg concatenates two 64-bit floating-point registers $F_D[rs1]$ and $F_D[rs2]$, to form a 128-bit (16-byte) intermediate value. The contents of the first source operand form the more-significant 8 bytes of the intermediate value, and the contents of the second source operand form the less significant 8 bytes of the intermediate value. Bytes in the intermediate value are numbered from most significant (byte 0) to least significant (byte 15). Eight bytes are extracted from the intermediate value and stored in the 64-bit floating-point destination register, $F_D[rd]$. GSR.align specifies the number of the most significant byte to extract (and, therefore, the least significant byte extracted is numbered GSR.align+7).

Programming Note FALIGNDATAg relies on a single GSR.align field, effectively limiting its utility to a single stream of misaligned data. FALIGNDATAi is much more suitable for handling multiple streams of misaligned data.

GSR.align is normally set by a previous ALIGNADDRESS instruction.




FIGURE 7-6 FALIGNDATAg

A byte-aligned 64-bit load can be performed as shown below.

```

alignaddr  UnalignedAddress, Offset, Address!set GSR.align
ldd        [Address], %d0
ldd        [Address + 8], %d2
faligndata %d0, %d2, %d4      !use GSR.align to select bytes
    
```

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FALIGNDATAg instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

See Also Align Address on page 117
Align Data (using Integer register) on page 155

FALIGNDATAi

7.28 Align Data (using Integer register)

The FALIGNDATAi instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
FALIGNDATAi	0 0100 1001	Perform data alignment for misaligned data, using alignment indicated in an integer register VIS 4	<code>faligndata reg_{rs1}, freg_{rd}, freg_{rs2}, freg_{rd}</code>	C1	OSA 2015



Description FALIGNDATAi concatenates two 64-bit floating-point registers, $F_D[rd]$ and $F_D[rs2]$, to form a 128-bit (16-byte) intermediate value. The contents of $F_D[rd]$ form the more-significant 8 bytes of the intermediate value, and the contents of $F_D[rs2]$ form the less significant 8 bytes of the intermediate value. Bytes in the intermediate value are numbered from most significant (byte 0) to least significant (byte 15). Eight bytes are extracted from the intermediate value and stored in the 64-bit floating-point destination register specified by $F_D[rd]$. $R[rs1]\{2:0\}$ specifies the number of the most significant byte to extract (and, therefore, the least significant byte extracted is numbered $R[rs1]\{2:0\}+7$).

Programming Note FALIGNDATAi is a “destructive” instruction, in that its second and fourth operands are identical, so register $F_D[rd]$ serves as both a source and a destination register ($F_D[rd]$ is read-modify-write). The requirement that the second and fourth assembly-language operands refer to the same 64-bit (double-precision) floating-point register is enforced by the assembler.

Programming Note Any integer register can be used to select the result bytes for FALIGNDATAi. Thus, it can be used to byte-align multiple streams of misaligned data — in contrast with the original FALIGNDATAg instruction, which works best with only one misaligned data stream.




FIGURE 7-7 FALIGNDATAi

FALIGNDATAi

A byte-aligned 64-bit load can be performed as shown below.

```
andn      UnalignedAddress, 0x07, Address !get aligned address
and       UnalignedAddress, 0x07, %o5    !set alignment in %o5
ldd       [Address], %d0
ldd       [Address + 8], %d2
faligndata %o5, %d0, %d2, %d0 !%o5 selects alignment offset
```

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FALIGNDATAg instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

■ *See Also* Align Data (using gsr.align) on page 154

7.29 Branch on Floating-Point Condition Codes (FBfcc)

Opcode	cond	Operation	fcc Test	Assembly Language Syntax	Class
FBA ^D	1000	Branch Always	1	fba{, a} label	A1
FBN ^D	0000	Branch Never	0	fbn{, a} label	A1
FBU ^D	0111	Branch on Unordered	U	fbu{, a} label	A1
FBG ^D	0110	Branch on Greater	G	fbg{, a} label	A1
FBUG ^D	0101	Branch on Unordered or Greater	G or U	fbug{, a} label	A1
FBL ^D	0100	Branch on Less	L	fb1{, a} label	A1
FBUL ^D	0011	Branch on Unordered or Less	L or U	fbul{, a} label	A1
FBLG ^D	0010	Branch on Less or Greater	L or G	fb1g{, a} label	A1
FBNE ^D	0001	Branch on Not Equal	L or G or U	fbne [†] {, a} label	A1
FBE ^D	1001	Branch on Equal	E	fbe [‡] {, a} label	A1
FBUE ^D	1010	Branch on Unordered or Equal	E or U	fbue{, a} label	A1
FBGE ^D	1011	Branch on Greater or Equal	E or G	fbge{, a} label	A1
FBUGE ^D	1100	Branch on Unordered or Greater or Equal	E or G or U	fbuge{, a} label	A1
FBLE ^D	1101	Branch on Less or Equal	E or L	fb1e{, a} label	A1
FBULE ^D	1110	Branch on Unordered or Less or Equal	E or L or U	fbule{, a} label	A1
FBO ^D	1111	Branch on Ordered	E or L or G	fbo{, a} label	A1

[†] synonym: fbnz [‡] synonym: fbnz



Programming Note To set the annul (a) bit for FBfcc instructions, append “, a” to the opcode mnemonic. For example, use “fb1, a label”. In the preceding table, braces around “, a” signify that “, a” is optional.

Description Unconditional and Fcc branches are described below:

- **Unconditional branches** (FBA, FBN) — If its annul field is 0, an FBN (Branch Never) instruction acts like a NOP. If its annul field is 1, the following (delay) instruction is annulled (not executed) when the FBN is executed. In neither case does a transfer of control take place.
 FBA (Branch Always) causes a PC-relative, delayed control transfer to the address “PC + (4 × sign_ext(dispatch22))” regardless of the value of the floating-point condition code bits. If the annul field of the branch instruction is 1, the delay instruction is annulled (not executed). If the annul (a) bit is 0, the delay instruction is executed.
- **Fcc-conditional branches** — Conditional FBfcc instructions (except FBA and FBN) evaluate floating-point condition code zero (fcc0) according to the cond field of the instruction. Such evaluation produces either a TRUE or FALSE result. If TRUE, the branch is taken, that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign_ext(dispatch22))”. If FALSE, the branch is not taken.

FBfcc

If a conditional branch is taken, the delay instruction is always executed, regardless of the value of the annul (a) bit. If a conditional branch is not taken and the annul bit is 1 (a = 1), the delay instruction is annulled (not executed).

Note | The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

Annulment, delay instructions, and delayed control transfers are described further in Chapter 6.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FBfcc instruction causes an *fp_disabled* exception.


If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20), PSTATE.tct = 1, and the FBfcc instruction will cause a transfer of control (FBA or taken conditional branch), then FBfcc generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the FBfcc instruction) is stored in TPC[TL] and the value of NPC from before the FBfcc was executed is stored in TNPC[TL]. Note that FBN never causes a *control_transfer_instruction* exception.

Exceptions *fp_disabled*
control_transfer_instruction (impl. dep. #450-S20)

7.30 Branch on Floating-Point Condition Codes with Prediction (FBPfcc)

Instruction	cond	Operation	fcc Test	Assembly Language Syntax	Class
FBPA	1000	Branch Always	1	<code>fb{a}, {a}{,pt ,pn} %fccn, label</code>	A1
FBPN	0000	Branch Never	0	<code>fbn{,a}{,pt ,pn} %fccn, label</code>	A1
FBPU	0111	Branch on Unordered	U	<code>fbu{,a}{,pt ,pn} %fccn, label</code>	A1
FBPG	0110	Branch on Greater	G	<code>fbg{,a}{,pt ,pn} %fccn, label</code>	A1
FBPUG	0101	Branch on Unordered or Greater	G or U	<code>fbug{,a}{,pt ,pn} %fccn, label</code>	A1
FBPL	0100	Branch on Less	L	<code>fb{,a}{,pt ,pn} %fccn, label</code>	A1
FBPUL	0011	Branch on Unordered or Less	L or U	<code>fbul{,a}{,pt ,pn} %fccn, label</code>	A1
FBPLG	0010	Branch on Less or Greater	L or G	<code>fb{,a}{,pt ,pn} %fccn, label</code>	A1
FBPNE	0001	Branch on Not Equal	L or G or U	<code>fbne[†]{,a}{,pt ,pn} %fccn, label</code>	A1
FBPE	1001	Branch on Equal	E	<code>fbe[‡]{,a}{,pt ,pn} %fccn, label</code>	A1
FBPUE	1010	Branch on Unordered or Equal	E or U	<code>fbue{,a}{,pt ,pn} %fccn, label</code>	A1
FBPGE	1011	Branch on Greater or Equal	E or G	<code>fbge{,a}{,pt ,pn} %fccn, label</code>	A1
FBPUGE	1100	Branch on Unordered or Greater or Equal	E or G or U	<code>fbug{,a}{,pt ,pn} %fccn, label</code>	A1
FBPLE	1101	Branch on Less or Equal	E or L	<code>fb{,a}{,pt ,pn} %fccn, label</code>	A1
FBPULE	1110	Branch on Unordered or Less or Equal	E or L or U	<code>fbule{,a}{,pt ,pn} %fccn, label</code>	A1
FBPO	1111	Branch on Ordered	E or L or G	<code>fbo{,a}{,pt ,pn} %fccn, label</code>	A1

† synonym: `fbnz` ‡ synonym: `fbz`



cc1	cc0	Condition Code
0	0	<code>fcc0</code>
0	1	<code>fcc1</code>
1	0	<code>fcc2</code>
1	1	<code>fcc3</code>

Programming Note To set the annul (*a*) bit for FBPfcc instructions, append “, *a*” to the opcode mnemonic. For example, use “`fb{,a} %fcc3, label`”. In the preceding table, braces signify that the “, *a*” is optional. To set the branch prediction bit, append either “, *pt*” (for predict taken) or “, *pn*” (for predict not taken) to the opcode mnemonic. If neither “, *pt*” nor “, *pn*” is specified, the assembler defaults to “, *pt*”. To select the appropriate floating-point condition code, include “`%fcc0`”, “`%fcc1`”, “`%fcc2`”, or “`%fcc3`” before the label.

Description Unconditional branches and Fcc-conditional branches are described below.

FBPfcc

- **Unconditional branches (FBPA, FBPN)** — If its annul field is 0, an FBPN (Floating-Point Branch Never with Prediction) instruction acts like a NOP. If the Branch Never’s annul field is 0, the following (delay) instruction is executed; if the annul (a) bit is 1, the following instruction is annulled (not executed). In no case does an FBPN cause a transfer of control to take place.

FBPA (Floating-Point Branch Always with Prediction) causes an unconditional PC-relative, delayed control transfer to the address “PC + (4 × sign_ext (disp19))”. If the annul field of the branch instruction is 1, the delay instruction is annulled (not executed). If the annul (a) bit is 0, the delay instruction is executed.

- **Fcc-conditional branches** — Conditional FBPfcc instructions (except FBPA and FBPN) evaluate one of the four floating-point condition codes (fcc0, fcc1, fcc2, fcc3) as selected by cc0 and cc1, according to the cond field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken, that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign_ext (disp19))”. If FALSE, the branch is not taken.

If a conditional branch is taken, the delay instruction is always executed, regardless of the value of the annul (a) bit. If a conditional branch is not taken and the annul bit is 1 (a = 1), the delay instruction is annulled (not executed).

Note | The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

The predict bit (p) gives the hardware a hint about whether the branch is expected to be taken. A 1 in the p bit indicates that the branch is expected to be taken. A 0 indicates that the branch is expected not to be taken.

Annulment, delay instructions, and delayed control transfers are described further in Chapter 6, *Instruction Set Overview*.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FBPfcc instruction causes an *fp_disabled* exception.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20), PSTATE.tct = 1, and the FBPfcc instruction will cause a transfer of control (FBPA or taken conditional branch), then FBPfcc generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the FBPfcc instruction) is stored in TPC[TL] and the value of NPC from before the FBPfcc was executed is stored in TNPC[TL]. Note that FBPN never causes a *control_transfer_instruction* exception.


Exceptions *fp_disabled*
control_transfer_instruction (impl. dep. #450-S20)

FCHKSM16

7.31 Checksum VIS 3

The Checksum instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FCHKSM16	0 0100 0100	16-bit partitioned checksum	f64	f64	f64	<code>fchkasm16 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	C1	OSA 2011



Description This instruction performs four 16-bit additions between the corresponding integer values contained in its 64-bit source operands ($F_D[rs1]$, $F_D[rs2]$). The carry out of each addition is added to the least significant bit of each intermediate sum to produce the final sum.

Programming Note This is useful for networking applications requiring 16-bit TCP/UDP checksums.

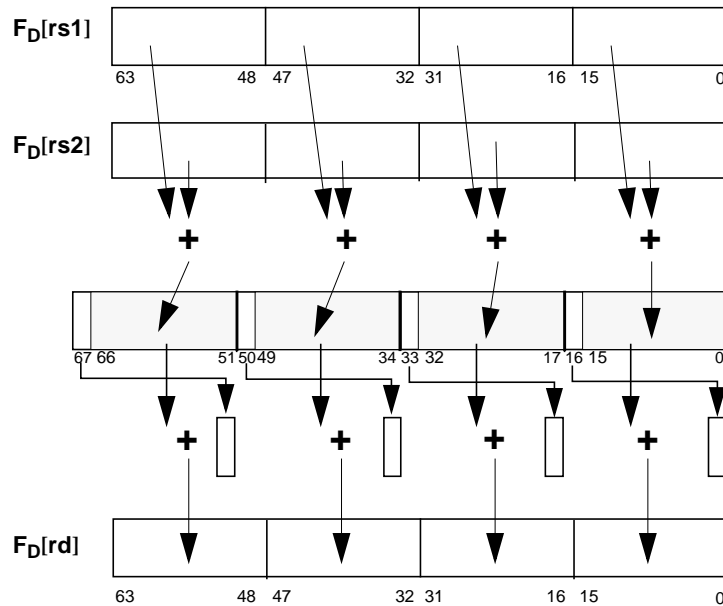


FIGURE 7-8 FCHKSM16 Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FCHKSUM16 instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

7.32 Floating-Point Compare

Instruction	opf	Operation	Assembly Language Syntax	Class
FCMPs	0 0101 0001	Compare Single	<code>fcmps %fccn, freg_{rs1}, freg_{rs2}</code>	A1
FCMPd	0 0101 0010	Compare Double	<code>fcmpd %fccn, freg_{rs1}, freg_{rs2}</code>	A1
FCMPq	0 0101 0011	Compare Quad	<code>fcmpq %fccn, freg_{rs1}, freg_{rs2}</code>	C3
FCMPEs	0 0101 0101	Compare Single and Exception if Unordered	<code>fcmpes %fccn, freg_{rs1}, freg_{rs2}</code>	A1
FCMPEd	0 0101 0110	Compare Double and Exception if Unordered	<code>fcmped %fccn, freg_{rs1}, freg_{rs2}</code>	A1
FCMPEq	0 0101 0111	Compare Quad and Exception if Unordered	<code>fcmp eq %fccn, freg_{rs1}, freg_{rs2}</code>	C3

[Description]



cc1	cc0	Condition Code
0	0	<code>fcc0</code>
0	1	<code>fcc1</code>
1	0	<code>fcc2</code>
1	1	<code>fcc3</code>

These instructions compare F[rs1] with F[rs2], and set the selected floating-point condition code (`fccn`) as follows

Relation	Resulting fcc value
$freg_{rs1} = freg_{rs2}$	0
$freg_{rs1} < freg_{rs2}$	1
$freg_{rs1} > freg_{rs2}$	2
$freg_{rs1} ? freg_{rs2}$ (unordered)	3

The “?” in the preceding table means that the compared values are unordered. The unordered condition occurs when one or both of the operands to the comparison is a signalling or quiet NaN

The “compare and cause exception if unordered” (FCMPEs, FCMPEd, and FCMPEq) instructions cause an invalid (NV) exception if either operand is a NaN.

FCMP<s|d|q>, FCMPE<s|d|q>

FCMP causes an invalid (NV) exception if either operand is a signalling NaN.

V8 Compatibility Note	Unlike the SPARC V8 architecture, SPARC V9 and the Oracle SPARC Architecture do not require an instruction between a floating-point compare operation and a floating-point branch (FBfcc, FBPfcc). SPARC V8 floating-point compare instructions are required to have rd = 0. In SPARC V9 and the Oracle SPARC Architecture, bits 26 and 25 of the instruction (rd{1:0}) specify the floating-point condition code to be set. Legal SPARC V8 code will work on SPARC V9 and the Oracle SPARC Architecture because the zeroes in the R[rd] field are interpreted as fcc0 and the FBfcc instruction branches based on the value of fcc0.
------------------------------	--

An attempt to execute an FCMP instruction when instruction bits 29:27 are nonzero causes an *illegal_instruction* exception.

Note	Oracle SPARC Architecture 2015 processors do not implement in hardware the instructions that refer to quad-precision floating-point registers. An attempt to execute FCMPq or FCMPEq generates an <i>illegal_instruction</i> exception, which causes a trap, allowing privileged software to emulate the instruction.
-------------	---

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FCMP or FCMPE instruction causes an *fp_disabled* exception.

An attempt to execute an FCMPq or FCMPEq instruction when (rs1{1} ≠ 0) or (rs2{1} ≠ 0) causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.


For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

<i>Exceptions</i>	<i>illegal_instruction</i> <i>fp_disabled</i> <i>fp_exception_ieee_754</i> (NV) <i>fp_exception_other</i> (FSR.ftt = invalid_fp_register (FCMPq, FCMPEq only))
-------------------	---

See Also	Partitioned Signed Compare } on page 207 Partitioned Unsigned Compare on page 210
-----------------	--

7.33 Floating-Point Divide

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FDIVs	11 0100	0 0100 1101	Divide Single	<code>fdivs <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1
FDIVd	11 0100	0 0100 1110	Divide Double	<code>fdivd <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1
FDIVq	11 0100	0 0100 1111	Divide Quad	<code>fdivq <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C3



Description The floating-point divide instructions divide the contents of the floating-point register(s) specified by the rs1 field by the contents of the floating-point register(s) specified by the rs2 field. The instructions then write the quotient into the floating-point register(s) specified by the rd field.

Rounding is performed as specified by FSR.rd.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware the instructions that refer to quad-precision floating-point registers. An attempt to execute an FDIVq instruction generates an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FCMP or FCMPE instruction causes an *fp_disabled* exception.

An attempt to execute an FADDq instruction when (rs1{1} ≠ 0) or (rs2{1} ≠ 0) causes an *fp_exception_other* (FSR.ftt = *invalid_fp_register*) exception.

Note For FDIVs and FDIVd, an *fp_exception_other* with FSR.ftt = *unfinished_FPop* can occur if the divide unit detects unusual, implementation-specific conditions.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions

- illegal_instruction*
- fp_disabled*
- fp_exception_other* (FSR.ftt = *invalid_fp_register* (FDIVq only))
- fp_exception_other* (FSR.ftt = *unfinished_FPop* (FDIVs, FDIV))
- fp_exception_ieee_754* (OF, UF, DZ, NV, NX)

FEXPAND

7.34 FEXPAND VIS 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FEXPAND	0 0100 1101	Four 16-bit expands	—	f32	f64	<code>fexpand <i>reg_rs2</i>, <i>reg_rd</i></code>	B1



Description FEXPAND takes four 8-bit unsigned integers from $F_S[rs2]$, converts each integer to a 16-bit fixed-point value, and stores the four resulting 16-bit values in a 64-bit floating-point register $F_D[rd]$. FIGURE 7-10 illustrates the operation.




FIGURE 7-9 FEXPAND Operation

This operation is carried out as follows:

1. Left-shift each 8-bit value by 4 and zero-extend that result on the left to a 16-bit value.
2. Store the result in the destination register, $F_D[rd]$.

Programming Note | FEXPAND performs the inverse of the FPACK16 operation.

An attempt to execute an FEXPAND instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FEXPAND instruction causes an *fp_disabled* exception.


Exceptions *illegal_instruction*
fp_disabled

See Also FPMERGE on page 216
FPACK on page 197

FHADD<s|d>

7.35 Floating-point Add and Halve VIS 3

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FHADDs	11 0100	0 0110 0001	Single-precision Add and Halve	<code>fhadds <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i></code>	C2
FHADDd	11 0100	0 0110 0010	Double-precision Add and Halve	<code>fhaddd <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i></code>	C2



Description The FHADD<s|d> instructions are used to perform an addition of two floating-point values and halve the result (divide it by 2) in a single operation. One benefit of this operation is that it cannot produce an arithmetic overflow.

Exceptions *fp_disabled*
fp_exception_ieee_754 (UF, NX, NV)

See Also FHSUB on page 167
FNHADD on page 194

FHSUB<s|d>

7.36 Floating-point Subtract and Halve VIS 3

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FHSUBs	11 0100	0 0110 0101	Single-precision Subtract and Halve	<code>fhsubs <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C2
FHSUBd	11 0100	0 0110 0110	Double-precision Subtract and Halve	<code>fhsupd <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C2



Description The FHSUB<s|d> instructions are used to take the difference of two floating-point values and halve the result (divide it by 2) in a single operation. One benefit of this operation is that it cannot produce an arithmetic overflow { is this statement TRUE, for the subtract version? }.


Exceptions *fp_disabled*
fp_exception_ieee_754 (UF, NX, NV)

See Also FHADD on page 166
 FNHADD on page 194

FiTO<s|d|q>

7.37 Convert 32-bit Integer to Floating Point

Instruction	op3	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FiTOs	11 0100	0 1100 0100	Convert 32-bit Integer to Single	—	f32	f32	fitos $freq_{rs2}, freq_{rd}$	A1
FiTOd	11 0100	0 1100 1000	Convert 32-bit Integer to Double	—	f32	f64	fitod $freq_{rs2}, freq_{rd}$	A1
FiTOq	11 0100	0 1100 1100	Convert 32-bit Integer to Quad	—	f32	f128	fitoq $freq_{rs2}, freq_{rd}$	C3



Description FiTOs, FiTOd, and FiTOq convert the 32-bit signed integer operand in floating-point register $F_S[rs2]$ into a floating-point number in the destination format. All write their result into the floating-point register(s) specified by rd.

The value of FSR.rd determines how rounding is performed by FiTOs.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute a FiTOq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FiTO<s|d|q> instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FiTO<s|d|q> instruction causes an *fp_disabled* exception.

An attempt to execute an FiTOq instruction when rd{1} ≠ 0 causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.

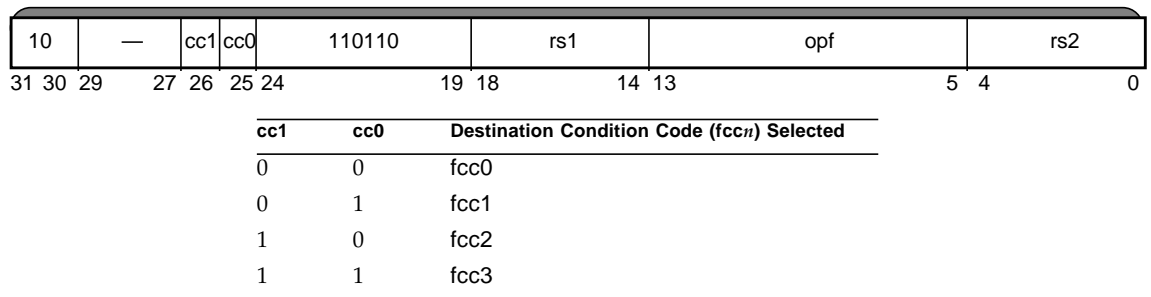
For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (FSR.ftt = invalid_fp_register (FiTOq))
fp_exception_ieee_754 (NX (FiTOs only))

7.38 Floating-Point Lexicographic Compare VIS 3

The FLCMP instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FLCMPs	1 0101 0001	Single-precision lexicographic compare	f32	f32	FSR.fccn	flcmps %fccn, freg _{rs1} , freg _{rs2}	C1
FLCMPd	1 0101 0010	Double-precision lexicographic compare	f64	f64	FSR.fccn	flcmpd %fccn, freg _{rs1} , freg _{rs2}	C1



Description Lexicographic comparisons differ from standard floating-point comparisons in that -0 and $+0$ are treated as distinct values, with -0 comparing as less than $+0$. The handling of NaN operands and the encoding of the comparison results in the fcc bits is also different.

These instructions compare the floating-point register(s) specified by the rs1 field with the floating-point register(s) specified by the rs2 field, and set the selected floating-point condition code (fcc_n) according to the following table:

Relation	Value Written to fcc _n
$F[rs1] \geq F[rs2]$, neither operand is NaN	0
$F[rs1] < F[rs2]$, neither operand is NaN	1
$F[rs1]$ is NaN, $F[rs2]$ is not NaN	2
$F[rs2]$ is NaN, regardless of $F[rs1]$	3

Programming Note A lexicographic compare instruction is not an FPop, so it:
 1) leaves all bits of FSR.cexc and FSR.aexc unchanged
 2) never generates an *fp_exception_other* exception

Programming Note The encoding for the results of FP lexicographic compares (written to fcc_n) differ from those for regular SPARC V9 floating-point compares. This implies that a move or branch instruction dependent on a floating-point lexicographic compare might function differently from what its instruction mnemonic implies.

FLCMP

TABLE 7-10 Effect of fcc Result Encoding of Floating-Point Lexicographic Compares on Dependent Conditional Moves/Branches

Relation Specified in Branch Opcode	Actual Functionality When fcc Source for Branch/move was Set by FLCMPs or FLCMPd
Always	Always
Never	Never
Unordered	rs2 source operand of FLCMP was NaN
Greater	rs1 source operand of FLCMP was NaN rs2 source operand of FLCMP was not NaN
Unordered or Greater	One or both the source operands rs1,rs2 of FLCMP was NaN
Less	Less
Unordered or Less	Less or rs2 source operand of FLCMP was NaN
Less or Greater	Less or (rs1 source operand of FLCMP was NaN and rs2 source operand was not NaN)
Not equal	Less Than, or one or both the source operands rs1,rs2 of FLCMP was NaN
Equal	Greater Than or Equal
Unordered or Equal	Greater Than or Equal or rs2 source of FLCMP was NaN
Greater or Equal	Greater Than or Equal or (rs1 source of FLCMP was NaN and rs2 source was not NaN)
Unordered or Greater or Equal	Greater Than or Equal or one or both the sources rs1, rs2 of FLCMP was NaN
Less or Equal	Less Than or Greater Than or Equal (i.e., neither source of FLCMP was NaN)
Unordered or Less oR Equal	Less Than or Greater Than or Equal or rs2 source of FLCMP was NaN
Ordered	Less Than or Greater Than or Equal or ((rs1 source of FLMCP was NaN) and (rs2 source was not NaN)) (that is, rs2 source was not NaN)

An attempt to execute an FLCMP instruction when instruction bits 29:27 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FLCMP instruction causes an *fp_disabled* exception.

Exceptions *illegal_instruction*
 fp_disabled

7.39 Flush Instruction Memory

Instruction	op3	Operation	Assembly Language Syntax†	Class
FLUSH	11 1011	Flush Instruction Memory	<code>flush [address]</code>	A1

† The original assembly language syntax for a FLUSH instruction (“`flush address`”) has been deprecated because of inconsistency with other SPARC assembly language syntax. Over time, assemblers will support the new syntax for this instruction. In the meantime, some existing assemblers may only recognize the original syntax.



Description FLUSH ensures that the aligned doubleword specified by the effective address is consistent across any local caches and, in a multiprocessor system, will eventually (impl. dep. #122-V9) become consistent everywhere.

The SPARC V9 instruction set architecture does not guarantee consistency between instruction memory and data memory. When software writes¹ to a memory location that may be executed as an instruction (self-modifying code²), a potential memory consistency problem arises, which is addressed by the FLUSH instruction. Use of FLUSH after instruction memory has been modified ensures that instruction and data memory are synchronized for the processor that issues the FLUSH instruction.

The virtual processor waits until all previous (cacheable) stores have completed before issuing a FLUSH instruction. For the purpose of memory ordering, a FLUSH instruction behaves like a store instruction.

In the following discussion P_{FLUSH} refers to the virtual processor that executed the FLUSH instruction.

FLUSH causes a synchronization within a virtual processor which ensures that instruction fetches from the specified effective address by P_{FLUSH} appear to execute after any loads, stores, and atomic load-stores to that address issued by P_{FLUSH} prior to the FLUSH. In a multiprocessor system, FLUSH also ensures that these values will eventually become visible to the instruction fetches of all other virtual processors in the system. With respect to MEMBAR-induced orderings, FLUSH behaves as if it is a store operation (see *Memory Barrier* on page 269).

Given any store S_A to address A , that precedes in memory order a FLUSH F_A to address A , that in turn precedes in memory order a store S_B to address B ; if any instruction I_B fetched from address B executes the instruction created by store S_B , then any instruction I_A that fetched from address A and that follows I_B in program order cannot execute any version of the instruction from address A that existed prior to the store S_A .

The preceding statement defines an ordering requirement to which Oracle SPARC Architecture processors comply. By using a FLUSH instruction between two stores that modify instructions, atomicity between the two stores is guaranteed such that any virtual processor executing the instruction modified by the later store will never fetch and/or execute the instruction before it was modified by the earlier store.

If $i = 0$, the effective address operand for the FLUSH instruction is “ $R[\text{rs1}] + R[\text{rs2}]$ ”; if $i = 1$, it is “ $R[\text{rs1}] + \text{sign_ext}(\text{simm13})$ ”. The three least-significant bits of the effective address are ignored; that is, the effective address always refers to an aligned doubleword.

¹ this includes use of store instructions (executed on the same or another virtual processor) that write to instruction memory, or any other means of writing into instruction memory (for example, DMA transfer)

² practiced, for example, by software such as debuggers and dynamic linkers

FLUSH

See implementation-specific documentation for details on specific implementations of the FLUSH instruction.

On an Oracle SPARC Architecture processor:

- A FLUSH instruction causes a synchronization within the virtual processor on which the FLUSH is executed, which flushes its instruction pipeline to ensure that no instruction already fetched has subsequently been modified in memory. Any other virtual processors on the same physical processor are unaffected by a FLUSH.
- Coherency between instruction and data memories may or may not be maintained by hardware.

IMPL. DEP. #409-S10: The implementation of the FLUSH instruction is implementation dependent. If the implementation automatically maintains consistency between instruction and data memory,

- (1) the FLUSH address is ignored and
- (2) the FLUSH instruction cannot cause any data access exceptions, because its effective address operand is not translated or used by the MMU.

On the other hand, if the implementation does *not* maintain consistency between instruction and data memory, the FLUSH address is used to access the MMU and the FLUSH instruction can cause data access exceptions.

Programming Note For portability across all SPARC V9 implementations, software must always supply the target effective address in FLUSH instructions.

- If the implementation contains instruction prefetch buffers:
 - the instruction prefetch buffer(s) are invalidated
 - instruction prefetching is suspended, but may resume starting with the instruction immediately following the FLUSH

Programming Notes

1. Typically, FLUSH is used in self-modifying code. The use of self-modifying code is discouraged.
2. If a program includes self-modifying code, to be portable it *must* issue a FLUSH instruction for each modified doubleword of instructions (or make a call to privileged software that has an equivalent effect) after storing into the instruction stream.
3. The order in which memory is modified can be controlled by means of FLUSH and MEMBAR instructions interspersed appropriately between stores and atomic load-stores. FLUSH is needed only between a store and a subsequent instruction fetch from the modified location. When multiple processes may concurrently modify live (that is, potentially executing) code, the programmer must ensure that the order of update maintains the program in a semantically correct form at all times.
4. The memory model guarantees in a uniprocessor that *data* loads observe the results of the most recent store, even if there is no intervening FLUSH.
5. FLUSH may be a time-consuming operation. (see the Implementation Note below)
6. In a multiprocessor system, the effects of a FLUSH operation will be globally visible before any subsequent store becomes globally visible.

FLUSH

7. FLUSH is designed to act on a doubleword. On some implementations, FLUSH may trap to system software. For these reasons, system software should provide a service routine, callable by nonprivileged software, for flushing arbitrarily-sized regions of memory. On some implementations, this routine would issue a series of FLUSH instructions; on others, it might issue a single trap to system software that would then flush the entire region.

8. FLUSH operates using the current (implicit) context. Therefore, a FLUSH executed in privileged mode will use the nucleus context and will not necessarily affect instruction cache lines containing data from a user (nonprivileged) context.

Implementation Note In a multiprocessor configuration, FLUSH requires all processors that may be referencing the addressed doubleword to flush their instruction caches, which is a potentially disruptive activity.

V9 Compatibility Note The effect of a FLUSH instruction as observed from the virtual processor on which FLUSH executes is immediate. Other virtual processors in a multiprocessor system eventually will see the effect of the FLUSH, but the latency is implementation dependent.

An attempt to execute a FLUSH instruction when instruction bits 29:25 are nonzero causes an *illegal_instruction* exception.

An attempt to execute a FLUSH instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
 DAE_nfo_page

FLUSHW

7.40 Flush Register Windows

Instruction	op3	Operation	Assembly Language Syntax	Class
FLUSHW	10 1011	Flush Register Windows	<code>flushw</code>	A1



Description FLUSHW causes all active register windows except the current window to be flushed to memory at locations determined by privileged software. FLUSHW behaves as a NOP if there are no active windows other than the current window. At the completion of the FLUSHW instruction, the only active register window is the current one.

Programming Note The FLUSHW instruction can be used by application software to flush register windows to memory so that it can switch memory stacks or examine register contents from previous stack frames.

FLUSHW acts as a NOP if $CANSAVE = N_REG_WINDOWS - 2$. Otherwise, there is more than one active window, so FLUSHW causes a spill exception. The trap vector for the spill exception is based on the contents of OTHERWIN and WSTATE. The spill trap handler is invoked with the CWP set to the window to be spilled (that is, $(CWP + CANSAVE + 2) \bmod N_REG_WINDOWS$). See *Register Window Management Instructions* on page 90.

Programming Note Typically, the spill handler saves a window on a memory stack and returns to reexecute the FLUSHW instruction. Thus, FLUSHW traps and reexecutes until all active windows other than the current window have been spilled.

An attempt to execute a FLUSHW instruction when instruction bits 29:25 or 18:0 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
spill_n_normal
spill_n_other

7.41 Floating-Point Multiply-Add and Multiply-Subtract (fused)

Instruction	op5	Operation	Assembly Language Syntax		Class	Added
FMADDs	00 01	Multiply-Add Single	<code>fmaddds</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FMADDd	00 10	Multiply-Add Double	<code>fmaddd</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FMSUBs	01 01	Multiply-Subtract Single	<code>fmsubs</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FMSUBd	01 10	Multiply-Subtract Double	<code>fmsubd</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FNMSUBs	10 01	Negative Multiply-Subtract Single	<code>fnmsubs</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FNMSUBd	10 10	Negative Multiply-Subtract Double	<code>fnmsubd</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FNMADDs	11 01	Negative Multiply-Add Single	<code>fnmadds</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007
FNMADDd	11 10	Negative Multiply-Add Double	<code>fnmaddd</code>	<code>freg_{rs1}, freg_{rs2}, freg_{rs3}, freg_{rd}</code>	B1	UA 2007



Instruction	Implementation
Multiply-Add (fused)	$F[rd] \leftarrow (F[rs1] \times F[rs2]) + F[rs3]$
Multiply-Subtract (fused)	$F[rd] \leftarrow (F[rs1] \times F[rs2]) - F[rs3]$
Negative Multiply-Add (fused)	$F[rd] \leftarrow -((F[rs1] \times F[rs2]) + F[rs3])$
Negative Multiply-Subtract (fused)	$F[rd] \leftarrow -((F[rs1] \times F[rs2]) - F[rs3])$

Description

The fused floating-point multiply-add instructions, FMADD<*s*|*d*>, perform a floating-point multiplication of the values in the floating-point registers specified by *rs1* and *rs2*, perform a floating-point addition of the resulting product to the value in the register(s) specified by *rs3*, round the resulting sum, and write the result into the floating-point register(s) specified by *rd*.

The fused floating-point multiply-subtract instructions, FMSUB<*s*|*d*>, perform a floating-point multiplication of the values in the floating-point registers specified by *rs1* and *rs2*, perform a floating-point subtraction from the resulting product of the value in the register(s) specified by *rs3*, round the resulting difference, and write the result into the floating-point register(s) specified by *rd*.

The fused floating-point negative multiply-add instructions, FNMADD<*s*|*d*>, perform a floating-point multiplication of the values in the floating-point registers specified by *rs1* and *rs2*, perform a floating-point addition of the resulting product to the value in the register(s) specified by *rs3*, negate the resulting sum, round the result, and write the result into the floating-point register(s) specified by *rd*.

The fused floating-point negative multiply-subtract instructions, FNMSUB<*s*|*d*>, perform a floating-point multiplication of the values in the floating-point registers specified by *rs1* and *rs2*, perform a floating-point subtraction from the resulting product of the value in the register(s) specified by *rs3*, negate the resulting difference, round the result, and write the result into the floating-point register(s) specified by *rd*.

All of the above instructions are “fused” operations; no rounding is performed between the multiplication operation and the subsequent addition (or subtraction). Therefore, at most one rounding step occurs per instruction executed.

The negative fused multiply-add/subtract instructions (FNM*) treat NaN values as follows:

- A source QNaN propagates with its sign bit unchanged

FMAf

- A generated (default response) QNaN result has a sign bit of zero
- A source SNaN that is converted to a QNaN result retains the sign bit of the source SNaN

Exceptions. If an FMAf instruction is not implemented in hardware, it generates an *illegal_instruction* exception, so that privileged software can emulate the instruction.

If the FPU is not enabled ($\text{FPRS.fef} = 0$ or $\text{PSTATE.pef} = 0$) or if no FPU is present, an attempt to execute an FMAf instruction causes an *fp_disabled* exception.

Overflow, underflow, and inexact exception bits within FSR.cexc and FSR.aexc are updated based on the final result of the operation and not on the intermediate result of the multiplication. The invalid operation exception bits within FSR.cexc and FSR.aexc are updated as if the multiplication and the addition/subtraction were performed using two individual instructions. An invalid operation exception is detected when any of the following conditions are true:

- A source operand ($\text{F}[\text{rs1}]$, $\text{F}[\text{rs2}]$, or $\text{F}[\text{rs3}]$) is a SNaN
- $\infty \times 0$
- $\infty - \infty$

If the instruction generates an IEEE-754 exception or exceptions for which the corresponding trap enable mask (FSR.tem) bits are set, an *fp_exception_ieee_754* exception and subsequent trap is generated.

If either the multiply or the add/subtract operation detects an unfinished_FPop condition (for example, due to a subnormal operand or final result), the Multiply-Add/Subtract instruction generates an *fp_exception_other* exception with $\text{FSR.ftt} = \text{unfinished_FPop}$. An *fp_exception_other* exception with $\text{FSR.ftt} = \text{unfinished_FPop}$ always takes precedence over an *fp_exception_ieee_754* exception. That is, if an *fp_exception_other* exception occurs due to an unfinished_FPop condition, the FSR.cexc and FSR.aexc fields remain unchanged even if a floating point IEEE 754 exception occurs during the multiply operation (regardless whether traps are enabled, via FSR.tem , for the IEEE exception) and the unfinished_FPop condition occurs during the subsequent add/subtract operation.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Semantic Definitions

FMADD:

- (1) $\text{tmp} \leftarrow \text{F}[\text{rs1}] \times \text{F}[\text{rs2}]$
- (2) $\text{tmp} \leftarrow \text{tmp} + \text{F}[\text{rs3}]$
- (3) $\text{F}[\text{rd}] \leftarrow \text{round}(\text{tmp})$

FMADD:

- (1) $\text{tmp} \leftarrow \text{F}[\text{rs1}] \times \text{F}[\text{rs2}]$
- (2) $\text{tmp} \leftarrow \text{tmp} + \text{F}[\text{rs3}]$
- (3) $\text{tmp} \leftarrow -\text{tmp}$
- (4) $\text{F}[\text{rd}] \leftarrow \text{round}(\text{tmp})$

FMSUB:

- (1) $\text{tmp} \leftarrow \text{F}[\text{rs1}] \times \text{F}[\text{rs2}]$
- (2) $\text{tmp} \leftarrow \text{tmp} - \text{F}[\text{rs3}]$
- (3) $\text{F}[\text{rd}] \leftarrow \text{round}(\text{tmp})$

FMSUB:

- (1) $\text{tmp} \leftarrow \text{F}[\text{rs1}] \times \text{F}[\text{rs2}]$
- (2) $\text{tmp} \leftarrow \text{tmp} - \text{F}[\text{rs3}]$
- (3) $\text{tmp} \leftarrow -\text{tmp}$
- (4) $\text{F}[\text{rd}] \leftarrow \text{round}(\text{tmp})$

Exceptions

fp_disabled
fp_exception_ieee_754 (OF, UF, NX, NV)
fp_exception_other ($\text{FSR.ftt} = \text{unfinished_FPop}$)

See Also

FMUL on page 190
FADD on page 153
FSUB on page 219

7.42 Partitioned Mean VIS 3

The FMEAN16 instructions is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FMEAN16	0 0100 0000	16-bit partitioned average	f64	f64	f64	fmean16 <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2011



Description These instructions perform four 16-bit partitioned arithmetic averages between pairs of corresponding signed integer values contained in the source operands ($F_D[rs1]$, $F_D[rs2]$).

The average of each pair of 16-bit operand values is calculated as

$$\text{trunc}((\text{sign_ext}(16\text{-bit operand1}) + \text{sign_ext}(16\text{-bit operand2}) + 1) \div 2)$$

where the **trunc** function returns the integer portion of the result of the division. FIGURE 7-10 shows the division operation as a right shift by one bit position.




FIGURE 7-10 FMEAN16 Operation

Programming Note This operation in effect performs integer division round-to-nearest (round up in case of a tie) for positive numbers, which is beneficial for MPEG/DVD motion compensation and motion estimation. It can be used for half-pel interpolation as well as combining motion prediction data and inverse discrete cosine transform data.

Note Since 17 bits are provided to hold the intermediate sum before right shifting, no overflow can occur.

FMEAN16


If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FMEAN16 instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

FMOV

7.43 Floating-Point Move

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FMOV _s	11 0100	0 0000 0001	Move (copy) Single	fmovs <i>freq_{rs2}, freq_{rd}</i>	A1
FMOV _d	11 0100	0 0000 0010	Move (copy) Double	fmovd <i>freq_{rs2}, freq_{rd}</i>	A1
FMOV _q	11 0100	0 0000 0011	Move (copy) Quad	fmovq <i>freq_{rs2}, freq_{rd}</i>	C3



Description FMOV copies the source floating-point register(s) to the destination floating-point register(s), unaltered.

FMOV_s, FMOV_d, and FMOV_q perform 32-bit, 64-bit, and 128-bit operations, respectively.

These instructions clear (set to 0) both FSR.cexc and FSR.ftt. They do not round, do not modify FSR.aexc, and do not treat floating-point NaN values differently from other floating-point values.

Programming Note If a 64-bit floating-point register-to-register copy is desired and simultaneous modification of FSR is not required, use of FSRC2d is strongly recommended over FMOV_d. FSRC2d is at least as fast as, and on many implementations much faster than, FMOV_d and FSRC1d.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FMOV_q instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FMOV instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FMOV instruction causes an *fp_disabled* exception.

An attempt to execute an FMOV_q instruction when rs2{1} ≠ 0 or rd{1} ≠ 0 causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.


Exceptions *illegal_instruction*
fp_disabled *fp_exception_other* (FSR.ftt = invalid_fp_register (FMOV_q only))

See Also *f Register Logical Operate (2 operand)* on page 226

FMOVcc

7.44 Move Floating-Point Register on Condition (FMOVcc)

Instruction	opf_low	Operation	Assembly Language Syntax	Class
FMOVSicc	00 0001	Move Floating-Point Single, based on 32-bit integer condition codes	<i>fmovsicc</i> %icc, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FMOVDicc	00 0010	Move Floating-Point Double, based on 32-bit integer condition codes	<i>fmovdicc</i> %icc, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FMOVQicc	00 0011	Move Floating-Point Quad, based on 32-bit integer condition codes	<i>fmovqicc</i> %icc, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C3
FMOVSxcc	00 0001	Move Floating-Point Single, based on 64-bit integer condition codes	<i>fmovsxcc</i> %xcc, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FMOVDxcc	00 0010	Move Floating-Point Double, based on 64-bit integer condition codes	<i>fmovdxcc</i> %xcc, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FMOVQxcc	00 0011	Move Floating-Point Quad, based on 64-bit integer condition codes	<i>fmovqxcc</i> %xcc, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C3
FMOVsfcc	00 0001	Move Floating-Point Single, based on floating-point condition codes	<i>fmovsfcc</i> %fccn, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FMOVDfcc	00 0010	Move Floating-Point Double, based on floating-point condition codes	<i>fmovdfcc</i> %fccn, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	A1
FMOVQfcc	00 0011	Move Floating-Point Quad, based on floating-point condition codes	<i>fmovqfcc</i> %fccn, <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C3



FMOVcc

Encoding of the *cond* Field for F.P. Moves Based on Integer Condition Codes (*icc* or *xcc*)

cond	Operation	icc / xcc Test	<i>icc/xcc</i> name(s) in Assembly Language Mnemonics
1000	Move Always	1	a
0000	Move Never	0	n
1001	Move if Not Equal	not Z	ne (or nz)
0001	Move if Equal	Z	e (or z)
1010	Move if Greater	not (Z or (N xor V))	g
0010	Move if Less or Equal	Z or (N xor V)	le
1011	Move if Greater or Equal	not (N xor V)	ge
0011	Move if Less	N xor V	l
1100	Move if Greater Unsigned	not (C or Z)	gu
0100	Move if Less or Equal Unsigned	(C or Z)	leu
1101	Move if Carry Clear (Greater or Equal, Unsigned)	not C	cc (or geu)
0101	Move if Carry Set (Less than, Unsigned)	C	cs (or lu)
1110	Move if Positive	not N	pos
0110	Move if Negative	N	neg
1111	Move if Overflow Clear	not V	vc
0111	Move if Overflow Set	V	vs

Encoding of the *cond* Field for F.P. Moves Based on Floating-Point Condition Codes (*fccn*)

cond	Operation	fccn Test	<i>fcc</i> name(s) in Assembly Language Mnemonics
1000	Move Always	1	a
0000	Move Never	0	n
0111	Move if Unordered	U	u
0110	Move if Greater	G	g
0101	Move if Unordered or Greater	G or U	ug
0100	Move if Less	L	l
0011	Move if Unordered or Less	L or U	ul
0010	Move if Less or Greater	L or G	lg
0001	Move if Not Equal	L or G or U	ne (or nz)
1001	Move if Equal	E	e (or z)
1010	Move if Unordered or Equal	E or U	ue
1011	Move if Greater or Equal	E or G	ge
1100	Move if Unordered or Greater or Equal	E or G or U	uge
1101	Move if Less or Equal	E or L	le
1110	Move if Unordered or Less or Equal	E or L or U	ule
1111	Move if Ordered	E or L or G	o

FMOVcc

Encoding of `opf_cc` Field (also see TABLE A-9 on page 491)

<code>opf_cc</code>	Instruction	Condition Code to be Tested
100 ₂	FMOV<s d q>icc	icc
110 ₂	FMOV<s d q>xcc	xcc
000 ₂	FMOV<s d q>fcc	fcc0
001 ₂		fcc1
010 ₂		fcc2
011 ₂		fcc3
101 ₂	<i>(illegal_instruction</i> exception)	
111 ₂		

Description

The FMOVcc instructions copy the floating-point register(s) specified by `rs2` to the floating-point register(s) specified by `rd` if the condition indicated by the `cond` field is satisfied by the selected floating-point condition code field in `FSR`. The condition code used is specified by the `opf_cc` field of the instruction. If the condition is `FALSE`, then the destination register(s) are not changed.

These instructions read, but do not modify, any condition codes.

These instructions clear (set to 0) both `FSR.cexc` and `FSR.ftt`. They do not round, do not modify `FSR.aexc`, and do not treat floating-point NaN values differently from other floating-point values.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FMOVQicc, FMOVQxcc, or FMOVQfcc instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FMOVcc instruction when instruction bit 18 is nonzero or `opf_cc` = 101₂ or 111₂ causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef` = 0 or `PSTATE.pef` = 0) or if no FPU is present, an attempt to execute an FMOVQicc, FMOVQxcc, or FMOVQfcc instruction causes an *fp_disabled* exception.

An attempt to execute an FMOVQicc, FMOVQxcc, or FMOVQfcc instruction when `rs2{1}` ≠ 0 or `rd{1}` ≠ 0 causes an *fp_exception_other* (`FSR.ftt` = `invalid_fp_register`) exception.

FMOVcc

Programming Note Branches cause the performance of most implementations to degrade significantly. Frequently, the MOVcc and FMOVcc instructions can be used to avoid branches. For example, the following C language segment:

```
double A, B, X;  
if (A > B) then X = 1.03; else X = 0.0;
```

can be coded as

```
! assume A is in %f0; B is in %f2; %xx points to  
! constant area  
    ldd    [%xx+C_1.03],%f4    ! X = 1.03  
    fcmpd %fcc3,%f0,%f2      ! A > B  
    fble,a %fcc3,label  
    ! following instruction only executed if the  
    ! preceding branch was taken  
    fsubd %f4,%f4,%f4        ! X = 0.0  
label:...
```

This code takes four instructions including a branch.

With FMOVcc, this could be coded as

```
    ldd    [%xx+C_1.03],%f4    ! X = 1.03  
    fsubd %f4,%f4,%f6        ! X' = 0.0  
    fcmpd %fcc3,%f0,%f2      ! A > B  
    fmovdle %fcc3,%f6,%f4    ! X = 0.0
```


This code also takes four instructions but requires no branches and may boost performance significantly. Use MOVcc and FMOVcc instead of branches wherever these instructions would improve performance.

Exceptions *illegal_instruction*
 fp_disabled
 fp_exception_other (FSR.ftt = invalid_fp_register (FMOVQ instructions))

FMOVR

7.45 Move Floating-Point Register on Integer Register Condition (FMOVR)

Instruction	rcond	opf_low	Operation	Test	Class
—	000	0 0101	<i>Reserved</i>	—	—
FMOVRsZ	001	0 0101	Move Single if Register = 0	R[rs1] = 0	A1
FMOVRsLEZ	010	0 0101	Move Single if Register ≤ 0	R[rs1] ≤ 0	A1
FMOVRsLZ	011	0 0101	Move Single if Register < 0	R[rs1] < 0	A1
—	100	0 0101	<i>Reserved</i>	—	—
FMOVRsNZ	101	0 0101	Move Single if Register ≠ 0	R[rs1] ≠ 0	A1
FMOVRsGZ	110	0 0101	Move Single if Register > 0	R[rs1] > 0	A1
FMOVRsGEZ	111	0 0101	Move Single if Register ≥ 0	R[rs1] ≥ 0	A1
—	000	0 0110	<i>Reserved</i>	—	—
FMOVRdZ	001	0 0110	Move Double if Register = 0	R[rs1] = 0	A1
FMOVRdLEZ	010	0 0110	Move Double if Register ≤ 0	R[rs1] ≤ 0	A1
FMOVRdLZ	011	0 0110	Move Double if Register < 0	R[rs1] < 0	A1
—	100	0 0110	<i>Reserved</i>	—	—
FMOVRdNZ	101	0 0110	Move Double if Register ≠ 0	R[rs1] ≠ 0	A1
FMOVRdGZ	110	0 0110	Move Double if Register > 0	R[rs1] > 0	A1
FMOVRdGEZ	111	0 0110	Move Double if Register ≥ 0	R[rs1] ≥ 0	A1
—	000	0 0111	<i>Reserved</i>	—	—
FMOVRqZ	001	0 0111	Move Quad if Register = 0	R[rs1] = 0	C3
FMOVRqLEZ	010	0 0111	Move Quad if Register ≤ 0	R[rs1] ≤ 0	C3
FMOVRqLZ	011	0 0111	Move Quad if Register < 0	R[rs1] < 0	C3
—	100	0 0111	<i>Reserved</i>	—	—
FMOVRqNZ	101	0 0111	Move Quad if Register ≠ 0	R[rs1] ≠ 0	C3
FMOVRqGZ	110	0 0111	Move Quad if Register > 0	R[rs1] > 0	C3
FMOVRqGEZ	111	0 0111	Move Quad if Register ≥ 0	R[rs1] ≥ 0	C3



Assembly Language Syntax

```

fmovr{s,d,q}z  regrs1, fregrs2, fregrd      (synonym: fmovr{s,d,q}e)
fmovr{s,d,q}lez regrs1, fregrs2, fregrd
fmovr{s,d,q}lz  regrs1, fregrs2, fregrd
fmovr{s,d,q}nz  regrs1, fregrs2, fregrd      (synonym: fmovr{s,d,q}ne)
fmovr{s,d,q}gz  regrs1, fregrs2, fregrd
fmovr{s,d,q}gez regrs1, fregrs2, fregrd

```


FMOVR

Description

If the contents of integer register R[rs1] satisfy the condition specified in the rcond field, these instructions copy the contents of the floating-point register(s) specified by the rs2 field to the floating-point register(s) specified by the rd field. If the contents of R[rs1] do not satisfy the condition, the floating-point register(s) specified by the rd field are not modified.

These instructions treat the integer register contents as a signed integer value; they do not modify any condition codes.

These instructions clear (set to 0) both FSR.cexc and FSR.ftt. They do not round, do not modify FSR.aexc, and do not treat floating-point NaN values differently from other floating-point values.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FMOVRq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FMOVR instruction when instruction bit 13 is nonzero or rcond = 000₂ or 100₂ causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FMOVR instruction causes an *fp_disabled* exception.

An attempt to execute an FMOVRq instruction when rs2{1} ≠ 0 or rd{1} ≠ 0 causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.

Implementation Note If this instruction is implemented by tagging each register value with an N (negative) and a Z (zero) condition bit, use the following table to determine whether rcond is TRUE:

Branch	Test
FMOVRNZ	not Z
FMOVRZ	Z
FMOVRGEZ	not N
FMOVRLZ	N
FMOVRLEZ	N or Z
FMOVRGZ	N nor Z


Exceptions

illegal_instruction
fp_disabled *fp_exception_other* (FSR.ftt = invalid_fp_register (FMOVRq instructions))

FMUL (partitioned)

7.46 Partitioned Multiply Instructions VIS 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FMUL8x16	0 0011 0001	Unsigned 8-bit by signed 16-bit partitioned product	f32	f64	f64	<code>fmul8x16 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FMUL8x16AU	0 0011 0011	Unsigned 8-bit by signed 16-bit upper α partitioned product	f32	f32	f64	<code>fmul8x16au freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FMUL8x16AL	0 0011 0101	Unsigned 8-bit by signed 16-bit lower α partitioned product	f32	f32	f64	<code>fmul8x16al freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FMUL8SUX16	0 0011 0110	Signed upper 8-bit by signed 16-bit partitioned product	f64	f64	f64	<code>fmul8sux16 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FMUL8ULX16	0 0011 0111	Unsigned lower 8-bit by signed 16-bit partitioned product	f64	f64	f64	<code>fmul8ulx16 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FMULD8SUX16	0 0011 1000	Signed upper 8-bit by signed 16-bit partitioned product	f32	f32	f64	<code>fmuld8sux16 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FMULD8ULX16	0 0011 1001	Unsigned lower 8-bit by signed 16-bit partitioned product	f32	f32	f64	<code>fmuld8ulx16 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1



Programming Note | When software emulates an 8-bit unsigned by 16-bit signed multiply, the unsigned value must be zero-extended and the 16-bit value sign-extended before the multiplication.

Description | The following sections describe the versions of partitioned multiplies.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an partitioned multiply instruction causes an *fp_disabled* exception.

Exceptions | *fp_disabled*

7.46.1 FMUL8x16 Instruction

FMUL8x16 multiplies each unsigned 8-bit value (for example, a pixel component) in the 32-bit floating-point register $F_S[rs1]$ by the corresponding (signed) 16-bit fixed-point integer in the 64-bit floating-point register $F_D[rs2]$. It rounds the 24-bit product (assuming binary point between bits 7 and 8) and stores the most significant 16 bits of the result into the corresponding 16-bit field in the 64-bit floating-point destination register $F_D[rd]$. FIGURE 7-11 illustrates the operation.

Note | This instruction treats the pixel component values as fixed-point with the binary point to the left of the most significant bit. Typically, this operation is used with filter coefficients as the fixed-point *rs2* value and image data as the *rs1* pixel value. Appropriate scaling of the coefficient allows various fixed-point scaling to be realized.

FMUL (partitioned)

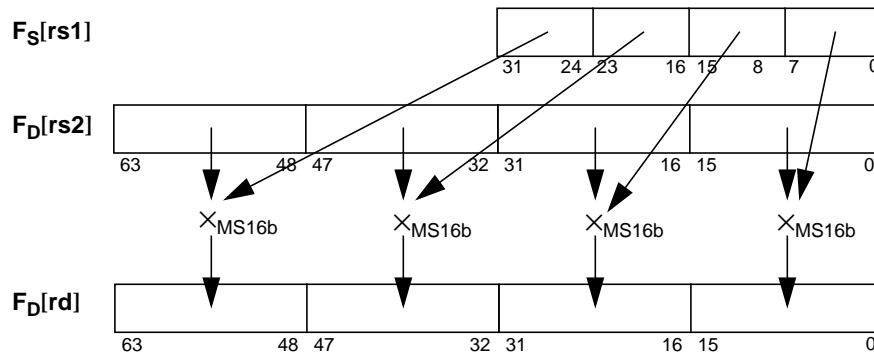


FIGURE 7-11 FMUL8x16 Operation

7.46.2 FMUL8x16AU Instruction

FMUL8x16AU is the same as FMUL8x16, except that one 16-bit fixed-point value is used as the multiplier for all four multiplies. This multiplier is the most significant (“upper”) 16 bits of the 32-bit register $F_S[rs2]$ (typically an α pixel component value). FIGURE 7-12 illustrates the operation.




FIGURE 7-12 FMUL8x16AU Operation

7.46.3 FMUL8x16AL Instruction

FMUL8x16AL is the same as FMUL8x16AU, except that the least significant (“lower”) 16 bits of the 32-bit register $F_S[rs2]$ register are used as a multiplier. FIGURE 7-13 illustrates the operation.




FIGURE 7-13 FMUL8x16AL Operation

FMUL (partitioned)

7.46.4 FMUL8SUx16 Instruction

FMUL8SUx16 multiplies the most significant (“upper”) 8 bits of each 16-bit signed value in the 64-bit floating-point register $F_D[rs1]$ by the corresponding signed, 16-bit, fixed-point, signed integer in the 64-bit floating-point register $F_D[rs2]$. It rounds the 24-bit product toward the nearest representable value and then stores the most significant 16 bits of the result into the corresponding 16-bit field of the 64-bit floating-point destination register $F_D[rd]$. If the product is exactly halfway between two integers, the result is rounded toward positive infinity. FIGURE 7-14 illustrates the operation.




FIGURE 7-14 FMUL8SUx16 Operation

7.46.5 FMUL8ULx16 Instruction

FMUL8ULx16 multiplies the unsigned least significant (“lower”) 8 bits of each 16-bit value in the 64-bit floating-point register $F_D[rs1]$ by the corresponding fixed-point signed 16-bit integer in the 64-bit floating-point register $F_D[rs2]$. Each 24-bit product is sign-extended to 32 bits. The most significant (“upper”) 16 bits of the sign-extended value are rounded to nearest and then stored in the corresponding 16-bit field of the 64-bit floating-point destination register $F_D[rd]$. If the result is exactly halfway between two integers, the result is rounded toward positive infinity. FIGURE 7-15 illustrates the operation; CODE EXAMPLE 7-1 exemplifies the operation.

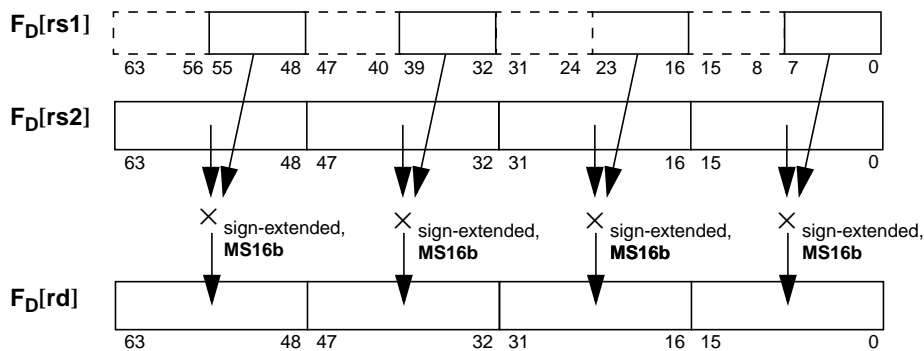


FIGURE 7-15 FMUL8ULx16 Operation

CODE EXAMPLE 7-1 16-bit × 16-bit 16-bit Multiply

<code>fmul8sux16</code>	<code>%f0, %f2, %f4</code>
<code>fmul8ulx16</code>	<code>%f0, %f2, %f6</code>
<code>fpadd16</code>	<code>%f4, %f6, %f8</code>

FMUL (partitioned)

7.46.6 FMULD8SUx16 Instruction

FMULD8SUx16 multiplies the most significant (“upper”) 8 bits of each 16-bit signed value in F[rs1] by the corresponding signed 16-bit fixed-point value in F[rs2]. Each 24-bit product is shifted left by 8 bits to generate a 32-bit result, which is then stored in the 64-bit floating-point register specified by rd. FIGURE 7-16 illustrates the operation.




FIGURE 7-16 FMULD8SUx16 Operation

7.46.7 FMULD8ULx16 Instruction

FMULD8ULx16 multiplies the unsigned least significant (“lower”) 8 bits of each 16-bit value in F[rs1] by the corresponding 16-bit fixed-point signed integer in F[rs2]. Each 24-bit product is sign-extended to 32 bits and stored in the corresponding half of the 64-bit floating-point register specified by rd. FIGURE 7-17 illustrates the operation; CODE EXAMPLE 7-2 exemplifies the operation.




FIGURE 7-17 FMULD8ULx16 Operation


CODE EXAMPLE 7-2 16-bit by 16-bit 32-bit Multiply

```

fmuld8sux16  %f0, %f2, %f4
fmuld8ulx16  %f0, %f2, %f6
fpadd32      %f4, %f6, %f8
    
```

7.47 Floating-Point Multiply

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FMULs	11 0100	0 0100 1001	Multiply Single	fmuls <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	A1
FMULd	11 0100	0 0100 1010	Multiply Double	fmuld <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	A1
FMULq	11 0100	0 0100 1011	Multiply Quad	fmulq <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C3
FsMULd	11 0100	0 0110 1001	Multiply Single to Double	fsmuld <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	A1
FdMULq	11 0100	0 0110 1110	Multiply Double to Quad	fdmulq <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C3



Description The floating-point multiply instructions multiply the contents of the floating-point register(s) specified by the rs1 field by the contents of the floating-point register(s) specified by the rs2 field. The instructions then write the product into the floating-point register(s) specified by the rd field.

The FsMULd instruction provides the exact double-precision product of two single-precision operands, without underflow, overflow, or rounding error. Similarly, FdMULq provides the exact quad-precision product of two double-precision operands.

Rounding is performed as specified by FSR.rd.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FMULq or FdMULq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute any FMUL instruction causes an *fp_disabled* exception.

An attempt to execute an FMULq instruction when rs1{1} ≠ 0 or rs2{1} ≠ 0 or rd{1:0} ≠ 0 causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.

An attempt to execute an FdMULq instruction when rd{1} ≠ 0 causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.


For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (FSR.ftt = invalid_fp_register (FMULq and FdMULq only))
fp_exception_other (FSR.ftt = unfinished_FPop)
fp_exception_ieee_754 (any: NV; FMUL<s|d|q> only: OF, UF, NX)

■ **See Also** FMAf on page 175

7.48 Floating-Point Negative Add vis 3

Instruction	opf	Operation	Assembly Language Syntax	Class
FNADDs	0 0101 0001	Negative Add, Single	<code>fnadds <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1
FNADDd	0 0101 0010	Negative Add, Double	<code>fnaddd <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1



Instruction	Implementation
FNADDs	$F_S[rd] \leftarrow -(F_S[rs1] + F_S[rs2])$
FNADDd	$F_D[rd] \leftarrow -(F_D[rs1] + F_D[rs2])$

Description The floating-point negative add instructions, FNADD<s|d>, add the floating-point register(s) specified by rs1 and the floating-point register(s) specified by rs2, negate the sum, and write the result into the floating-point register(s) specified by rd.

No rounding is performed between the addition operation and the subsequent negation; at most one rounding step occurs.

The FNADD instructions treat NaN values as follows:

- A source QNaN propagates with its sign bit unchanged
- A generated (default response) QNaN result has a sign bit of zero
- A source SNaN that is converted to a QNaN result retains the sign bit of the source SNaN

Exceptions. If an FNADD instruction is not implemented in hardware, it generates an *illegal_instruction* exception, so that privileged software can emulate the instruction.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FNADD instruction causes an *fp_disabled* exception.

Overflow, underflow, and inexact exception bits within `FSR.cexc` and `FSR.aexc` are updated based on the final result of the operation and not on the intermediate result of the addition. The invalid operation exception bits within `FSR.cexc` and `FSR.aexc` are updated as if the addition and negation were performed using two individual instructions. An invalid operation exception is detected when any of the following conditions are true:

- A source operand (`F[rs1]` or `F[rs2]`) is a SNaN
- $\infty - \infty$

If the instruction generates an IEEE-754 exception or exceptions for which the corresponding trap enable mask (`FSR.tem`) bits are set, an *fp_exception_ieee_754* exception and subsequent trap is generated.

If the addition operation detects an `unfinished_FPop` condition (for example, due to a subnormal operand or intermediate result), the Negative Add instruction generates an *fp_exception_other* exception with `FSR.ftt = unfinished_FPop`. The `unfinished_FPop` trap takes priority over any potential IEEE-754 exception (regardless whether traps are enabled, via `FSR.tem`, for the IEEE exception) and the `FSR.cexc` and `FSR.aexc` fields remain unchanged.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

FNADD

Semantic Definitions

FNADD:

- (1) $\text{tmp} \leftarrow F[\text{rs1}] + F[\text{rs2}]$
- (2) $\text{tmp} \leftarrow -\text{tmp}$
- (3) $F[\text{rd}] \leftarrow \mathbf{round}(\text{tmp})$

Exceptions

fp_disabled

fp_exception_ieee_754 (OF, UF, NX, NV)

fp_exception_other (FSR.ftt = unfinished_FPop)

See Also


FMAf on page 175

FADD on page 153 FNEG on page 193

FNMUL on page 175

7.49 Floating-Point Negate

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FNEGs	11 0100	0 0000 0101	Negate Single	<code>fnegs <i>freq_{rs2}, freq_{rd}</i></code>	A1
FNEGd	11 0100	0 0000 0110	Negate Double	<code>fnegd <i>freq_{rs2}, freq_{rd}</i></code>	A1
FNEGq	11 0100	0 0000 0111	Negate Quad	<code>fnegq <i>freq_{rs2}, freq_{rd}</i></code>	C3



Description FNEG copies the source floating-point register(s) to the destination floating-point register(s), with the sign bit complemented.

These instructions clear (set to 0) both `FSR.cexc` and `FSR.ftt`. They do not round, do not modify `FSR.aexc`, and do not treat floating-point NaN values differently from other floating-point values.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FNEGq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FNEG instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FNEG instruction causes an *fp_disabled* exception.


An attempt to execute an FNEGq instruction when `rs2{1} ≠ 0` or `rd{1} ≠ 0` causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (`FSR.ftt = invalid_fp_register` (FNEGq only))

FNHADD<s|d>

7.50 Floating-point Negative Add and Halve VIS 3

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FNHADDs	11 0100	0 0111 0001	Single-precision Add, Halve, and Negate	<code>fnhadds <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C2
FNHADDd	11 0100	0 0111 0010	Double-precision Add, Halve, and Negate	<code>fnhaddd <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C2




Description *The FNHADD<s|d> instructions are used to perform an addition of two floating-point values, halve the result (divide it by 2), and then negate it, all in a single operation. One benefit of this operation is that it cannot produce an arithmetic overflow.*

Exceptions *fp_disabled*
fp_exception_ieee_754 (UF, NX, NV)

See Also FHADD on page 166
 FHSUB on page 167

7.51 Floating-Point Negative Multiply VIS 3

Instruction	opf	Operation	Assembly Language Syntax	Class
FNMULs	0 0101 1001	Negative Multiply, Single	<code>fnmuls <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1
FNMULD	0 0101 1010	Negative Multiply, Double	<code>fnmuld <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1
FNsMULD	0 0111 1001	Negative Multiply, Single to Double	<code>fnsmuld <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1



Instruction	Implementation
FNMULs	$F_S[rd] \leftarrow - (F_S[rs1] \times F_S[rs2])$
FNMULD	$F_D[rd] \leftarrow - (F_D[rs1] \times F_D[rs2])$
FNsMULD	$F_D[rd] \leftarrow - (F_S[rs1] \times F_S[rs2])$

Description The floating-point negative multiply instructions, FNMUL<sd>, multiply the floating-point register(s) specified by rs1 and the floating-point register(s) specified by rs2, negate the product, and write the result into the floating-point register(s) specified by rd.

The floating-point negative multiply single to double instruction, FNsMULD, multiplies the single-precision floating point registers $F_S[rs1]$ and $F_S[rs2]$, generates a double-precision product, negates the product, and writes the result into the double-precision floating-point register $F_D[rd]$.

No rounding is performed between the multiplication operation and the subsequent negation; at most one rounding step occurs.

The FNMUL instructions treat NaN values as follows:

- A source QNaN propagates with its sign bit unchanged
- A generated (default response) QNaN result has a sign bit of zero
- A source SNaN that is converted to a QNaN result retains the sign bit of the source SNaN

Exceptions. If an FNMUL instruction is not implemented in hardware, it generates an *illegal_instruction* exception, so that privileged software can emulate the instruction.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FNMUL instruction causes an *fp_disabled* exception.

Overflow, underflow, and inexact exception bits within `FSR.cexc` and `FSR.aexc` are updated based on the final result of the operation and not on the intermediate result of the multiplication. The invalid operation exception bits within `FSR.cexc` and `FSR.aexc` are updated as if the multiplication and negation were performed using two individual instructions. An invalid operation exception is detected when any of the following conditions are true:

- A source operand ($F[rs1]$ or $F[rs2]$) is a SNaN
- $\infty \times 0$

If the instruction generates an IEEE-754 exception or exceptions for which the corresponding trap enable mask (`FSR.tem`) bits are set, an *fp_exception_ieee_754* exception and subsequent trap is generated.

If the multiplication operation detects an `unfinished_FPop` condition (for example, due to a subnormal operand or intermediate result), the Negative Multiply instruction generates an *fp_exception_other* exception with `FSR.ftt = unfinished_FPop`. The `unfinished_FPop` trap takes priority over any potential IEEE-754 exception (regardless whether traps are enabled, via `FSR.tem`, for the IEEE exception) and the `FSR.cexc` and `FSR.aexc` fields remain unchanged.

FNMUL

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Semantic Definitions:

FNMUL:

- (1) $\text{tmp} \leftarrow F[\text{rs1}] \times F[\text{rs2}]$
- (2) $\text{tmp} \leftarrow -\text{tmp}$
- (3) $F[\text{rd}] \leftarrow \mathbf{round}(\text{tmp})$


Exceptions *fp_disabled*
fp_exception_ieee_754 (OF, UF, NX, NV)
fp_exception_other (FSR.ftt = unfinished_FPop)

See Also FMAf on page 175
 FMUL on page 190
 FNADD on page 175
 FNEG on page 193

FPAK

7.52 FPAK VIS 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FPAK16	0 0011 1011	Four 16-bit packs into 8 unsigned bits	—	f64	f32	<code>fpack16 freg_{rs2}, freg_{rd}</code>	B1
FPAK32	0 0011 1010	Two 32-bit packs into 8 unsigned bits	f64	f64	f64	<code>fpack32 freg_{rs1}, freg_{rs2}, freg_{rd}</code>	B1
FPAKFIX	0 0011 1101	Four 16-bit packs into 16 signed bits	—	f64	f32	<code>fpackfix freg_{rs2}, freg_{rd}</code>	B1



Description The FPAK instructions convert multiple values in a source register to a lower-precision fixed or pixel format and stores the resulting values in the destination register. Input values are clipped to the dynamic range of the output format. Packing applies a scale factor from `GSR.scale` to allow flexible positioning of the binary point. See the subsections on following pages for more detailed descriptions of the operations of these instructions.

An attempt to execute an FPAK16 or FPAKFIX instruction when `rs1 ≠ 0` causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute any FPAK instruction causes an *fp_disabled* exception.

Exceptions *illegal_instruction* *fp_disabled*

See Also FEXPAND on page 165
 FPMERGE on page 216

FPAK

7.52.1 FPAK16

FPAK16 takes four 16-bit fixed values from the 64-bit floating-point register $F_D[rs2]$, scales, truncates, and clips them into four 8-bit unsigned integers, and stores the results in the 32-bit destination register, $F_S[rd]$. FIGURE 7-18 illustrates the FPAK16 operation.

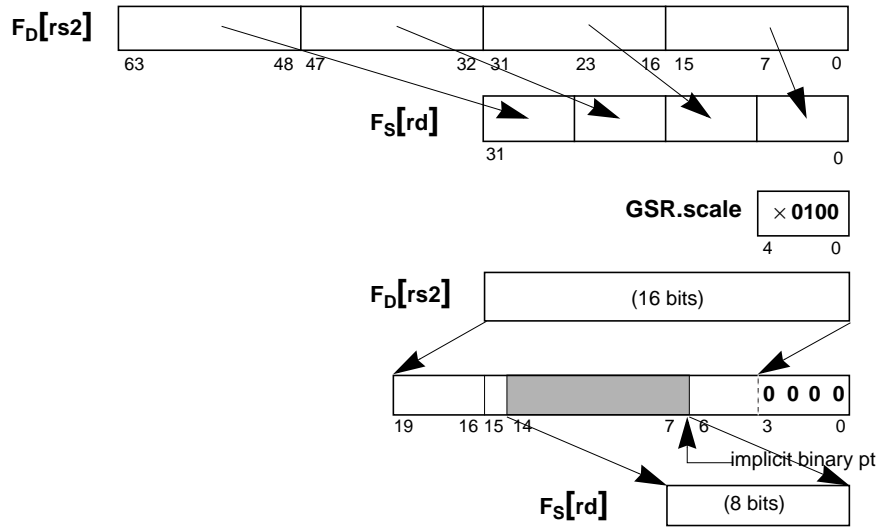


FIGURE 7-18 FPAK16 Operation

Note | FPAK16 ignores the most significant bit of $GSR.scale$ ($GSR.scale\{4\}$).

This operation is carried out as follows:

1. Left-shift the value from $F_D[rs2]$ by the number of bits specified in $GSR.scale$ while maintaining clipping information.
2. Truncate and clip to an 8-bit unsigned integer starting at the bit immediately to the left of the implicit binary point (that is, between bits 7 and 6 for each 16-bit word). Truncation converts the scaled value into a signed integer (that is, round toward negative infinity). If the resulting value is negative (that is, its most significant bit is set), 0 is returned as the clipped value. If the value is greater than 255, then 255 is delivered as the clipped value. Otherwise, the scaled value is returned as the result.
3. Store the result in the corresponding byte in the 32-bit destination register, $F_S[rd]$.

For each 16-bit partition, the sequence of operations performed is shown in the following example pseudo-code:

```
tmp ← source_operand{15:0} << GSR.scale;
// Pick off the bits from bit position 15+GSR.scale to
// bit position 7 from the shifted result
trunc_signed_value ← tmp{(15+GSR.scale):7};
If (trunc_signed_value < 0)
unsigned_8bit_result ← 0;
else if (trunc_signed_value > 255)
unsigned_8bit_result ← 255;
else
unsigned_8bit_result ← trunc_signed_value{14:7};
```

FPAK32

7.52.2 FPAK32

FPAK32 takes two 32-bit fixed values from the second source operand (64-bit floating-point register $F_D[rs2]$) and scales, truncates, and clips them into two 8-bit unsigned integers. The two 8-bit integers are merged at the corresponding least significant byte positions of each 32-bit word in the 64-bit floating-point register $F_D[rs1]$, left-shifted by 8 bits. The 64-bit result is stored in $F_D[rd]$. Thus, successive FPAK32 instructions can assemble two pixels by using three or four pairs of 32-bit fixed values. FIGURE 7-19 illustrates the FPAK32 operation.

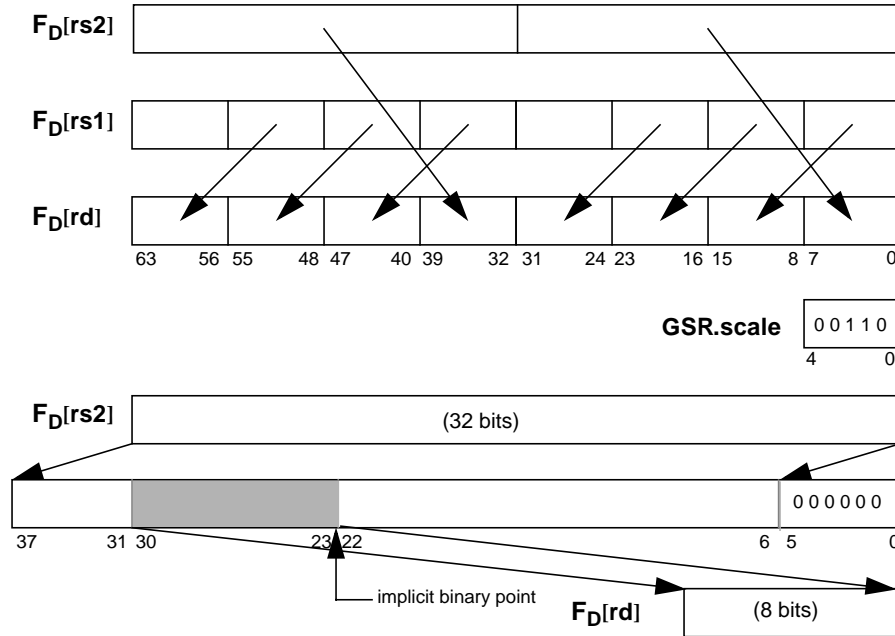


FIGURE 7-19 FPAK32 Operation

This operation, illustrated in FIGURE 7-19, is carried out as follows:

1. Left-shift each 32-bit value in $F_D[rs2]$ by the number of bits specified in $GSR.scale$, while maintaining clipping information.
2. For each 32-bit value, truncate and clip to an 8-bit unsigned integer starting at the bit immediately to the left of the implicit binary point (that is, between bits 23 and 22 for each 32-bit word). Truncation is performed to convert the scaled value into a signed integer (that is, round toward negative infinity). If the resulting value is negative (that is, the most significant bit is 1), then 0 is returned as the clipped value. If the value is greater than 255, then 255 is delivered as the clipped value. Otherwise, the scaled value is returned as the result.
3. Left-shift each 32-bit value from $F_D[rs1]$ by 8 bits.
4. Merge the two clipped 8-bit unsigned values into the corresponding least significant byte positions in the left-shifted $F_D[rs2]$ value.
5. Store the result in the 64-bit destination register $F_D[rd]$.

For each 32-bit partition, the sequence of operations performed is shown in the following pseudo-code:

```
tmp ← source_operand2{31:0} << GSR.scale;
// Pick off the bits from bit position 31+GSR.scale to
// bit position 23 from the shifted result
trunc_signed_value ← tmp{(31+GSR.scale):23};
if (trunc_signed_value < 0)
unsigned_8bit_value ← 0;
```

FPACK

```

else if (trunc_signed_value > 255)
unsigned_8bit_value ← 255;
else
unsigned_8bit_value ← trunc_signed_value{30:23};
Final_32bit_Result ← (source_operand1{31:0} << 8) |
(unsigned_8bit_value{7:0});

```

7.52.3 FPACKFIX

FPACKFIX takes two 32-bit fixed values from the 64-bit floating-point register $F_D[rs2]$, scales, truncates, and clips them into two 16-bit unsigned integers, and then stores the result in the 32-bit destination register $F_S[rd]$. FIGURE 7-20 illustrates the FPACKFIX operation.




FIGURE 7-20 FPACKFIX Operation

This operation is carried out as follows:

1. Left-shift each 32-bit value from $F_D[rs2]$ by the number of bits specified in $GSR.scale$, while maintaining clipping information.
2. For each 32-bit value, truncate and clip to a 16-bit unsigned integer starting at the bit immediately to the left of the implicit binary point (that is, between bits 16 and 15 for each 32-bit word). Truncation is performed to convert the scaled value into a signed integer (that is, round toward negative infinity). If the resulting value is less than -32768 , then -32768 is returned as the clipped value. If the value is greater than 32767 , then 32767 is delivered as the clipped value. Otherwise, the scaled value is returned as the result.
3. Store the result in the 32-bit destination register $F_S[rd]$.

For each 32-bit partition, the sequence of operations performed is shown in the following pseudo-code:

```


tmp ← source_operand{31:0} << GSR.scale;
// Pick off the bits from bit position 31+GSR.scale to
// bit position 16 from the shifted result
trunc_signed_value ← tmp{(31+GSR.scale):16};
if (trunc_signed_value < -32768) then signed_16bit_result ← -32768;
else if (trunc_signed_value > 32767) then signed_16bit_result ← 32767;
else signed_16bit_result ← trunc_signed_value{31:16};

```


7.53 Partitioned Add

The FPADD8 and FPADD64 instructions are new and not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPADD8	1 0010 0100	Eight 8-bit adds VIS 4	f64	f64	f64	<code>f_padd8 <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1	OSA 2015
FPADD16	0 0101 0000	Four 16-bit adds	f64	f64	f64	<code>f_padd16 <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1	VIS 1
FPADD32	0 0101 0010	Two 32-bit adds	f64	f64	f64	<code>f_padd32 <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1	VIS 1
FPADD64	0 0100 0010	64-bit integer add VIS 3B	f64	f64	f64	<code>f_padd64 <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C1	OSA 2011
FPADD16s	0 0101 0001	Two 16-bit adds	f32	f32	f32	<code>f_padd16s <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1	VIS 1
FPADD32s	0 0101 0011	One 32-bit add	f32	f32	f32	<code>f_padd32s <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1	VIS 1



Description

The 64-bit versions of these instructions perform partitioned integer addition of eight 8-bit (FPADD8), four 16-bit (FPADD16), two 32-bit (FPADD32) values, or one 64-bit value (FPADD64) value in $F_D[rs1]$ to corresponding value(s) in $F_D[rs2]$. The result values are written to the destination register, $F_D[rd]$.

The 32-bit versions of these instructions perform two 16-bit partitioned additions (FPADD16s) or one 32-bit (FPADD32s) partitioned addition, writing the result(s) to $F_S[rd]$.

Any carry out from each addition is discarded and a 2's-complement arithmetic result is produced.




FIGURE 7-21 FPADD8 Operation

FPADD




FIGURE 7-22 FPADD16 Operation




FIGURE 7-23 FPADD32 Operation




FIGURE 7-24 FPADD64 Operation

FPADD




FIGURE 7-25 FPADD16s Operation




FIGURE 7-26 FPADD32s Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FPADD instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*


See Also *Partitioned Add with Saturation* on page 204
 or *pstate.pef = 0 Partitioned Subtract* on page 219

FPADDS, FPADDUS

7.54 Partitioned Add with Saturation

The FPADDS8, FPADDUS8, and FPADDUS16 instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPADDS8 ^N	1 0010 0110	Eight 8-bit adds with saturation VIS 4	f64	f64	f64	<code>fpadds8 <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2015
FPADDUS8 ^N	1 0010 0111	Eight 8-bit unsigned adds with saturation VIS 4	f64	f64	f64	<code>fpadding8 <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2015
FPADDS16 ^N	0 0101 1000	Four 16-bit adds with saturation	f64	f64	f64	<code>fpadding16 <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2011
FPADDUS16 ^N	1 0010 0011	Four 16-bit unsigned adds with saturation VIS 4	f64	f64	f64	<code>fpadding16 <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2015
FPADDS32 ^N	0 0101 1010	Two 32-bit adds with saturation	f64	f64	f64	<code>fpadding32 <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2011
FPADDS16s ^N	0 0101 1001	Two 16-bit adds with saturation	f32	f32	f32	<code>fpadding16s <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2011
FPADDS32s ^N	0 0101 1011	One 32-bit add with saturation	f32	f32	f32	<code>fpadding32s <i>reg_{rs1}</i>, <i>reg_{rs2}</i>, <i>reg_{rd}</i></code>	C1	OSA 2011



Description The 64-bit versions of these instructions perform partitioned integer addition with saturation of eight 8-bit (FPADDS8, FPADDUS8), four 16-bit (FPADDS16, FPADDUS16), or two 32-bit (FPADDS32) values in $F_D[rs1]$ to corresponding values in $F_D[rs2]$. The result values are written to the destination register, $F_D[rd]$.

The 32-bit versions of these instructions perform partitioned integer addition of two 16-bit values (FPADDS16s) or one 32-bit value (FPADDS32s) in $F_S[rs1]$ to corresponding value(s) in $F_S[rs2]$. The result value(s) are written to $F_S[rd]$.

These instructions clip (saturate) overflow results, as indicated in TABLE 7-11.

TABLE 7-11 Clipping Values for FPADDS instructions

Element (Partition) Size	Signed Add with Saturation		Unsigned Add with Saturation	
	Negative Overflow Clipped in Result to	Positive Overflow Clipped in Result to	Negative Overflow Clipped in Result to	Positive Overflow Clipped in Result to
8 bits	-2^7 (-128)	2^7-1 (127)	(cannot occur)	2^8-1 (255)
16 bits	-2^{15} (-32,768)	$2^{15}-1$ (32,767)	(cannot occur)	$2^{16}-1$ (65,535)
32 bits	-2^{31}	$2^{31}-1$	—	—

FPADDS, FPADDUS




FIGURE 7-27 FPADDS8, FPADDUS8 Operation




FIGURE 7-28 FPADDS16, FPADDUS16 Operation




FIGURE 7-29 FPADDS32 Operation

FPADDS, FPADDUS




FIGURE 7-30 FPADDS16s Operation




FIGURE 7-31 FPADDS32s Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FPADDS instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

See Also *Partitioned Add* on page 201
 Partitioned Subtract with Saturation on page 222


7.55 Partitioned Signed Compare }

The FCMPLE8 and FCMPGT8 instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPCMPLE8	0 0011 0100	Eight 8-bit compares; set bit in R[rd] if $src1 \leq src2$ VIS 4	f64	f64	i64	<code>fpcmp1e8</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPCMPGT8	0 0011 1100	Eight 8-bit compares; set bit in R[rd] if $src1 > src2$ VIS 4	f64	f64	i64	<code>fpcmpgt8</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPCMPNE8	1 0010 0010	Eight 8-bit compares; set bit in R[rd] if $src1 \neq src2$	f64	f64	i64	<code>fpcmpne8</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fpcmpune8 ‡</code>)	C1	OSA 2011
FPCMPEQ8	1 0010 1010	Eight 8-bit compares; set bit in R[rd] if $src1 = src2$	f64	f64	i64	<code>fpcmp8eq8</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fpcmpueq8 ‡</code>)	C1	OSA 2011
FPCMPEQ16	0 0010 1010	Four 16-bit compares; set bit in R[rd] if $src1 = src2$	f64	f64	i64	<code>fpcmp16eq16</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fpcmpueq16 ‡</code>) (or <code>fcmpeq16 †</code>)	B1	VIS 1
FPCMPNE16	0 0010 0010	Four 16-bit compares; set bit in R[rd] if $src1 \neq src2$	f64	f64	i64	<code>fpcmp16ne16</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fpcmpune16 ‡</code>) (or <code>fcmpne16 †</code>)	B1	VIS 1
FPCMPLE16	0 0010 0000	Four 16-bit compares; set bit in R[rd] if $src1 \leq src2$	f64	f64	i64	<code>fpcmp16le16</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fcmple16 †</code>)	B1	VIS 1
FPCMPGT16	0 0010 1000	Four 16-bit compares; set bit in R[rd] if $src1 > src2$	f64	f64	i64	<code>fpcmp16gt16</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fcmpgt16 †</code>)	B1	VIS 1
FPCMPEQ32	0 0010 1110	Two 32-bit compares; set bit in R[rd] if $src1 = src2$	f64	f64	i64	<code>fpcmp32eq32</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fpcmpueq32 ‡</code>) (or <code>fcmpeq32 †</code>)	B1	VIS 1
FPCMPNE32	0 0010 0110	Two 32-bit compares; set bit in R[rd] if $src1 \neq src2$	f64	f64	i64	<code>fpcmp32ne32</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fpcmpune32 ‡</code>) (or <code>fcmpne32 †</code>)	B1	VIS 1
FPCMPLE32	0 0010 0100	Two 32-bit compares; set bit in R[rd] if $src1 \leq src2$	f64	f64	i64	<code>fpcmp32le32</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fcmple32 †</code>)	B1	VIS 1
FPCMPGT32	0 0010 1100	Two 32-bit compares; set bit in R[rd] if $src1 > src2$	f64	f64	i64	<code>fpcmp32gt32</code> <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or <code>fcmpgt32 †</code>)	B1	VIS 1

† the older, deprecated mnemonic for this instruction (still recognized by the assembler)

‡ a synonymous mnemonic, since signed and unsigned comparisons for equality (or inequality) are identical



Description Eight 8-bit, four 16-bit, or two 32-bit signed integer values in $F_D[rs1]$ and $F_D[rs2]$ are compared. The 8, 4 or 2 bits (respectively) of comparison results are stored in the least significant bits of integer register R[rd]. The least significant 8-bit, 16-bit or 32-bit comparison result corresponds to bit 0 of R[rd].

Note Bits not written with comparison results in destination integer register R[rd] are set to zero.

FPCMP

For FPCMPGT{8,16,32}, each bit in the result is set to 1 if the corresponding signed integer value in $F_D[rs1]$ is greater than the signed value in $F_D[rs2]$.

Programming Note | Less-than comparisons are made by swapping the source operands of FPCMPGT.

For FPCMPLE{8,16,32}, each bit in the result is set to 1 if the corresponding signed integer value in $F_D[rs1]$ is less than or equal to the signed value in $F_D[rs2]$.

Programming Note | Greater-than-or-equal comparisons are made by swapping the source operands of FPCMPLE.

For FPCMPPEQ16 and FPCMPPEQ32, each bit in the result is set to 1 if the corresponding signed integer value in $F_D[rs1]$ is equal to the signed value in $F_D[rs2]$.

For FPCMPNE16 and FPCMPNE32, each bit in the result is set to 1 if the corresponding signed integer value in $F_D[rs1]$ is not equal to the signed value in $F_D[rs2]$.

FIGURE 7-33, FIGURE 7-33 and FIGURE 7-34 illustrate 8-bit, 16-bit and 32-bit pixel comparison operations, respectively.




FIGURE 7-32 Eight 8-bit Signed Integer Partitioned Comparison Operations




FIGURE 7-33 Four 16-bit Signed Integer Partitioned Comparison Operations

FPCMP




FIGURE 7-34 Two 32-bit Signed Integer Partitioned Comparison Operation

In all comparisons, if a compare condition is not true, the corresponding bit in the result is set to 0.

Programming Note | The results of a Partitioned signed compare operation can be used directly by integer operations (for example, partial stores) and by the CMASK instruction.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute a SIMD signed compare instruction causes an *fp_disabled* exception.

Exception *fp_disabled*

See Also

- CMASK on page 139
- Floating-Point Compare on page 162
- FPCMPU on page 210
- STPARTIALF on page 347

7.56 Partitioned Unsigned Compare

The FPCMPU*16, and FPCMPU*32 instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPCMPUEQ8	1 0010 1010	Eight 8-bit compares VIS 3B ; set bit in R[rd] if $src1 = src2$	f64	f64	i64	fpcmpueq8 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or fpcmpeq8 ‡)	C1	OSA 2011
FPCMPUNE8	1 0010 0010	Eight 8-bit compares VIS 3B ; set bit in R[rd] if $src1 \neq src2$	f64	f64	i64	fpcmpune8 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or fpcmpne8 ‡)	C1	OSA 2011
FPCMPULE8	1 0010 0000	Eight 8-bit compares VIS 3B ; set bit in R[rd] if $src1 \leq src2$	f64	f64	i64	fpcmpule8 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2011
FPCMPUGT8	1 0010 1000	Eight 8-bit compares VIS 3B ; set bit in R[rd] if $src1 > src2$	f64	f64	i64	fpcmpugt8 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2011
FPCMPUEQ16	0 0010 1010	Four 16-bit compares; set bit in R[rd] if $src1 = src2$	f64	f64	i64	fpcmpueq16 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or fpcmpeq16 ‡)	B1	VIS 1
FPCMPUNE16	0 0010 0010	Four 16-bit compares; set bit in R[rd] if $src1 \neq src2$	f64	f64	i64	fpcmpune16 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or fpcmpne16 ‡)	B1	VIS 1
FPCMPULE16	1 0010 1110	Four 16-bit compares; set bit in R[rd] if $src1 \leq src2$ VIS 4	f64	f64	i64	fpcmpule16 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPCMPUGT16	1 0010 1011	Four 16-bit compares; set bit in R[rd] if $src1 > src2$ VIS 4	f64	f64	i64	fpcmpugt16 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPCMPUEQ32	0 0010 1110	Two 32-bit compares; set bit in R[rd] if $src1 = src2$	f64	f64	i64	fpcmpueq32 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or fpcmpeq32 ‡)	B1	VIS 1
FPCMPUNE32	0 0010 0110	Two 32-bit compares; set bit in R[rd] if $src1 \neq src2$	f64	f64	i64	fpcmpune32 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i> (or fpcmpne32 ‡)	B1	VIS 1
FPCMPULE32	1 0010 1111	Two 32-bit compares; set bit in R[rd] if $src1 \leq src2$ VIS 4	f64	f64	i64	fpcmpule32 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPCMPUGT32	1 0010 1100	Two 32-bit compares; set bit in R[rd] if $src1 > src2$ VIS 4	f64	f64	i64	fpcmpugt32 <i>freq_{rs1}, freq_{rs2}, reg_{rd}</i>	C1	OSA 2015

‡ a synonymous mnemonic, since signed and unsigned comparisons for equality (or inequality) are identical



Description Eight 8-bit, four 16-bit, or two 32-bit unsigned integer values in two 64-bit floating-point registers are compared. The 8 (or 4, or 2) bits, respectively, of comparison results are stored in the least significant bits of the integer register R[rd]. The least significant comparison result corresponds to bit zero of R[rd].

Note Bits not written with comparison results in destination integer register R[rd] are set to zero.

FPCMPU

For FPCMPUGT8 , FPCMPUGT16, and FPCMPUGT32, each bit in the result is set to 1 if the corresponding unsigned 8-bit, 16-bit, or 32-bit value in $F_D[rs1]$ is greater than the corresponding unsigned value in $F_D[rs2]$.

For FPCMPULE8, FPCMPULE16, and FPCMPULE32, each bit in the result is set to 1 if the corresponding unsigned 8-bit, 16-bit, or 32-bit value in $F_D[rs1]$ is less than or equal to the corresponding unsigned value in $F_D[rs2]$.

For FPCMPUNE8, each bit in the result is set to 1 if the corresponding unsigned 8-bit value in $F_D[rs1]$ is not equal to the corresponding unsigned value in $F_D[rs2]$.

For FPCMPUEQ8, each bit in the result is set to 1 if the corresponding unsigned 8-bit value in $F_D[rs1]$ is equal to the corresponding unsigned value in $F_D[rs2]$.

Programming Note | Less-than comparisons are made by swapping the source operands of FPCMPUGT*.

Programming Note | Greater-than-or-equal comparisons are made by swapping the source operands of FPCMPULE*.

Programming Note | The results of a SIMD unsigned compare operation can be used directly by integer operations (for example, partial stores) and by the CMASK instruction.

FIGURE 7-35 illustrates the 8-bit unsigned Partitioned comparison operations.




FIGURE 7-35 Eight 8-bit Unsigned Integer Partitioned Comparison Operations




FIGURE 7-36 Four 16-bit Unsigned Integer Partitioned Comparison Operations

FPCMPU




FIGURE 7-37 Two 32-bit Unsigned Integer Partitioned Comparison Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FPCMPU instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

See Also CMASK on page 139
 FPCMP* on page 208
 STPARTIALF on page 347

7.57 Integer Multiply-Add

The integer multiply-add instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	op5	Operation	Assembly Language Syntax	Class	Added
FPMADDX	00 00	Unsigned integer multiply-add for lower 64 bits	<code>fpmaddx <i>reg_{rs1}, reg_{rs2}, reg_{rs3}, reg_{rd}</i></code>	N1	IMA
FPMADDXHI	01 00	Unsigned integer multiply-add for upper 64 bits	<code>fpmaddxhi <i>reg_{rs1}, reg_{rs2}, reg_{rs3}, reg_{rd}</i></code>	N1	IMA



Description The Integer Multiply-Add instruction performs a fused multiply and add operation on unsigned 64-bit integer values residing in floating-point registers.

FPMADDX performs an unsigned integer multiplication of the 64-bit floating-point registers $F_D[rs1]$ and $F_D[rs2]$, adds the resulting product to the unsigned integer value from $F_D[rs3]$, then writes the *least* significant 8 bytes of the result into $F_D[rd]$.

FPMADDXHI performs an unsigned integer multiplication of the 64-bit floating-point registers $F_D[rs1]$ and $F_D[rs2]$, adds the resulting product to the unsigned integer value from $F_D[rs3]$, then writes the *most* significant 8 bytes of the result into $F_D[rd]$.

FPMADDX and FPMADDXHI do not modify any portion of FSR.

Exceptions `fp_disabled`

7.58 Partitioned Integer Maximum VIS 4

The Partitioned Maximum instructions are new in Oracle SPARC Architecture 2015 and are not guaranteed to be implemented on all implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPMAX8	1 0001 1101	Eight 8-bit maximums	f64	f64	f64	f _{pmax8} <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2015
FPMAX16	1 0001 1110	Four 16-bit maximums	f64	f64	f64	f _{pmax16} <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2015
FPMAX32	1 0001 1111	Two 32-bit maximums	f64	f64	f64	f _{pmax32} <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2015
FPMAXU8	1 0101 1101	Eight 8-bit unsigned maximums	f64	f64	f64	f _{pmaxu8} <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2015
FPMAXU16	1 0101 1110	Four 16-bit unsigned maximums	f64	f64	f64	f _{pmaxu16} <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2015
FPMAXU32	1 0101 1111	Two 32-bit unsigned maximums	f64	f64	f64	f _{pmaxu32} <i>freg_{rs1}</i> , <i>freg_{rs2}</i> , <i>freg_{rd}</i>	C1	OSA 2015



Description

FPMAX8, FPMAX16, and FPMAX32, respectively, compare eight 8-bit, four 16-bit, or two 32-bit corresponding partitioned signed integer values in the source operands ($F_D[rs1]$, $F_D[rs2]$). For each pair of corresponding fields, the numerically greater (maximum) of the two values is written to the corresponding field of the destination register, $F_D[rd]$.

FPMAXU8, FPMAXU16, and FPMAXU32 operate identically to FPMAX8, FPMAX16, and FPMAX32, except that the values being compared are treated as unsigned integers.




FIGURE 7-38 FPMAX8, FPMAXU8 Operation

FPMAX, FPMAXU




FIGURE 7-39 FPMAX16, FPMAXU16 Operation

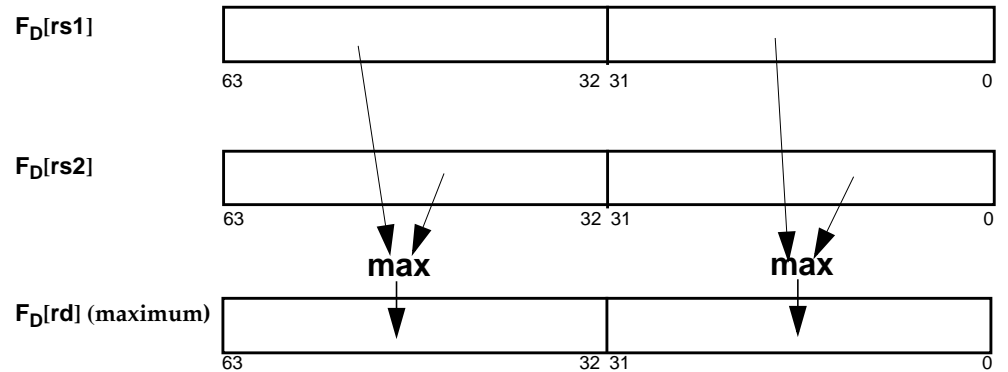


FIGURE 7-40 FPMAX32, FPMAXU32 Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an $FPMAX[U]$ instruction causes an *fp_disabled* exception.


Exceptions *fp_disabled*

■ *See Also* *Partitioned Integer Minimum* on page 217

FPMERGE

7.59 FPMERGE VIS 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FPMERGE	0 0100 1011	Two 32-bit merges	f32	f32	f64	<code>fpmerge <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	B1



Description FPMERGE interleaves eight 8-bit unsigned values in $F_S[rs1]$ and $F_S[rs2]$ to produce a 64-bit value in the destination register $F_D[rd]$. This instruction converts from packed to planar representation when it is applied twice in succession; for example, $R1G1B1A1, R3G3B3A3 \rightarrow R1R3G1G3A1A3 \rightarrow R1R2R3R4G1G2G3G4$.

FPMERGE also converts from planar to packed when it is applied twice in succession; for example, $R1R2R3R4, B1B2B3B4 \rightarrow R1B1R2B2R3B3R4B4 \rightarrow R1G1B1A1R2G2B2A2$.

FIGURE 7-41 illustrates the operation.




FIGURE 7-41 FPMERGE Operation

	<code>%d0</code>	<code>!R1</code>	<code>G1</code>	<code>B1</code>	<code>A1</code>	<code>R2</code>	<code>G2</code>	<code>B2</code>	<code>A2</code>	} <i>packed representation</i>
	<code>%d2</code>	<code>!R3</code>	<code>G3</code>	<code>B3</code>	<code>A3</code>	<code>R4</code>	<code>G4</code>	<code>B4</code>	<code>A4</code>	
<code>fpmerge %f0, %f2, %d4</code>	<code>!R1</code>	<code>R3</code>	<code>G1</code>	<code>G3</code>	<code>B1</code>	<code>B3</code>	<code>A1</code>	<code>A3</code>	} <i>intermediate</i>	
<code>fpmerge %f1, %f3, %d6</code>	<code>!R2</code>	<code>R4</code>	<code>G2</code>	<code>G4</code>	<code>B2</code>	<code>B4</code>	<code>A2</code>	<code>A4</code>		
<code>fpmerge %f4, %f6, %d0</code>	<code>!R1</code>	<code>R2</code>	<code>R3</code>	<code>R4</code>	<code>G1</code>	<code>G2</code>	<code>G3</code>	<code>G4</code>	} <i>planar representation</i>	
<code>fpmerge %f5, %f7, %d2</code>	<code>!B1</code>	<code>B2</code>	<code>B3</code>	<code>B4</code>	<code>A1</code>	<code>A2</code>	<code>A3</code>	<code>A4</code>		
<code>fpmerge %f0, %f2, %d4</code>	<code>!R1</code>	<code>B1</code>	<code>R2</code>	<code>B2</code>	<code>R3</code>	<code>B3</code>	<code>R4</code>	<code>B4</code>	} <i>intermediate</i>	
<code>fpmerge %f1, %f3, %d6</code>	<code>!G1</code>	<code>A1</code>	<code>G2</code>	<code>A2</code>	<code>G3</code>	<code>A3</code>	<code>G4</code>	<code>A4</code>		
<code>fpmerge %f4, %f6, %d0</code>	<code>!R1</code>	<code>G1</code>	<code>B1</code>	<code>A1</code>	<code>R2</code>	<code>G2</code>	<code>B2</code>	<code>A2</code>	} <i>packed representation</i>	
<code>fpmerge %f5, %f7, %d2</code>	<code>!R3</code>	<code>G3</code>	<code>B3</code>	<code>A3</code>	<code>R4</code>	<code>G4</code>	<code>B4</code>	<code>A4</code>		

CODE EXAMPLE 7-3 FPMERGE example

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FPMERGE instruction causes an *fp_disabled* exception.

Exceptions `fp_disabled`


See Also FPACK on page 197
FEXPAND on page 165

FPMIN, FPMINU

7.60 Partitioned Integer Minimum VIS 4

The Partitioned Minimum instructions are new in Oracle SPARC Architecture 2015 and are not guaranteed to be implemented on all implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPMIN8	1 0001 1010	Eight 8-bit minimums	f64	f64	f64	fpmin8 <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C1	OSA 2015
FPMIN16	1 0001 1011	Four 16-bit minimums	f64	f64	f64	fpmin16 <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C1	OSA 2015
FPMIN32	1 0001 1100	Two 32-bit minimums	f64	f64	f64	fpmin32 <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C1	OSA 2015
FPMINU8	1 0101 1010	Eight 8-bit unsigned minimums	f64	f64	f64	fpminu8 <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C1	OSA 2015
FPMINU16	1 0101 1011	Four 16-bit unsigned minimums	f64	f64	f64	fpminu16 <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C1	OSA 2015
FPMINU32	1 0101 1100	Two 32-bit unsigned minimums	f64	f64	f64	fpminu32 <i>freq_{rs1}, freq_{rs2}, freq_{rd}</i>	C1	OSA 2015



Description

FPMIN8, FPMIN16, and FPMIN32, respectively, compare eight 8-bit, four 16-bit, or two 32-bit corresponding partitioned signed integer values in the source operands ($F_D[rs1]$, $F_D[rs2]$). For each pair of corresponding fields, the numerically lesser (minimum) of the two values is written to the corresponding field of the destination register, $F_D[rd]$.

FPMINU8, FPMINU16, and FPMINU32 operate identically to FPMIN8, FPMIN16, and FPMIN32, except that the values being compared are treated as unsigned integers.

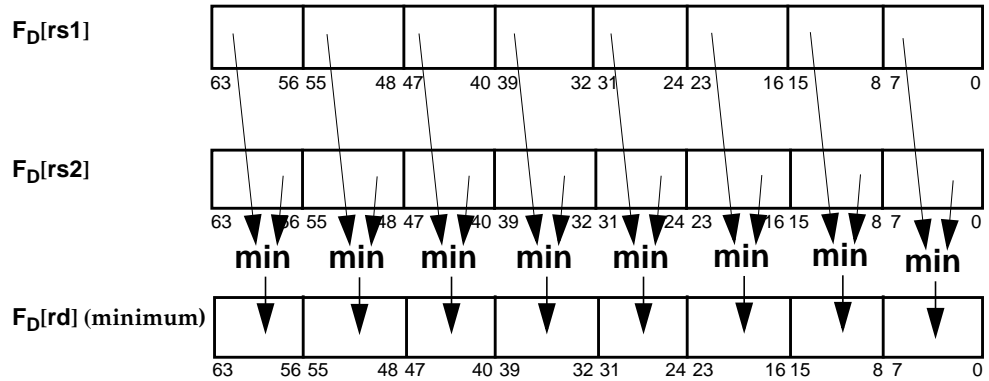


FIGURE 7-42 FPMIN8, FPMINU8 Operation

FPMIN, FPMINU




FIGURE 7-43 FPMIN16, FPMINU16 Operation




FIGURE 7-44 FPMIN32, FPMINU32 Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FP[U]MIN instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

■ *See Also* *Partitioned Integer Maximum* on page 214


FPSUB

7.61

or PSTATE.pef = 0 Partitioned Subtract

The FPSUB8 and FPSUB64 instructions are new and not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax		Class	Added
FPSUB8 ^N	1 0101 0100	Eight 8-bit subtracts VIS 4	f64	f64	f64	f _p sub8	<i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	C1	OSA 2015
FPSUB16	0 0101 0100	Four 16-bit subtracts	f64	f64	f64	f _p sub16	<i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1	VIS 1
FPSUB32	0 0101 0110	Two 32-bit subtracts	f64	f64	f64	f _p sub32	<i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1	VIS 1
FPSUB64	0 0100 0110	64-bit integer subtract VIS 3B	f64	f64	f64	f _p sub64	<i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	C1	OSA 2011
FPSUB16s	0 0101 0101	Two 16-bit subtracts	f32	f32	f32	f _p sub16s	<i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1	VIS 1
FPSUB32s	0 0101 0111	One 32-bit subtract	f32	f32	f32	f _p sub32s	<i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1	VIS 1



Description

The 64-bit versions of these instructions perform partitioned integer subtraction of eight 8-bit (FPSUB8), four 16-bit (FPSUB16), two 32-bit (FPSUB32) values, or one 64-bit value (FPSUB64) value in $F_D[rs2]$ from corresponding value(s) in $F_D[rs1]$. The result value(s) are written to the destination register, $F_D[rd]$.

The 32-bit versions of these instructions perform partitioned integer subtraction of two 16-bit values (FPSUB16s) or one 32-bit value (FPSUB32s) in $F_S[rs2]$ from the corresponding value(s) in $F_S[rs1]$. The result(s) are written to $F_S[rd]$.

Any borrow out from each subtraction is discarded and a 2's-complement arithmetic result is produced.




FIGURE 7-45 FPSUB8 Operation

FPSUB




FIGURE 7-46 FPSUB16 Operation




FIGURE 7-47 FPSUB32 Operation




FIGURE 7-48 FPSUB64 Operation

FPSUB




FIGURE 7-49 FPSUB16s Operation




FIGURE 7-50 FPSUB32s Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FPSUB instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*


See Also *Partitioned Add* on page 201
 Partitioned Subtract with Saturation on page 222

FPSUBS, FPSUBUS

7.62 Partitioned Subtract with Saturation

The FPSUBS8, FPSUBUS8, and FPSUBUS16 instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
FPSUBS8	1 0101 0110	Eight 8-bit subtracts with saturation VIS 4	f64	f64	f64	fpsubs8 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPSUBUS8	1 0101 0111	Eight 8-bit unsigned subtracts with saturation VIS 4	f64	f64	f64	fpsubus8 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPSUBS16	0 0101 1100	Four 16-bit subtracts with saturation	f64	f64	f64	fpsubs16 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2011
FPSUBUS16	1 0101 0011	Four 16-bit unsigned subtracts with saturation VIS 4	f64	f64	f64	fpsubus16 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2015
FPSUBS32	0 0101 1110	Two 32-bit subtracts with saturation	f64	f64	f64	fpsubs32 <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2011
FPSUBS16s	0 0101 1101	Two 16-bit subtracts with saturation	f32	f32	f32	fpsubs16s <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2011
FPSUBS32s	0 0101 1111	One 32-bit subtract with saturation	f32	f32	f32	fpsubs32s <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2011



Description

The 64-bit versions of these instructions perform partitioned integer subtraction with saturation of eight 8-bit (FPSUBS8, FPSUBUS8), four 16-bit (FPSUBS16, FPSUBUS16), or two 32-bit (FPSUBS32) values in $F_D[rs2]$ from the corresponding values in $F_D[rs1]$. The result values are written to the destination register, $F_D[rd]$.

The 32-bit versions of these instructions (FPSUBS16s and FPSUBS32s) perform partitioned integer subtraction of two 16-bit or one 32-bit value in $F_S[rs2]$ from corresponding values in $F_S[rs1]$. The result values are written to $F_S[rd]$.

These instructions clip (saturate) overflow results, as indicated in TABLE 7-12.

TABLE 7-12 Clipping Values for FPSUBS instructions

Element (Partition) Size	Signed Subtract with Saturation		Unsigned Subtract with Saturation	
	Negative Overflow Clipped in Result to	Postive Overflow Clipped in Result to	Negative Overflow Clipped in Result to	Postive Overflow Clipped in Result to
8 bits	-2^7 (-128)	2^7-1 (127)	0	(cannot occur)
16 bits	-2^{15} (-32,768)	$2^{15}-1$ (32,767)	0	(cannot occur)
32 bits	-2^{31}	$2^{31}-1$	—	—

FPSUBS, FPSUBUS




FIGURE 7-51 FPSUBS8 Operation




FIGURE 7-52 FPSUBS16 Operation




FIGURE 7-53 FPSUBS32 Operation

FPSUBS, FPSUBUS




FIGURE 7-54 FPSUBS16s Operation




FIGURE 7-55 FPSUBS32s Operation

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FPSUBS instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

See Also *Partitioned Add with Saturation* on page 204
 or *pstate.pef = 0 Partitioned Subtract* on page 219

F Register 1-operand Logical Ops

7.63 F Register Logical Operate (1 operand) VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
FZEROd	0 0110 0000	Zero fill	fzerod† <i>freg_{rd}</i>	A1
FZEROs	0 0110 0001	Zero fill, 32-bit	fzeros <i>freg_{rd}</i>	A1
FONEd	0 0111 1110	One fill	foned† <i>freg_{rd}</i>	A1
FONEs	0 0111 1111	One fill, 32-bit	fones <i>freg_{rd}</i>	A1

† this assembly-language instruction mnemonic is also accepted without a trailing “d”, for backward compatibility



Description FZEROd and FONEd fill the 64-bit destination register, $F_D[rd]$, with all ‘0’ bits or all ‘1’ bits (respectively).

FZEROs and FONEs fill the 32-bit destination register, $F_D[rd]$, with all ‘0’ bits or all ‘1’ bits (respectively).

An attempt to execute an FZERO or FONE instruction when instruction bits 18:14 or bits 4:0 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an FZERO or FONE instruction causes an *fp_disabled* exception.

Exceptions *illegal_instruction*
fp_disabled


See Also F Register 2-operand Logical Operations on page 226
F Register 3-operand Logical Operations on page 227

F Register 2-operand Logical Ops

7.64 F Register Logical Operate (2 operand) VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
FSRC1d	0 0111 0100	Copy $F_D[rs1]$ to $F_D[rd]$	<code>fsrc1d†</code> <i>freq_{rs1}, freq_{rd}</i>	A1
FSRC1s	0 0111 0101	Copy $F_S[rs1]$ to $F_S[rd]$, 32-bit	<code>fsrc1s</code> <i>freq_{rs1}, freq_{rd}</i>	A1
FSRC2d	0 0111 1000	Copy $F_D[rs2]$ to $F_D[rd]$	<code>fsrc2d†</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FSRC2s	0 0111 1001	Copy $F_S[rs2]$ to $F_S[rd]$, 32-bit	<code>fsrc2s</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FNOT1d	0 0110 1010	Negate (1's complement) $F_D[rs1]$	<code>fnot1d†</code> <i>freq_{rs1}, freq_{rd}</i>	A1
FNOT1s	0 0110 1011	Negate (1's complement) $F_S[rs1]$, 32-bit	<code>fnot1s</code> <i>freq_{rs1}, freq_{rd}</i>	A1
FNOT2d	0 0110 0110	Negate (1's complement) $F_D[rs2]$	<code>fnot2d†</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FNOT2s	0 0110 0111	Negate (1's complement) $F_S[rs2]$, 32-bit	<code>fnot2s</code> <i>freq_{rs2}, freq_{rd}</i>	A1

† this assembly-language instruction mnemonic is also accepted without a trailing “d”, for backward compatibility



Description The standard 64-bit versions of these instructions perform one of four 64-bit logical operations on the data from the 64-bit floating-point source register $F_D[rs1]$ (or $F_D[rs2]$) and store the result in the 64-bit floating-point destination register $F_D[rd]$.

The 32-bit (single-precision) versions of these instructions perform 32-bit logical operations on $F_S[rs1]$ (or $F_S[rs2]$) and store the result in $F_S[rd]$.

An attempt to execute an FSRC1 or FNOT1 instruction when instruction bits 4:0 are nonzero causes an *illegal_instruction* exception. An attempt to execute an FSRC2(s) or FNOT2(s) instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FSRC1<s|d>, FNOT1<s|d>, FSRC1<s|d>, or FNOT1<s|d> instruction causes an *fp_disabled* exception.

Programming Note FSRC1s (FSRC1d) and FSRC2s (FSRC2d) function similarly to FMOVs (FMOVd), except that the FSRC* instructions do not modify the FSR register while FMOVs (FMOVd) update some fields of FSR (see *Floating-Point Move* on page 179).
If a 64-bit floating-point register-to-register copy is desired and simultaneous modification of FSR is not required, use of FSRC2d is strongly recommended over FMOVd. FSRC2d is at least as fast as, and on many implementations much faster than, FMOVd and FSRC1d.

Exceptions *illegal_instruction*
fp_disabled


See Also *Floating-Point Move* on page 179
F Register 1-operand Logical Operations on page 225
F Register 3-operand Logical Operations on page 227

F Register 3-operand Logical Ops

7.65 F Register Logical Operate (3 operand) VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
FORd	0 0111 1100	Logical or	ford† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FORs	0 0111 1101	Logical or , 32-bit	fors <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FNORd	0 0110 0010	Logical nor	fnord† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FNORs	0 0110 0011	Logical nor , 32-bit	fnors <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FANDd	0 0111 0000	Logical and	fandd† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FANDs	0 0111 0001	Logical and , 32-bit	fands <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FNANDd	0 0110 1110	Logical nand	fnandd† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FNANDs	0 0110 1111	Logical nand , 32-bit	fnands <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FXORd	0 0110 1100	Logical xor	fxord† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FXORs	0 0110 1101	Logical xor , 32-bit	fxors <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FXNORd	0 0111 0010	Logical xnor	fxnord† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FXNORs	0 0111 0011	Logical xnor , 32-bit	fxnors <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FORNOT1d	0 0111 1010	(not F _D [rs1]) or F _D [rs2]	fornot1d† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FORNOT1s	0 0111 1011	(not F _S [rs1]) or F _S [rs2], 32-bit	fornot1s <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FORNOT2d	0 0111 0110	F _D [rs1] or (not F _D [rs2])	fornot2d† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FORNOT2s	0 0111 0111	F _S [rs1] or (not F _S [rs2]), 32-bit	fornot2s <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FANDNOT1d	0 0110 1000	(not F _D [rs1]) and F _D [rs2]	fandnot1d† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FANDNOT1s	0 0110 1001	(not F _S [rs1]) and F _S [rs2], 32-bit	fandnot1s <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FANDNOT2d	0 0110 0100	F _D [rs1] and (not F _D [rs2])	fandnot2d† <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1
FANDNOT2s	0 0110 0101	F _S [rs1] and (not F _S [rs2]), 32-bit	fandnot2s <i>freg_{rs1}, freg_{rs2}, freg_{rd}</i>	A1

† this assembly-language instruction mnemonic is also accepted without a trailing “d”, for backward compatibility



Description The standard 64-bit versions of these instructions perform one of ten 64-bit logical operations between the 64-bit floating-point registers F_D[rs1] and F_D[rs2]. The result is stored in the 64-bit floating-point destination register F_D[rd].

The 32-bit (single-precision) versions of these instructions perform 32-bit logical operations between F_S[rs1] and F_S[rs2], storing the result in F_S[rd].

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute any 3-operand F Register Logical Operate instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

See Also F Register 1-operand Logical Operations on page 225
F Register 2-operand Logical Operations on page 226

7.66 Partitioned Shift VIS 3

The Partitioned Shift instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	s3	Assembly Language Syntax			Class	Added
FSL16	0 0010 0001	16-bit partitioned shift left	f64	f64	f64	fsll16	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSRL16	0 0010 0011	16-bit partitioned shift right logical	f64	f64	f64	fsrl16	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSL32	0 0010 0101	32-bit partitioned shift left	f64	f64	f64	fsll32	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSRL32	0 0010 0111	32-bit partitioned shift right logical	f64	f64	f64	fsrl32	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSLAS16	0 0010 1001	16-bit partitioned shift left arithmetic with saturation	f64	f64	f64	fslas16	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSRA16	0 0010 1011	16-bit partitioned shift right arithmetic	f64	f64	f64	fsra16	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSLAS32	0 0010 1101	32-bit partitioned shift left arithmetic with saturation	f64	f64	f64	fslas32	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	
FSRA32	0 0010 1111	32-bit partitioned shift right arithmetic	f64	f64	f64	fsra32	<i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	N1	OSA 2011	



Description These operations shift data in each partition in $F_D[rs1]$ register according to the least significant bits of each corresponding partition in $F_D[rs2]$. 32-bit shift instructions use 5 bits in each of the two partitions in $F_D[rs2]$ to specify the shift amount. 16-bit shift instructions use 4 bits in each of the 4 partitions in $F_D[rs2]$ to specify the shift amount.




FIGURE 7-56 FSL16 Operation

FSL / FSRL / FSRA




FIGURE 7-57 FSL32 Operation

Shift right arithmetic and shift left arithmetic with saturation treat their first operand as a signed value. For FSLAS{16,32}, if a value opposite to the original sign is shifted through the sign position, the overflow result is clipped to the appropriate (same as the original sign) maximum positive or negative number representable by 16 (or 32) bits. Thus, the saturation values are $2^{15} - 1$ and -2^{15} for 16-bit partitioned saturating shifts and $2^{31} - 1$ and -2^{31} for 32-bit partitioned saturating shifts.


If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute any partitioned shift instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

FSQRT<s|d|q> Instructions

7.67 Floating-Point Square Root

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FSQRTs	11 0100	0 0010 1001	Square Root Single	<code>fsqrts</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	A1
FSQRTd	11 0100	0 0010 1010	Square Root Double	<code>fsqrtd</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	A1
FSQRTq	11 0100	0 0010 1011	Square Root Quad	<code>fsqrtq</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	C3



Description These SPARC V9 instructions generate the square root of the floating-point operand in the floating-point register(s) specified by the rs2 field and place the result in the destination floating-point register(s) specified by the rd field. Rounding is performed as specified by FSR.rd.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FSQRTq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FSQRT instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an FSQRT instruction causes an *fp_disabled* exception.

An attempt to execute an FSQRTq instruction when rs2{1} ≠ 0 or rd{1} ≠ 0 causes an *fp_exception_other* (FSR.ftt = *invalid_fp_register*) exception.

An *fp_exception_other* (with FSR.ftt = *unfinished_FPop*) can occur if the operand to the square root is positive and subnormal. See *FSR_floating-point_trap_type (ftt)* on page 55 for additional details.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (FSR.ftt = *invalid_fp_register* (FSQRTq only))
fp_exception_other (FSR.ftt = *unfinished_FPop*)
fp_exception_ieee_754 (IEEE_754_exception (NV, NX))

7.68 Convert Floating-Point to Integer

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FsTOx	0 1000 0001	Convert Single to 64-bit Integer	—	f32	f64	<code>fstox</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	A1
FdTOx	0 1000 0010	Convert Double to 64-bit Integer	—	f64	f64	<code>fdtox</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	A1
FqTOx	0 1000 0011	Convert Quad to 64-bit Integer	—	f128	f64	<code>fqtox</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	C3
FsTOi	0 1101 0001	Convert Single to 32-bit Integer	—	f32	f32	<code>fstoi</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	A1
FdTOi	0 1101 0010	Convert Double to 32-bit Integer	—	f64	f32	<code>fdtoi</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	A1
FqTOi	0 1101 0011	Convert Quad to 32-bit Integer	—	f128	f32	<code>fqtoi</code> <i>freq_{rs2}</i> , <i>freq_{rd}</i>	C3



Description FsTOx, FdTOx, and FqTOx convert the floating-point operand in the floating-point register(s) specified by `rs2` to a 64-bit integer in the floating-point register `FD[rd]`.

FsTOi, FdTOi, and FqTOi convert the floating-point operand in the floating-point register(s) specified by `rs2` to a 32-bit integer in the floating-point register `FS[rd]`.

The result is always rounded toward zero; that is, the rounding direction (`rd`) field of the FSR register is ignored.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute a FqTOx or FqTOi instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an F<s|d|q>TO<i|x> instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an F<s|d|q>TO<i|x> instruction causes an *fp_disabled* exception.

An attempt to execute an FqTOi or FqTOx instruction when `rs2{1} ≠ 0` causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

If the floating-point operand's value is too large to be converted to an integer of the specified size or is a NaN or infinity, then an *fp_exception_ieee_754* "invalid" exception occurs. The value written into the floating-point register(s) specified by `rd` in these cases is as defined in *Integer Overflow Definition* on page 393.


For more details regarding floating-point exceptions, see Chapter 8.

Exceptions

- illegal_instruction*
- fp_disabled*
- fp_exception_other* (`FSR.ftt = invalid_fp_register` (FqTOx and FqTOi only))
- fp_exception_other* (`FSR.ftt = unfinished_FPop`)
- fp_exception_ieee_754* (NV, NX)

7.69 Convert Between Floating-Point Formats

Instruction	op3	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FsTOd	11 0100	0 1100 1001	Convert Single to Double	—	f32	f64	<code>fstod</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FsTOq	11 0100	0 1100 1101	Convert Single to Quad	—	f32	f128	<code>fstoq</code> <i>freq_{rs2}, freq_{rd}</i>	C3
FdTOs	11 0100	0 1100 0110	Convert Double to Single	—	f64	f32	<code>fdtos</code> <i>freq_{rs2}, freq_{rd}</i>	A1
FdTOq	11 0100	0 1100 1110	Convert Double to Quad	—	f64	f128	<code>fdtoq</code> <i>freq_{rs2}, freq_{rd}</i>	C3
FqTOs	11 0100	0 1100 0111	Convert Quad to Single	—	f128	f32	<code>fqtos</code> <i>freq_{rs2}, freq_{rd}</i>	C3
FqTOd	11 0100	0 1100 1011	Convert Quad to Double	—	f128	f64	<code>fqtod</code> <i>freq_{rs2}, freq_{rd}</i>	C3



Description These instructions convert the floating-point operand in the floating-point register(s) specified by `rs2` to a floating-point number in the destination format. They write the result into the floating-point register(s) specified by `rd`.

The value of `FSR.rd` determines how rounding is performed by these instructions.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute a `FsTOq`, `FdTOq`, `FqTOs`, or `FqTOd` instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an `F<s|d|q>TO<s|d|q>` instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an `F<s|d|q>TO<s|d|q>` instruction causes an *fp_disabled* exception.

An attempt to execute an `FsTOq` or `FdTOq` instruction when `rd{1} ≠ 0` causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception. An attempt to execute an `FqTOs` or `FqTOd` instruction when `rs2{1} ≠ 0` causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

`FqTOd`, `FqTOs`, and `FdTOs` (the “narrowing” conversion instructions) can cause *fp_exception_ieee_754* OF, UF, and NX exceptions. `FdTOq`, `FsTOq`, and `FsTOd` (the “widening” conversion instructions) cannot.

Any of these six instructions can trigger an *fp_exception_ieee_754* NV exception if the source operand is a signalling NaN.

Note For `FdTOs`, and `FsTOd`, `FdTOq`, and `FqTOd`, an *fp_exception_other* with `FSR.ftt = unfinished_FPop` can occur if implementation-dependent conditions are detected during the conversion operation.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions

- illegal_instruction*
- fp_disabled*
- fp_exception_other* (`FSR.ftt = invalid_fp_register` (`FsTOq`, `FqTOs`, `FdTOq`, and `FqTOd` only))
- fp_exception_other* (`FSR.ftt = unfinished_FPop`)

F<sd|q>TO<sd|q>


fp_exception_ieee_754 (NV)

fp_exception_ieee_754 (OF, UF, NX (FqTOd, FqTOs, and FdTOs))

FSUB

7.70 Floating-Point Subtract

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FSUBs	11 0100	0 0100 0101	Subtract Single	<code>fsubs <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1
FSUBd	11 0100	0 0100 0110	Subtract Double	<code>fsubd <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	A1
FSUBq	11 0100	0 0100 0111	Subtract Quad	<code>fsubq <i>freg_{rs1}</i>, <i>freg_{rs2}</i>, <i>freg_{rd}</i></code>	C3



Description The floating-point subtract instructions subtract the floating-point register(s) specified by the `rs2` field from the floating-point register(s) specified by the `rs1` field. The instructions then write the difference into the floating-point register(s) specified by the `rd` field.

Rounding is performed as specified by `FSR.rd`.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute a `FSUBq` instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an `FSUB` instruction causes an *fp_disabled* exception.

An attempt to execute an `FSUBq` instruction when (`rs1{1} ≠ 0`) or (`rs2{1} ≠ 0`) or (`rd{1:0} ≠ 0`) causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

Note An *fp_exception_other* with `FSR.ftt = unfinished_FPop` can occur if the operation detects unusual, implementation-specific conditions (for `FSUBs` or `FSUBd`).

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (`FSR.ftt = invalid_fp_register` (`FSUBq` only))
fp_exception_other (`FSR.ftt = unfinished_FPop`)
fp_exception_ieee_754 (`OF`, `UF`, `NX`, `NV`)

■ **See Also** `FMAf` on page 175

7.71 Convert 64-bit Integer to Floating Point

Instruction	op3	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FxTOs	11 0100	0 1000 0100	Convert 64-bit Integer to Single	—	i64	f32	<code>fxtos <i>freq_{rs2}, freq_{rd}</i></code>	A1
FxTOd	11 0100	0 1000 1000	Convert 64-bit Integer to Double	—	i64	f64	<code>fxtod <i>freq_{rs2}, freq_{rd}</i></code>	A1
FxTOq	11 0100	0 1000 1100	Convert 64-bit Integer to Quad	—	i64	f128	<code>fxtoq <i>freq_{rs2}, freq_{rd}</i></code>	C3



Description FxTOs, FxTOd, and FxTOq convert the 64-bit signed integer operand in the floating-point register $F_D[rs2]$ into a floating-point number in the destination format.

All write their result into the floating-point register(s) specified by `rd`.

The value of `FSR.rd` determines how rounding is performed by FxTOs and FxTOd.

Note Oracle SPARC Architecture 2015 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute a FxTOq instruction causes an *illegal_instruction* exception, allowing privileged software to emulate the instruction.

An attempt to execute an FxTO<s|d|q> instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an FxTO<s|d|q> instruction causes an *fp_disabled* exception.

An attempt to execute an FxTOq instruction when `rd{1} ≠ 0` causes an *fp_exception_other* (`FSR.ftt = invalid_fp_register`) exception.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015*.

Exceptions *illegal_instruction*
fp_disabled
fp_exception_other (`FSR.ftt = invalid_fp_register` (FxTOq))
fp_exception_ieee_754 (NX (FxTOs and FxTOd only))

ILLTRAP

7.72 Illegal Instruction Trap

Instruction	op	op2	Operation	Assembly Language Syntax	Class
ILLTRAP	00	000	<i>illegal_instruction</i> trap	<code>illtrap const22</code>	A1



Description The ILLTRAP instruction causes an *illegal_instruction* exception. The `const22` value in the instruction is ignored by the virtual processor; specifically, this field is *not* reserved by the architecture for any future use.

V9 Compatibility | Except for its name, this instruction is identical to the SPARC V8
Note | UNIMP instruction.


An attempt to execute an ILLTRAP instruction when reserved instruction bits 29:25 are nonzero (also) causes an *illegal_instruction* exception. However, software should not rely on this behavior, because a future version of the architecture may use nonzero values of bits 29:25 to encode other functions.

Exceptions *illegal_instruction*

INVALW

7.73 Mark Register Window Sets as “Invalid”

Instruction	Operation	Assembly Language Syntax	Class
INVALW ^P	Mark all register window sets as “invalid”	invalw	A1



Description The INVALW instruction marks all register window sets as “invalid”; specifically, it atomically performs the following operations:

CANSAVE \leftarrow ($N_REG_WINDOWS - 2$)
 CANRESTORE \leftarrow 0
 OTHERWIN \leftarrow 0

Programming Notes INVALW marks all windows as invalid; after executing INVALW, $N_REG_WINDOWS-2$ SAVES can be performed without generating a spill trap.

An attempt to execute an INVALW instruction when instruction bits 18:0 are nonzero causes an *illegal_instruction* exception.

An attempt to execute an INVALW instruction in nonprivileged mode ($PSTATE.priv = 0$) causes a *privileged_opcode* exception.


Exceptions *illegal_instruction*
privileged_opcode

See Also ALLCLEAN on page 118
 NORMALW on page 295
 OTHERW on page 297
 RESTORED on page 317
 SAVED on page 324

JMPL

7.74 Jump and Link

Instruction	op3	Operation	Assembly Language Syntax	Class
JMPL	11 1000	Jump and Link	<code>jmp1 address, reg_{rd}</code>	A1



Description The JMPL instruction causes a register-indirect delayed control transfer to the address given by “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + **sign_ext**(simm13)” if $i = 1$.

The JMPL instruction copies the PC, which contains the address of the JMPL instruction, into register R[rd].

An attempt to execute a JMPL instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If either of the low-order two bits of the jump address is nonzero, a *mem_address_not_aligned* exception occurs.

Programming Notes A JMPL instruction with $rd = 15$ functions as a register-indirect call using the standard link register.

JMPL with $rd = 0$ can be used to return from a subroutine. The typical return address is “r[31] + 8” if a nonleaf routine (one that uses the SAVE instruction) is entered by a CALL instruction, or “R[15] + 8” if a leaf routine (one that does not use the SAVE instruction) is entered by a CALL instruction or by a JMPL instruction with $rd = 15$.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20) and `PSTATE.tct = 1`, then JMPL generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the JMPL instruction) is stored in `TPC[TL]` and the value of NPC from before the JMPL was executed is stored in `TNPC[TL]`.

When `PSTATE.am = 1`, the more-significant 32 bits of the target instruction address are masked out (set to 0) before being sent to the memory system or being written into R[rd] (or, if a *control_transfer_instruction* trap occurs, into `TPC[TL]`). (closed impl. dep. #125-V9-Cs10)


Exceptions *illegal_instruction*
mem_address_not_aligned
control_transfer_instruction (impl. dep. #450-S20)

See Also CALL on page 130
Bicc on page 123
BPCC on page 128

7.75 Load Integer

Instruction	op3	Operation	Assembly Language Syntax		Class
LDSB	00 1001	Load Signed Byte	ldsb	[address], reg _{rd}	A1
LDSH	00 1010	Load Signed Halfword	ldsh	[address], reg _{rd}	A1
LDSW	00 1000	Load Signed Word	ldsw	[address], reg _{rd}	A1
LDUB	00 0001	Load Unsigned Byte	ldub	[address], reg _{rd}	A1
LDUH	00 0010	Load Unsigned Halfword	lduh	[address], reg _{rd}	A1
LDUW	00 0000	Load Unsigned Word	lduw [†]	[address], reg _{rd}	A1
LDX	00 1011	Load Extended Word	ldx	[address], reg _{rd}	A1

[†] synonym: ld



Description The load integer instructions copy a byte, a halfword, a word, or an extended word from memory. All copy the fetched value into R[rd]. A fetched byte, halfword, or word is right-justified in the destination register R[rd]; it is either sign-extended or zero-filled on the left, depending on whether the opcode specifies a signed or unsigned operation, respectively.

Load integer instructions access memory using the implicit ASI (see page 83). The effective address is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + **sign_ext**(simm13)” if $i = 1$.

A successful load (notably, load extended) instruction operates atomically.

An attempt to execute a load integer instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Any of the following conditions cause a *mem_address_not_aligned* exception:

- an attempt to execute LDUH or LDSH when the effective address is not halfword-aligned
- an attempt to execute LDUW or LDSW when the effective address is not word-aligned
- an attempt to execute LDX when the effective address is not doubleword-aligned

V8 Compatibility Note The SPARC V8 LD instruction was renamed LDUW in the SPARC V9 architecture. The LDSW instruction was new in the SPARC V9 architecture.

A load integer twin word (LDTW) instruction exists, but is deprecated; see *Load Integer Twin Word* on page 257 for details.

Exceptions *illegal_instruction*
mem_address_not_aligned (all except LDSB, LDUB)
VA_watchpoint
DAE_privilege_violation
DAE_nfo_page

7.76 Load Integer from Alternate Space

Instruction	op3	Operation	Assembly Language Syntax		Class
LDSBA ^{PASI}	01 1001	Load Signed Byte from Alternate Space	ldsba	[regaddr] imm_asi, reg_rd ldsba [reg_plus_imm] %asi, reg_rd	A1
LDSHA ^{PASI}	01 1010	Load Signed Halfword from Alternate Space	ldsha	[regaddr] imm_asi, reg_rd ldsha [reg_plus_imm] %asi, reg_rd	A1
LDSWA ^{PASI}	01 1000	Load Signed Word from Alternate Space	ldswa	[regaddr] imm_asi, reg_rd ldswa [reg_plus_imm] %asi, reg_rd	A1
LDUBA ^{PASI}	01 0001	Load Unsigned Byte from Alternate Space	lduba	[regaddr] imm_asi, reg_rd lduba [reg_plus_imm] %asi, reg_rd	A1
LDUHA ^{PASI}	01 0010	Load Unsigned Halfword from Alternate Space	lduha	[regaddr] imm_asi, reg_rd lduha [reg_plus_imm] %asi, reg_rd	A1
LDUWA ^{PASI}	01 0000	Load Unsigned Word from Alternate Space	lduwa†	[regaddr] imm_asi, reg_rd lduwa [reg_plus_imm] %asi, reg_rd	A1
LDXA ^{PASI}	01 1011	Load Extended Word from Alternate Space	ldxa	[regaddr] imm_asi, reg_rd ldxa [reg_plus_imm] %asi, reg_rd	A1

† synonym: lda



Description The load integer from alternate space instructions copy a byte, a halfword, a word, or an extended word from memory. All copy the fetched value into R[rd]. A fetched byte, halfword, or word is right-justified in the destination register R[rd]; it is either sign-extended or zero-filled on the left, depending on whether the opcode specifies a signed or unsigned operation, respectively.

The load integer from alternate space instructions contain the address space identifier (ASI) to be used for the load in the imm_asi field if $i = 0$, or in the ASI register if $i = 1$. The effective address for these instructions is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + sign_ext(simm13)” if $i = 1$.

A successful load (notably, load extended) instruction operates atomically.

A load integer twin word from alternate space (LDTWA) instruction exists, but is deprecated; see *Load Integer Twin Word from Alternate Space* on page 259 for details.

Any of the following conditions cause a *mem_address_not_aligned* exception:

- an attempt to execute LDUHA or LDSHA when the effective address is not halfword-aligned
- an attempt to execute LDUWA or LDSWA when the effective address is not word-aligned
- an attempt to execute LDXA when the effective address is not doubleword-aligned

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, these instructions cause a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, these instructions cause a *privileged_action* exception.

LDA

LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, and LDUWA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with these instructions causes a *DAE_invalid_asi* exception.

ASIs valid for LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, and LDUWA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
12 ₁₆ (ASI_MONITOR_AS_IF_USER_PRIMARY)	13 ₁₆ (ASI_MONITOR_AS_IF_USER_SECONDARY)
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
15 ₁₆ ASI_REAL_IO	1D ₁₆ ASI_REAL_IO_LITTLE
<hr/>	
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE
82 ₁₆ ASI_PRIMARY_NO_FAULT	8A ₁₆ ASI_PRIMARY_NO_FAULT_LITTLE
83 ₁₆ ASI_SECONDARY_NO_FAULT	8B ₁₆ ASI_SECONDARY_NO_FAULT_LITTLE
84 ₁₆ (ASI_MONITOR_PRIMARY)	85 ₁₆ (ASI_MONITOR_SECONDARY)

LDXA can be used with any ASI (including, but not limited to, the above two lists), unless it either (a) violates the privilege mode rules described for the *privileged_action* exception above or (b) is used with any of the following ASIs, which causes a *DAE_invalid_asi* exception.

ASIs invalid for LDXA (cause <i>DAE_invalid_asi</i> exception)	
22 ₁₆ (ASI_TWIX_AIUP)	2A ₁₆ (ASI_TWIX_AIUP_L)
23 ₁₆ (ASI_TWIX_AIUS)	2B ₁₆ (ASI_TWIX_AIUS_L)
26 ₁₆ (ASI_TWIX_REAL)	2E ₁₆ (ASI_TWIX_REAL_L)
27 ₁₆ (ASI_TWIX_N)	2F ₁₆ (ASI_TWIX_NL)
ASI_BLOCK_AS_IF_USER_PRIMARY	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE
ASI_BLOCK_AS_IF_USER_SECONDARY	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE
C0 ₁₆ ASI_PST8_PRIMARY *	C8 ₁₆ ASI_PST8_PRIMARY_LITTLE *
C1 ₁₆ ASI_PST8_SECONDARY *	C9 ₁₆ ASI_PST8_SECONDARY_LITTLE *
C2 ₁₆ ASI_PST16_PRIMARY *	CA ₁₆ ASI_PST16_PRIMARY_LITTLE *
C3 ₁₆ ASI_PST16_SECONDARY *	CB ₁₆ ASI_PST16_SECONDARY_LITTLE *
C4 ₁₆ ASI_PST32_PRIMARY *	CC ₁₆ ASI_PST32_PRIMARY_LITTLE *
C5 ₁₆ ASI_PST32_SECONDARY *	CD ₁₆ ASI_PST32_SECONDARY_LITTLE *
D0 ₁₆ ASI_FL8_PRIMARY	D8 ₁₆ ASI_FL8_PRIMARY_LITTLE
D1 ₁₆ ASI_FL8_SECONDARY	D9 ₁₆ ASI_FL8_SECONDARY_LITTLE
D2 ₁₆ ASI_FL16_PRIMARY	DA ₁₆ ASI_FL16_PRIMARY_LITTLE
D3 ₁₆ ASI_FL16_SECONDARY	DB ₁₆ ASI_FL16_SECONDARY_LITTLE
E0 ₁₆ ASI_BLOCK_COMMIT_PRIMARY	E1 ₁₆ ASI_BLOCK_COMMIT_SECONDARY
E2 ₁₆ (ASI_TWIX_P)	EA ₁₆ (ASI_TWIX_PL)
E3 ₁₆ (ASI_TWIX_S)	EB ₁₆ (ASI_TWIX_SL)
F0 ₁₆ ASI_BLOCK_PRIMARY	F8 ₁₆ ASI_BLOCK_PRIMARY_LITTLE
F1 ₁₆ ASI_BLOCK_SECONDARY	F9 ₁₆ ASI_BLOCK_SECONDARY_LITTLE

Exceptions

mem_address_not_aligned (all except LDSBA and LDUBA)
privileged_action
VA_watchpoint
DAE_invalid_asi
DAE_privilege_violation
DAE_nfo_page
DAE_side_effect_page

LDA

| *See Also*

LD on page 239
STA on page 335


LDBLOCKF

7.77 Block Load VIS 1

The LDBLOCKF instructions are deprecated and should not be used in new software. A sequence of LDDF instructions should be used instead.

The LDBLOCKF instruction is intended to be a processor-specific instruction, which may or may not be implemented in future Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruc-tion	ASI Value	Operation	Assembly Language Syntax	Class
LDBLOCKF ^D	16 ₁₆	64-byte block load from primary address space, user privilege	ldda [regaddr] #ASI_BLK_AIUP, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	17 ₁₆	64-byte block load from secondary address space, user privilege	ldda [regaddr] #ASI_BLK_AIUS, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	1E ₁₆	64-byte block load from primary address space, little-endian, user privilege	ldda [regaddr] #ASI_BLK_AIUPL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	1F ₁₆	64-byte block load from secondary address space, little-endian, user privilege	ldda [regaddr] #ASI_BLK_AIUSL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	F0 ₁₆	64-byte block load from primary address space	ldda [regaddr] #ASI_BLK_P, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	F1 ₁₆	64-byte block load from secondary address space	ldda [regaddr] #ASI_BLK_S, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	F8 ₁₆	64-byte block load from primary address space, little-endian	ldda [regaddr] #ASI_BLK_PL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2
LDBLOCKF ^D	F9 ₁₆	64-byte block load from secondary address space, little-endian	ldda [regaddr] #ASI_BLK_SL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	D2



Description A block load (LDBLOCKF) instruction uses one of several special block-transfer ASIs. Block transfer ASIs allow block loads to be performed accessing the same address space as normal loads. Little-endian ASIs (those with an 'L' suffix) access data in little-endian format; otherwise, the access is assumed to be big-endian. Byte swapping is performed separately for each of the eight 64-bit (double-precision) F registers used by the instruction.

A block load instruction loads 64 bytes of data from a 64-byte aligned memory area into the eight double-precision floating-point registers specified by rd. The lowest-addressed eight bytes in memory are loaded into the lowest-numbered 64-bit (double-precision) destination F register.

A block load only guarantees atomicity for each 64-bit (8-byte) portion of the 64 bytes it accesses.

Programming Note The block load instruction, LDBLOCKF^D, and its companion, STBLOCKF^D, were originally defined to provide a fast mechanism for block-copy operations. However, in modern implementations they are rarely much faster than a sequence of regular loads and stores, so are now deprecated.

LDBLOCKF

Programming Note LDBLOCKF^D is intended to be a processor-specific instruction (see the warning at the top of page 243). If LDBLOCKF^D *must* be used in software intended to be portable across current and previous processor implementations, then it must be coded to work in the face of any implementation variation that is permitted by implementation dependency #410-S10, described below.

IMPL. DEP. #410-S10: The following aspects of the behavior of block load (LDBLOCKF) instructions are implementation dependent:

- What memory ordering model is used by LDBLOCKF^D (LDBLOCKF^D is not required to follow TSO memory ordering)
- Whether LDBLOCKF^D follows memory ordering with respect to stores (including block stores), including whether the virtual processor detects read-after-write and write-after-read hazards to overlapping addresses
- Whether LDBLOCKF^D appears to execute out of order, or follow LoadLoad ordering (with respect to older loads, younger loads, and other LDBLOCKFs)
- Whether LDBLOCKF^D follows register-dependency interlocks, as do ordinary load instructions
- Whether *VA_watchpoint* exceptions are recognized on accesses to all 64 bytes of a LDBLOCKF^D (the recommended behavior), or only on the first eight bytes
- Whether the MMU ignores the side-effect bit (TTE.e) for LDBLOCKF^D accesses

Programming Note If ordering with respect to earlier stores is important (for example, a block load that overlaps a previous store) and read-after-write hazards are not detected, there must be a MEMBAR #StoreLoad instruction between earlier stores and a block load.

If ordering with respect to later stores is important, there must be a MEMBAR #LoadStore instruction between a block load and subsequent stores.

If LoadLoad ordering with respect to older or younger loads or other block load instructions is important and is not provided by an implementation, an intervening MEMBAR #LoadLoad is required.

For further restrictions on the behavior of the block load instruction, see implementation-specific processor documentation.

Implementation Note In all Oracle SPARC Architecture implementations, the MMU ignores the side-effect bit (TTE.e) for LDBLOCKF^D accesses (impl. dep. #410-S10).

Exceptions. An *illegal_instruction* exception occurs if LDBLOCKF's floating-point destination registers are not aligned on an eight-double-precision register boundary.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an LDBLOCKF^D instruction causes an *fp_disabled* exception.

If the least significant 6 bits of the effective memory address in an LDBLOCKF^D instruction are nonzero, a *mem_address_not_aligned* exception occurs.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0 (ASIs 16₁₆, 17₁₆, 1E₁₆, and 1F₁₆), LDBLOCKF^D causes a *privileged_action* exception.

An access caused by LDBLOCKF^D may trigger a *VA_watchpoint* exception (impl. dep. #410-S10).

An attempted access by an LDBLOCKF^D instruction to noncacheable memory causes an *DAE_nc_page* exception.

LDBLOCKF

Implementation | LDBLOCKF^D shares an opcode with LDDFA and LDSHORTEF;
Note | they are distinguished by the ASI used.

Exceptions

- illegal_instruction*
- fp_disabled*
- mem_address_not_aligned*
- privileged_action*
- VA_watchpoint* (impl. dep. #410-S10)
- DAE_privilege_violation*
- DAE_nc_page*

See Also

- LDDF on page 246
- STBLOCKF^D on page 337

7.78 Load Floating-Point Register

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
LDF	10 0000	0–31	Load Floating-Point Register	ld [address], freg _{rd}	A1
LDDF	10 0011	‡	Load Double Floating-Point Register	ldd [address], freg _{rd}	A1
LDQF	10 0010	‡	Load Quad Floating-Point Register	ldq [address], freg _{rd}	C3

‡ Encoded floating-point register value, as described on page 51.



Description The load single floating-point instruction (LDF) copies a word from memory into 32-bit floating-point destination register $F_S[rd]$.

The load doubleword floating-point instruction (LDDF) copies a word-aligned doubleword from memory into a 64-bit floating-point destination register, $F_D[rd]$. The unit of atomicity for LDDF is 4 bytes (one word).

The load quad floating-point instruction (LDQF) copies a word-aligned quadword from memory into a 128-bit floating-point destination register, $F_Q[rd]$. The unit of atomicity for LDQF is 4 bytes (one word).

These load floating-point instructions access memory using the implicit ASI (see page 83).

If $i = 0$, the effective address for these instructions is “ $R[rs1] + R[rs2]$ ” and if $i = 1$, the effective address is “ $R[rs1] + \text{sign_ext}(simm13)$ ”.

Exceptions. An attempt to execute an LDF, LDDF, or LDQF instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an LDF, LDDF, or LDQF instruction causes an *fp_disabled* exception.

Any of the following conditions cause an exception:

- an attempt to execute LDF, LDDF, or LDQF when the effective address is not word-aligned causes a *mem_address_not_aligned* exception
- an attempt to execute LDDF when the effective address is word-aligned but not doubleword-aligned causes an *LDDF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the LDDF instruction and return ((impl. dep. #109-V9-Cs10(a))).
- an attempt to execute LDQF when the effective address is word-aligned but not quadword-aligned causes an *LDQF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the LDQF instruction and return (impl. dep. #111-V9-Cs10(a)).

Programming Note Some compilers issued sequences of single-precision loads for SPARC V8 processor targets when the compiler could not determine whether doubleword or quadword operands were properly aligned. For SPARC V9 processors, since emulation of misaligned loads is expected to be fast, compilers should issue sets of single-precision loads only when they can determine that doubleword or quadword operands are *not* properly aligned.

LDF / LDDF / LDQF

An attempt to execute an LDQF instruction when $rd\{1\} \neq 0$ causes an *fp_exception_other* (FSR.ftt = *invalid_fp_register*) exception.

Implementation Note	Since Oracle SPARC Architecture 2015 processors do not implement in hardware instructions (including LDQF) that refer to quad-precision floating-point registers, the <i>LDQF_mem_address_not_aligned</i> and <i>fp_exception_other</i> (with FSR.ftt = <i>invalid_fp_register</i>) exceptions do not occur in hardware. However, their effects must be emulated by software when the instruction causes an <i>illegal_instruction</i> exception and subsequent trap.
----------------------------	--

Destination Register(s) when Exception Occurs. If a load floating-point instruction generates an exception that causes a *precise* trap, the destination floating-point register(s) remain unchanged.

IMPL. DEP. #44-V8-Cs10(a)(1): If a load floating-point instruction generates an exception that causes a *non-precise* trap, the contents of the destination floating-point register(s) remain unchanged or are undefined.

Exceptions

illegal_instruction
fp_disabled
LDDF_mem_address_not_aligned
LDQF_mem_address_not_aligned (not used in Oracle SPARC Architecture 2015)
mem_address_not_aligned
fp_exception_other (FSR.ftt = *invalid_fp_register* (LDQF only))
VA_watchpoint
DAE_privilege_violation
DAE_nfo_page

See Also

Load Floating-Point from Alternate Space on page 248
Load Floating-Point State Register (Lower) on page 251
Store Floating-Point on page 340

7.79 Load Floating-Point from Alternate Space

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
L DFA ^{P_{ASI}}	11 0000	0–31	Load Floating-Point Register from Alternate Space	l da [regaddr] imm_asi, freg _{rd} l da [reg_plus_imm] %asi, freg _{rd}	A1
LDDFA ^{P_{ASI}}	11 0011	‡	Load Double Floating-Point Register from Alternate Space	l dda [regaddr] imm_asi, freg _{rd} l dda [reg_plus_imm] %asi, freg _{rd}	A1
LDQFA ^{P_{ASI}}	11 0010	‡	Load Quad Floating-Point Register from Alternate Space	l dqa [regaddr] imm_asi, freg _{rd} l dqa [reg_plus_imm] %asi, freg _{rd}	C3

‡ Encoded floating-point register value, as described in *Floating-Point Register Number Encoding* on page 51.



Description The load single floating-point from alternate space instruction (L DFA) copies a word from memory into 32-bit floating-point destination register $F_S[rd]$.

The load double floating-point from alternate space instruction (LDDFA) copies a word-aligned doubleword from memory into a 64-bit floating-point destination register, $F_D[rd]$. The unit of atomicity for LDDFA is 4 bytes (one word).

The load quad floating-point from alternate space instruction (LDQFA) copies a word-aligned quadword from memory into a 128-bit floating-point destination register, $F_Q[rd]$. The unit of atomicity for LDQFA is 4 bytes (one word).

If $i = 0$, these instructions contain the address space identifier (ASI) to be used for the load in the `imm_asi` field and the effective address for the instruction is “ $R[rs1] + R[rs2]$ ”. If $i = 1$, the ASI to be used is contained in the ASI register and the effective address for the instruction is “ $R[rs1] + \text{sign_ext}(\text{simm13})$ ”.

Exceptions. If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an L DFA, LDDFA, or LDQFA instruction causes an *fp_disabled* exception.

V9 Compatibility | L DFA, LDDFA, and LDQFA cause a *privileged_action* exception
Note | if `PSTATE.priv = 0` and bit 7 of the ASI is 0.

Any of the following conditions cause an exception:

- an attempt to execute L DFA, LDDFA, or LDQFA when the effective address is not word-aligned causes a *mem_address_not_aligned* exception
- an attempt to execute LDDFA when the effective address is word-aligned but not doubleword-aligned causes an *LDDF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the LDDFA instruction and return ((impl. dep. #109-V9-Cs10(a))).
- an attempt to execute LDQFA when the effective address is word-aligned but not quadword-aligned causes an *LDQF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the LDQFA instruction and return (impl. dep. #111-V9-Cs10(a)).

An attempt to execute an LDQFA instruction when `rd{1} ≠ 0` causes an *fp_exception_other* (with `FSR.ftt = invalid_fp_register`) exception.

LDFA / LDDFA / LDQFA

Implementation Note Since Oracle SPARC Architecture 2015 processors do not implement in hardware instructions (including LDQFA) that refer to quad-precision floating-point registers, the *LDQF_mem_address_not_aligned* and *fp_exception_other* (with *FSR.ftt = invalid_fp_register*) exceptions do not occur in hardware. However, their effects must be emulated by software when the instruction causes an *illegal_instruction* exception and subsequent trap.

Programming Note Some compilers issued sequences of single-precision loads for SPARC V8 processor targets when the compiler could not determine whether doubleword or quadword operands were properly aligned. For SPARC V9 processors, since emulation of misaligned loads is expected to be fast, compilers should issue sets of single-precision loads only when they can determine that doubleword or quadword operands are *not* properly aligned.

In nonprivileged mode (*PSTATE.priv = 0*), if bit 7 of the ASI is 0, this instruction causes a *privileged_action* exception. In privileged mode (*PSTATE.priv = 1*), if the ASI is in the range 30_{16} to $7F_{16}$, this instruction causes a *privileged_action* exception.

LDFA and LDQFA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with these instructions causes a *DAE_invalid_asi* exception.

ASIs valid for LDFA and LDQFA	
04_{16} ASI_NUCLEUS	$0C_{16}$ ASI_NUCLEUS_LITTLE
10_{16} ASI_AS_IF_USER_PRIMARY	18_{16} ASI_AS_IF_USER_PRIMARY_LITTLE
11_{16} ASI_AS_IF_USER_SECONDARY	19_{16} ASI_AS_IF_USER_SECONDARY_LITTLE
14_{16} ASI_REAL	$1C_{16}$ ASI_REAL_LITTLE
15_{16} ASI_REAL_IO	$1D_{16}$ ASI_REAL_IO_LITTLE
80_{16} ASI_PRIMARY	88_{16} ASI_PRIMARY_LITTLE
81_{16} ASI_SECONDARY	89_{16} ASI_SECONDARY_LITTLE
82_{16} ASI_PRIMARY_NO_FAULT	$8A_{16}$ ASI_PRIMARY_NO_FAULT_LITTLE
83_{16} ASI_SECONDARY_NO_FAULT	$8B_{16}$ ASI_SECONDARY_NO_FAULT_LITTLE

LDDFA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with the LDDFA instruction causes a *DAE_invalid_asi* exception.

ASIs valid for LDDFA	
04_{16} ASI_NUCLEUS	$0C_{16}$ ASI_NUCLEUS_LITTLE
10_{16} ASI_AS_IF_USER_PRIMARY	18_{16} ASI_AS_IF_USER_PRIMARY_LITTLE
11_{16} ASI_AS_IF_USER_SECONDARY	19_{16} ASI_AS_IF_USER_SECONDARY_LITTLE
14_{16} ASI_REAL	$1C_{16}$ ASI_REAL_LITTLE
15_{16} ASI_REAL_IO	$1D_{16}$ ASI_REAL_IO_LITTLE
80_{16} ASI_PRIMARY	88_{16} ASI_PRIMARY_LITTLE
81_{16} ASI_SECONDARY	89_{16} ASI_SECONDARY_LITTLE
82_{16} ASI_PRIMARY_NO_FAULT	$8A_{16}$ ASI_PRIMARY_NO_FAULT_LITTLE
83_{16} ASI_SECONDARY_NO_FAULT	$8B_{16}$ ASI_SECONDARY_NO_FAULT_LITTLE

Behavior with Block-Store-with-Commit ASIs. ASIs $E0_{16}$ and $E1_{16}$ are only defined for use in Block Store with Commit operations (see page 337). Neither ASI $E0_{16}$ nor $E1_{16}$ should be used with LDDFA; however, if it is used, the LDDFA behaves as follows:

LDFA / LDDFA / LDQFA

1. If an LDDFA opcode is used with an ASI of $E0_{16}$ or $E1_{16}$ and a destination register number rd is specified which is not a multiple of 8 (“misaligned” rd), an Oracle SPARC Architecture 2015 virtual processor generates an *illegal_instruction* exception (impl. dep. #255-U3-Cs10).
2. **IMPL. DEP. #256-U3:** If an LDDFA opcode is used with an ASI of $E0_{16}$ or $E1_{16}$ and a memory address is specified with less than 64-byte alignment, the virtual processor generates an exception. It is implementation dependent whether the exception generated is *DAE_invalid_asi*, *mem_address_not_aligned*, or *LDDF_mem_address_not_aligned*.
3. If both rd and the memory address are correctly aligned, a *DAE_invalid_asi* exception occurs.

Behavior with Partial Store ASIs. ASIs $C0_{16}$ – $C5_{16}$ and $C8_{16}$ – CD_{16} are only defined for use in Partial Store operations (see page 347). None of them should be used with LDDFA; however, if any of those ASIs is used with LDDFA, the LDDFA behaves as follows:

1. **IMPL. DEP. #257-U3:** If an LDDFA opcode is used with an ASI of $C0_{16}$ – $C5_{16}$ or $C8_{16}$ – CD_{16} (Partial Store ASIs, which are an illegal combination with LDDFA) and a memory address is specified with less than 8-byte alignment, the virtual processor generates an exception. It is implementation dependent whether the generated exception is a *DAE_invalid_asi*, *mem_address_not_aligned*, or *LDDF_mem_address_not_aligned* exception.
2. If the memory address is correctly aligned, the virtual processor generates a *DAE_invalid_asi*.

Destination Register(s) when Exception Occurs. If a load floating-point alternate instruction generates an exception that causes a precise trap, the destination floating-point register(s) remain unchanged.

IMPL. DEP. #44-V8-Cs10(b): If a load floating-point alternate instruction generates an exception that causes a non-precise trap, it is implementation dependent whether the contents of the destination floating-point register(s) are undefined or are guaranteed to remain unchanged.

Implementation Note | LDDFA shares an opcode with the LDBLOCKF^D and LDSHORTF instructions; they are distinguished by the ASI used.

Exceptions

illegal_instruction
fp_disabled
LDDF_mem_address_not_aligned
LDQF_mem_address_not_aligned (not generated in Oracle SPARC Architecture 2015)
mem_address_not_aligned
fp_exception_other (FSR.ftt = invalid_fp_register (LDQFA only))
privileged_action
VA_watchpoint
DAE_invalid_asi
DAE_privilege_violation
DAE_nfo_page
DAE_side_effect_page

See Also


Load Floating-Point Register on page 246
Block Load on page 243
Store Short Floating-Point on page 350
Store Floating-Point into Alternate Space on page 342

LDFSR (Deprecated)

7.80 Load Floating-Point State Register (Lower)

The LDFSR instruction is deprecated and should not be used in new software. The LDXFSR instruction should be used instead.

Opcode	op3	rd	Operation	Assembly Language Syntax	Class
LDFSR ^D	10 0001	0	Load Floating-Point State Register (Lower)	ld [address], %fsr	D2
	10 0001	1-31	(see page 264)		



Description The Load Floating-point State Register (Lower) instruction (LDFSR) waits for all FPop instructions that have not finished execution to complete and then loads a word from memory into the less significant 32 bits of the FSR. The more-significant 32 bits of FSR are unaffected by LDFSR. LDFSR does not alter the *ver*, *ftt*, *qne*, *reserved*, or *unimplemented* (for example, *ns*) fields of FSR (see page 42).

Programming Note For future compatibility, software should only issue an LDFSR instruction with a zero value (or a value previously read from the same field) in any reserved field of FSR.

LDFSR accesses memory using the implicit ASI (see page 83).

An attempt to execute an LDFSR instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (*FPRS.fef* = 0 or *PSTATE.pef* = 0) or if no FPU is present, an attempt to execute an LDFSR instruction causes an *fp_disabled* exception.

LDFSR causes a *mem_address_not_aligned* exception if the effective memory address is not word-aligned.

V8 Compatibility Note The SPARC V9 architecture supports two different instructions to load the FSR: the (deprecated) SPARC V8 LDFSR instruction is defined to load only the less-significant 32 bits of the FSR, whereas LDXFSR allows SPARC V9 programs to load all 64 bits of the FSR.

Implementation Note LDFSR shares an opcode with the LDXFSR and LDXEFSR instructions (and possibly with other implementation-dependent instructions); they are differentiated by the instruction *rd* field. An attempt to execute the *op* = 11₂, *op3* = 10 0001₂ opcode with an invalid *rd* value causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
fp_disabled
mem_address_not_aligned
VA_watchpoint

LDFSR (Deprecated)

DAE_privilege_violation
DAE_nfo_page

See Also

Load Floating-Point Register on page 246
Load Floating-Point State Register on page 264
Store Floating-Point on page 340


DAE_invalid_aside

LDSHORTF

7.81

DAE_invalid_asi Load Short Floating-Point VIS 1

Instruction	ASI Value	Operation	Assembly Language Syntax		Class
LDSHORTF	D0 ₁₆	8-bit load from primary address space	ldda	[regaddr] #ASI_FL8_P, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	D1 ₁₆	8-bit load from secondary address space	ldda	[regaddr] #ASI_FL8_S, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	D8 ₁₆	8-bit load from primary address space, little-endian	ldda	[regaddr] #ASI_FL8_PL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	D9 ₁₆	8-bit load from secondary address space, little-endian	ldda	[regaddr] #ASI_FL8_SL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	D2 ₁₆	16-bit load from primary address space	ldda	[regaddr] #ASI_FL16_P, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	D3 ₁₆	16-bit load from secondary address space	ldda	[regaddr] #ASI_FL16_S, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	DA ₁₆	16-bit load from primary address space, little-endian	ldda	[regaddr] #ASI_FL16_PL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1
LDSHORTF	DB ₁₆	16-bit load from secondary address space, little-endian	ldda	[regaddr] #ASI_FL16_SL, freg _{rd} ldda [reg_plus_imm] %asi, freg _{rd}	B1



Description Short floating-point load instructions allow an 8- or 16-bit value to be loaded from memory into a 64-bit floating-point register.

An 8-bit load places the loaded value in the least significant byte of $F_D[rd]$ and zeroes in the most-significant seven bytes of $F_D[rd]$. An 8-bit LDSHORTF can be performed from an arbitrary byte address.

A 16-bit load places the loaded value in the least significant halfword of $F_D[rd]$ and zeroes in the more-significant six bytes of $F_D[rd]$.

Little-endian ASIs transfer data in little-endian format from memory; otherwise, memory is assumed to be in big-endian byte order.

Programming Note LDSHORTF is typically used with the FALIGNDATA_g or FALIGNDATA_i instruction (see *Align Data (using gsr.align)* on page 154) to assemble or store 64 bits from noncontiguous components.

Implementation Note LDSHORTF shares an opcode with the LDBLOCKF^D and LDDFA instructions; they are distinguished by the ASI used.

Exceptions. If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an LDSHORTF instruction causes an *fp_disabled* exception.

An attempt to execute a 16-bit (halfword) LDSHORTF instruction when the effective address is not halfword-aligned causes a *mem_address_not_aligned* exception.

LDSHORTF


Exceptions *fp_disabled*
mem_address_not_aligned
VA_watchpoint
DAE_privilege_violation
DAE_nfo_page

See Also *STSHORTF* on page 350
Align Data (using gsr.align) on page 154
Align Data (using Integer register) on page 155

LDSTUB

7.82 Load-Store Unsigned Byte

Instruction	op3	Operation	Assembly Language Syntax	Class
LDSTUB	00 1101	Load-Store Unsigned Byte	ldstub [address], reg _{rd}	A1



Description The load-store unsigned byte instruction copies a byte from memory into R[rd], then rewrites the addressed byte in memory to all 1's. The fetched byte is right-justified in the destination register R[rd] and zero-filled on the left.

The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more virtual processors executing LDSTUB, LDSTUBA, CASA, CASXA, SWAP, or SWAPA instructions addressing all or parts of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial, order.

LDSTUB accesses memory using the implicit ASI (see page 83). The effective address for this instruction is "R[rs1] + R[rs2]" if $i = 0$, or "R[rs1] + **sign_ext**(simm13)" if $i = 1$.

The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent (impl. dep. #120-V9).

An attempt to execute an LDSTUB instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
VA_watchpoint
DAE_privilege_violation
DAE_nc_page
DAE_nfo_page

See Also CASA on page 134
LDSTUBA on page 256
SWAP on page 359

LDSTUBA

7.83 Load-Store Unsigned Byte to Alternate Space

Instruction	op3	Operation	Assembly Language Syntax	Class
LDSTUBA ^{PASI}	01 1101	Load-Store Unsigned Byte into Alternate Space	ldstuba [<i>regaddr</i>] <i>imm_asi</i> , <i>reg_rd</i> ldstuba [<i>reg_plus_imm</i>] % <i>asi</i> , <i>reg_rd</i>	A1



Description The load-store unsigned byte into alternate space instruction copies a byte from memory into R[rd], then rewrites the addressed byte in memory to all 1's. The fetched byte is right-justified in the destination register R[rd] and zero-filled on the left.

The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more virtual processors executing LDSTUB, LDSTUBA, CASA, CASXA, SWAP, or SWAPA instructions addressing all or parts of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial, order.

If $i = 0$, LDSTUBA contains the address space identifier (ASI) to be used for the load in the *imm_asi* field. If $i = 1$, the ASI is found in the ASI register. In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30_{16} to $7F_{16}$, this instruction causes a *privileged_action* exception.

LDSTUBA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with this instruction causes a *DAE_invalid_asi* exception.

ASIs valid for LDSTUBA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE

Exceptions *privileged_action*
VA_watchpoint
DAE_invalid_asi
DAE_privilege_violation
DAE_nc_page
DAE_nfo_page

See Also CASA on page 134
LDSTUB on page 255
SWAP on page 359
SWAPA on page 360


LDTW (Deprecated)

7.84 Load Integer Twin Word

The LDTW instruction is deprecated and should not be used in new software. It is provided only for compatibility with previous versions of the architecture. The LDX instruction should be used instead.

Instruction	op3	Operation	Assembly Language Syntax †	Class
LDTW ^D	00 0011	Load Integer Twin Word	ldtw [address], reg _{rd}	D2

† The original assembly language syntax for this instruction used an “ldd” instruction mnemonic, which is now deprecated. Over time, assemblers will support the new “ldtw” mnemonic for this instruction. In the meantime, some existing assemblers may only recognize the original “ldd” mnemonic.



Description The load integer twin word instruction (LDTW) copies two words (with doubleword alignment) from memory into a pair of R registers. The word at the effective memory address is copied into the least significant 32 bits of the even-numbered R register. The word at the effective memory address + 4 is copied into the least significant 32 bits of the following odd-numbered R register. The most significant 32 bits of both the even-numbered and odd-numbered R registers are zero-filled.

Note Execution of an LDTW instruction with $rd = 0$ modifies only R[1].

Load integer twin word instructions access memory using the implicit ASI (see page 83). If $i = 0$, the effective address for these instructions is “R[rs1] + R[rs2]” and if $i = 1$, the effective address is “R[rs1] + `sign_ext(simm13)`”.

With respect to little endian memory, an LDTW instruction behaves as if it comprises two 32-bit loads, each of which is byte-swapped independently before being written into its respective destination register.

IMPL. DEP. #107-V9a: It is implementation dependent whether LDTW is implemented in hardware. If not, an attempt to execute an LDTW instruction will cause an *unimplemented_LDTW* exception.

Programming Note LDTW is provided for compatibility with existing SPARC V8 software. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties.

SPARC V9 Compatibility Note LDTW was (inaccurately) named LDD in the SPARC V8 and SPARC V9 specifications. It does not load a doubleword; it loads two words (into two registers), and has been renamed accordingly.

The least significant bit of the *rd* field in an LDTW instruction is unused and should always be set to 0 by software. An attempt to execute an LDTW instruction that refers to a misaligned (odd-numbered) destination register causes an *illegal_instruction* exception.

An attempt to execute an LDTW instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the effective address is not doubleword-aligned, an attempt to execute an LDTW instruction causes a *mem_address_not_aligned* exception.

LDTW (Deprecated)

A successful LDTW instruction operates atomically.

Programming Notes	<p>LDTW is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties. Therefore, software should avoid using LDTW.</p> <p>If LDTW is emulated in software, an LDX instruction should be used for the memory access in the emulation code to preserve atomicity. Emulation software should examine TSTATE[TL].pstate.cle (and, if appropriate, TTE.ie) to determine the endianness of the emulated memory access.</p> <p>Note that the value of TTE.ie is not saved during a trap. Therefore, if it is examined in the emulation trap handler, that should be done as quickly as possible, to minimize the window of time during which the value of TTE.ie could possibly be changed from the value it had at the time of the attempted execution of LDTW.</p>
--------------------------	--

Exceptions

- unimplemented_LDTW* (not used in Oracle SPARC Architecture 2015)
- illegal_instruction*
- mem_address_not_aligned*
- VA_watchpoint*
- DAE_privilege_violation*
- DAE_nfo_page*

See Also

- LDW/LDX on page 239
- STTW on page 352

LDTWA (Deprecated)

7.85 Load Integer Twin Word from Alternate Space

The LDTWA instruction is deprecated and should not be used in new software. The LDXA instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
LDTWA ^{D, P_{ASI}}	01 0011	Load Integer Twin Word from Alternate Space	ldtwa [<i>regaddr</i>] <i>imm_asi</i> , <i>reg_{rd}</i> ldtwa [<i>reg_plus_imm</i>] % <i>asi</i> , <i>reg_{rd}</i>	D2, Y3‡

† The original assembly language syntax for this instruction used an “ldda” instruction mnemonic, which is now deprecated. Over time, assemblers will support the new “ldtwa” mnemonic for this instruction. In the meantime, some assemblers may only recognize the original “ldda” mnemonic.

‡ Y3 for restricted ASIs (00₁₆-7F₁₆); D2 for unrestricted ASIs (80₁₆-FF₁₆)



Description The load integer twin word from alternate space instruction (LDTWA) copies two 32-bit words from memory (with doubleword memory alignment) into a pair of R registers. The word at the effective memory address is copied into the least significant 32 bits of the even-numbered R register. The word at the effective memory address + 4 is copied into the least significant 32 bits of the following odd-numbered R register. The most significant 32 bits of both the even-numbered and odd-numbered R registers are zero-filled.

Note Execution of an LDTWA instruction with *rd* = 0 modifies only R[1].

If *i* = 0, the LDTWA instruction contains the address space identifier (ASI) to be used for the load in its *imm_asi* field and the effective address for the instruction is “R[*rs1*] + R[*rs2*]”. If *i* = 1, the ASI to be used is contained in the ASI register and the effective address for the instruction is “R[*rs1*] + *sign_ext*(*simm13*)”.

With respect to little endian memory, an LDTWA instruction behaves as if it is composed of two 32-bit loads, each of which is byte-swapped independently before being written into its respective destination register.

IMPL. DEP. #107-V9b: It is implementation dependent whether LDTWA is implemented in hardware. If not, an attempt to execute an LDTWA instruction will cause an *unimplemented_LDTW* exception so that it can be emulated.

Programming Notes LDTWA is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties. Therefore, software should avoid using LDTWA.

If LDTWA is emulated in software, an LDXA instruction should be used for the memory access in the emulation code to preserve atomicity. Emulation software should examine TSTATE[TL].pstate.cle (and, if appropriate, TTE.ie) to determine the endianness of the emulated memory access.

LDTWA (Deprecated)

Note that the value of TTE.ie is not saved during a trap. Therefore, if it is examined in the emulation trap handler, that should be done as quickly as possible, to minimize the window of time during which the value of TTE.ie could possibly be changed from the value it had at the time of the attempted execution of LDTWA.

SPARC V9 Compatibility Note | LDTWA was (inaccurately) named LDDA in the SPARC V8 and SPARC V9 specifications.

The least significant bit of the rd field in an LDTWA instruction is unused and should always be set to 0 by software. An attempt to execute an LDTWA instruction that refers to a misaligned (odd-numbered) destination register causes an *illegal_instruction* exception.

If the effective address is not doubleword-aligned, an attempt to execute an LDTWA instruction causes a *mem_address_not_aligned* exception.

A successful LDTWA instruction operates atomically.

LDTWA causes a *mem_address_not_aligned* exception if the address is not doubleword-aligned.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, these instructions cause a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, these instructions cause a *privileged_action* exception.

LDTWA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with this instruction causes a *DAE_invalid_asi* exception (impl. dep. #300-U4-Cs10).

ASIs valid for LDTWA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
15 ₁₆ ASI_REAL_IO	1D ₁₆ ASI_REAL_IO_LITTLE
22 ₁₆ ‡ (ASI_TWINK_AIUP)	2A ₁₆ ‡ (ASI_TWINK_AIUP_L)
23 ₁₆ ‡ (ASI_TWINK_AIUS)	2B ₁₆ ‡ (ASI_TWINK_AIUS_L)
26 ₁₆ ‡ (ASI_TWINK_REAL)	2E ₁₆ ‡ (ASI_TWINK_REAL_L)
27 ₁₆ ‡ (ASI_TWINK_N)	2F ₁₆ ‡ (ASI_TWINK_NL)
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE
82 ₁₆ ASI_PRIMARY_NO_FAULT	8A ₁₆ ASI_PRIMARY_NO_FAULT_LITTLE
83 ₁₆ ASI_SECONDARY_NO_FAULT	8B ₁₆ ASI_SECONDARY_NO_FAULT_LITTLE
E2 ₁₆ ‡ (ASI_TWINK_P)	EA ₁₆ ‡ (ASI_TWINK_PL)
E3 ₁₆ ‡ (ASI_TWINK_S)	EB ₁₆ ‡ (ASI_TWINK_SL)

‡ If this ASI is used with the opcode for LDTWA and i = 0, the LDTXA instruction is executed instead of LDTWA. For behavior of LDTXA, see *Load Integer Twin Extended Word from Alternate Space* on page 262.
 If this ASI is used with the opcode for LDTWA and i = 1, a *DAE_invalid_asi* exception occurs.

Programming Note | Nontranslating ASIs (see page 423) should only be accessed using LDXA (not LDTWA) instructions. If an LDTWA referencing a nontranslating ASI is executed, per the above table, it generates a *DAE_invalid_asi* exception (impl. dep. #300-U4-Cs10).

LDTWA (Deprecated)

Implementation Note | The deprecated instruction LDTWA shares an opcode with LDTXA. LDTXA is *not* deprecated and has different address alignment requirements than LDTWA. See *Load Integer Twin Extended Word from Alternate Space* on page 262.

Exceptions *unimplemented_LDTW* (not used in Oracle SPARC Architecture 2015)
illegal_instruction
mem_address_not_aligned
privileged_action
VA_watchpoint
DAE_invalid_asi
DAE_privilege_violation
DAE_nfo_page
DAE_side_effect_page

See Also LDWA/LDXA on page 240
LDTXA on page 262
STTWA on page 354

LDTXA

7.86 Load Integer Twin Extended Word from Alternate Space VIS 2+

The LDTXA instructions are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	ASI Value	Operation	Assembly Language Syntax †	Class
LDTXA ^N	22 ₁₆	Load Integer Twin Extended Word, as if user (nonprivileged), Primary address space	<code>ldtxa [regaddr] #ASI_TWIX_AIUP, regrd</code>	N-
	23 ₁₆	Load Integer Twin Extended Word, as if user (nonprivileged), Secondary address space	<code>ldtxa [regaddr] #ASI_TWIX_AIUS, regrd</code>	N-
	26 ₁₆	Load Integer Twin Extended Word, real address	<code>ldtxa [regaddr] #ASI_TWIX_REAL, regrd</code>	N-
	27 ₁₆	Load Integer Twin Extended Word, nucleus context	<code>ldtxa [regaddr] #ASI_TWIX_N, regrd</code>	N-
	2A ₁₆	Load Integer Twin Extended Word, as if user (nonprivileged), Primary address space, little endian	<code>ldtxa [regaddr] #ASI_TWIX_AIUP_L, regrd</code>	N-
	2B ₁₆	Load Integer Twin Extended Word, as if user (nonprivileged), Secondary address space, little endian	<code>ldtxa [regaddr] #ASI_TWIX_AIUS_L, regrd</code>	N-
	2E ₁₆	Load Integer Twin Extended Word, real address, little endian	<code>ldtxa [regaddr] #ASI_TWIX_REAL_L, regrd</code>	N-
	2F ₁₆	Load Integer Twin Extended Word, nucleus context, little-endian	<code>ldtxa [regaddr] #ASI_TWIX_NL, regrd</code>	N-
LDTXA ^N	E2 ₁₆	Load Integer Twin Extended Word, Primary address space	<code>ldtxa [regaddr] #ASI_TWIX_P, regrd</code>	N-
	E3 ₁₆	Load Integer Twin Extended Word, Secondary address space	<code>ldtxa [regaddr] #ASI_TWIX_S, regrd</code>	N-
	EA ₁₆	Load Integer Twin Extended Word, Primary address space, little endian	<code>ldtxa [regaddr] #ASI_TWIX_PL, regrd</code>	N-
	EB ₁₆	Load Integer Twin Extended Word, Secondary address space, little-endian	<code>ldtxa [regaddr] #ASI_TWIX_SL, regrd</code>	N-

† The original assembly language syntax for these instructions used the “`ldda`” instruction mnemonic. That syntax is now deprecated. Over time, assemblers will support the new “`ldtxa`” mnemonic for this instruction. In the meantime, some existing assemblers may only recognize the original “`ldda`” mnemonic.



Description ASIs 26₁₆, 2E₁₆, E2₁₆, E3₁₆, EA₁₆, and EB₁₆ are used with the LDTXA instruction to atomically read a 128-bit data item into a pair of 64-bit R registers (a “twin extended word”). The data are placed in an even/odd pair of 64-bit registers. The lowest-address 64 bits are placed in the even-numbered register; the highest-address 64 bits are placed in the odd-numbered register.

Note | Execution of an LDTXA instruction with `rd = 0` modifies only R[1].

LDTXA

ASIs E2₁₆, E3₁₆, EA₁₆, and EB₁₆ perform an access using a virtual address, while ASIs 26₁₆ and 2E₁₆ use a real address.

An LDTXA instruction that performs a little-endian access behaves as if it comprises two 64-bit loads (performed atomically), each of which is byte-swapped independently before being written into its respective destination register.

Exceptions. An attempt to execute an LDTXA instruction with an odd-numbered destination register ($rd\{0\} = 1$) causes an *illegal_instruction* exception.

An attempt to execute an LDTXA instruction with an effective memory address that is not aligned on a 16-byte boundary causes a *mem_address_not_aligned* exception.

IMPL. DEP. #413-S10: It is implementation dependent whether *VA_watchpoint* exceptions are recognized on accesses to all 16 bytes of a LDTXA instruction (the recommended behavior) or only on accesses to the first 8 bytes.

An attempted access by an LDTXA instruction to noncacheable memory causes an a *DAE_nc_page* exception (impl. dep. #306-U4-Cs10).

Programming Note	A key use for this instruction is to read a full TTE entry (128 bits, tag and data) in a TSB directly, without using software interlocks. The “real address” variants can perform the access using a real address, bypassing the VA-to-RA translation.
-------------------------	--

The virtual processor MMU does not provide virtual-to-real translation for ASIs 26₁₆ and 2E₁₆; the effective address provided with either of those ASIs is interpreted directly as a real address.

Compatibility Note	ASIs 27 ₁₆ , 2F ₁₆ , 26 ₁₆ , and 2E ₁₆ are now standard ASIs that replace (respectively) ASIs 24 ₁₆ , 2C ₁₆ , 34 ₁₆ , and 3C ₁₆ that were supported in some previous UltraSPARC implementations.
---------------------------	--

Implementation Note	LDTXA shares an opcode with the “i = 0” variant of the (deprecated) LDTWA instruction; they are differentiated by the combination of the i instruction field and the ASI used in the instruction. See <i>Load Integer Twin Word from Alternate Space</i> on page 259.
----------------------------	---

<i>Exceptions</i>	<i>illegal_instruction</i> <i>mem_address_not_aligned</i> <i>privileged_action</i> <i>VA_watchpoint</i> (impl. dep. #413-S10) <i>DAE_nc_page</i> <i>DAE_nfo_page</i>
-------------------	---

<i>See Also</i>	LDTWA on page 259
-----------------	-------------------

7.87 Load Floating-Point State Register

Instruction	op3	rd	Operation	Assembly Language Syntax	Class	Added
	10 0001	0	(see page 251)			
LDXFSR	10 0001	1	Load Floating-Point State Register	ldx [address], %fsr	A1	
—	10 0001	2	Reserved			
LDXEFSR	10 0001	3	Load Entire Floating-Point State Register VIS 3B	ldx [address], %efsr	C1	OSA 2011
—	10 0001	4–31	Reserved			



Description A load floating-point state register instruction (LDXFSR or LDXEFSR) waits for all FPop instructions that have not finished execution to complete and then loads a doubleword from memory into the FSR.

LDXFSR does not alter the *ver*, *ftt*, *qne*, reserved, or unimplemented (for example, *ns*) fields of FSR (see page 42).

An LDXEFSR instruction loads from memory the “entire” FSR, including FSR.ftt. However, it does not alter the *ver*, *qne*, reserved, or unimplemented (for example, *ns*) fields of FSR; writes to those fields are ignored.

Programming Note For future compatibility, software should only issue an LDXFSR or LDXEFSR instruction with a zero value (or a value previously read from the same field) written into any reserved field of FSR.

LDXFSR and LDXEFSR access memory using the implicit ASI (see page 83).

If *i* = 0, the effective address for these instructions is “R[rs1] + R[rs2]” and if *i* = 1, the effective address is “R[rs1] + *sign_ext*(simm13)”.

Exceptions. An attempt to execute an instruction encoded as *op* = 2 and *op3* = 21₁₆ when any of the following conditions exist causes an *illegal_instruction* exception:

- *i* = 0 and instruction bits 12:5 are nonzero
- (*rd* = 2 or *rd* ≠ 4)

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute an LDXFSR or LDXEFSR instruction causes an *fp_disabled* exception.

If the effective address is not doubleword-aligned, an attempt to execute an LDXFSR or LDXEFSR instruction causes a *mem_address_not_aligned* exception.

Destination Register(s) when Exception Occurs. If a load floating-point state register instruction generates an exception that causes a *precise* trap, the destination register (FSR) remains unchanged.

IMPL. DEP. #44-V8-Cs10(a)(2): If an LDXFSR or LDXEFSR instruction generates an exception that causes a *non-precise* trap, it is implementation dependent whether the contents of the destination register (FSR) is undefined or is guaranteed to remain unchanged.

LDXFSR / LDXFSR

Implementation Note LDXFSR and LDXEFSR share an opcode with the (deprecated) LDFSR instruction (and possibly with other implementation-dependent instructions); they are differentiated by the instruction rd field. An attempt to execute the $op = 11_2$, $op3 = 10\ 0001_2$ opcode with an invalid rd value causes an *illegal_instruction* exception.

Exceptions

- illegal_instruction*
- fp_disabled*
- mem_address_not_aligned*
- VA_watchpoint*
- DAE_privilege_violation*
- DAE_nfo_page*

See Also

- Load Floating-Point Register on page 246*
- Load Floating-Point State Register (Lower) on page 251*
- Store Floating-Point State Register on page 356*


LZCNT

7.88 Leading Zeroes Count VIS 3

The LZCNT instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
LZCNT	0 0001 0111	Leading zeroes count (from left end)	—	i64	i64	lzcmt <i>reg_{rs2}</i> , <i>reg_{rd}</i> (or lzd ‡)	C1	OSA 2011

‡ The original assembly language syntax for this instruction used an “lzd” instruction mnemonic, which is now deprecated. Over time, assemblers will support the new “lzcmt” mnemonic for this instruction. In the meantime, some assemblers may only recognize the original “lzd” mnemonic.



Description The number of leading zeros (0 to 64) in R[rs2] is counted and written as the value *n* in the least significant 7 bits of the destination register, R[rd]. The most significant 57 bits of R[rd] bits are written with zeros.




FIGURE 7-58 Leading Zeroes Count

An attempt to execute an LZCNT instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

Programming Note LZCNT can also be used for “leading 1” detection, by taking the 1’s complement of the source operand before executing the LZCNT instruction.

Programming Note See the Programming Note on page 301, regarding how to a “Trailing Zeroes Count” operation can be synthesized using POPC.

Historical Note LZCNT’s original name was “LZD” for “Leading Zero Detect”. That was a misnomer, so it was renamed LZCNT. Software tools will continue recognizing the original assembly-language mnemonic.


Exceptions *illegal_instruction*

■ *See Also* *Population Count on page 301*

7.89 MD5 Hash Operation Crypto

The Hash instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class
MD5	1 0100 0000	MD5 operation on a single block	md5	N1



Description The MD5 instruction implements the MD5 algorithm, which is available as RFC-1321 at <http://nist.gov>.

The hash instructions operate on 64-bit floating-point registers and process an entire 512-bit block. The locations of the Initialization Vector (IV), Data, and Result in the floating-point register file are described below. To compute the hash over multiple blocks, the result from the previous hash instruction is used as the IV for the next block. Software must appropriately pad the final block as specified by the given algorithm.

MD5 : IV{127:0} = (F_D[0] :: F_D[2])
 Data{511:0} = (F_D[8] :: F_D[10] :: F_D[12] :: F_D[14] :: F_D[16] :: F_D[18] :: F_D[20] :: F_D[22])
 Result{127:0} = (F_D[0] :: F_D[2])
 Result{511:0} = (F_D[0] :: F_D[2] :: F_D[4] :: F_D[6] :: F_D[8] :: F_D[10] :: F_D[12] :: F_D[14])

If $rd \neq 0$, $rs1 \neq 0$, or $rs2 \neq 0$, an attempt to execute an MD5 instruction causes an *illegal_instruction* exception.

If $CFR.md5 = 0$, an attempt to execute an MD5 instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that $CFR.md5 = 1$ before executing the MD5 instruction. If $CFR.md5 = 0$, then software should assume that an attempt to execute the MD5 instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.


Therefore, if $CFR.md5 = 0$, software should perform the MD5 operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*

MEMBAR

7.90 Memory Barrier

Instruction	op3	Operation	Assembly Language Syntax	Class
MEMBAR	10 1000	Memory Barrier	membar <i>membar_mask</i>	A1



Description The memory barrier instruction, MEMBAR, has two complementary functions: to express order constraints between memory references and to provide explicit control of memory-reference completion. The *membar_mask* field in the suggested assembly language is the concatenation of the *cmask* and *mmask* instruction fields.

MEMBAR introduces an order constraint between classes of memory references appearing before the MEMBAR and memory references following it in a program. The particular classes of memory references are specified by the *mmask* field. Memory references are classified as loads (including instructions LDSTUB[A], SWAP[A], CASA, and CASX[A] and stores (including instructions LDSTUB[A], SWAP[A], CASA, CASXA, and FLUSH). The *mmask* field specifies the classes of memory references subject to ordering, as described below. MEMBAR applies to all memory operations in all address spaces referenced by the issuing virtual processor, but it has no effect on memory references by other virtual processors. When the *cmask* field is nonzero, completion as well as order constraints are imposed, and the order imposed can be more stringent than that specifiable by the *mmask* field alone.

A load has been performed when the value loaded has been transmitted from memory and cannot be modified by another virtual processor. A store has been performed when the value stored has become visible, that is, when the previous value can no longer be read by any virtual processor. In specifying the effect of MEMBAR, instructions are considered to be executed as if they were processed in a strictly sequential fashion, with each instruction completed before the next has begun.

The *mmask* field is encoded in bits 3 through 0 of the instruction. TABLE 7-13 specifies the order constraint that each bit of *mmask* (selected when set to 1) imposes on memory references appearing before and after the MEMBAR. From zero to four mask bits may be selected in the *mmask* field.

TABLE 7-13 MEMBAR *mmask* Encodings

Mask Bit	Assembly Language Name	Description
mmask{3}	#StoreStore	The effects of all stores appearing prior to the MEMBAR instruction must be visible to all virtual processors before the effect of any stores following the MEMBAR.
mmask{2}	#LoadStore	All loads appearing prior to the MEMBAR instruction must have been performed before the effects of any stores following the MEMBAR are visible to any other virtual processor.
mmask{1}	#StoreLoad	The effects of all stores appearing prior to the MEMBAR instruction must be visible to all virtual processors before loads following the MEMBAR may be performed.
mmask{0}	#LoadLoad	All loads appearing prior to the MEMBAR instruction must have been performed before any loads following the MEMBAR may be performed.

MEMBAR

The `cmask` field is encoded in bits 6 through 4 of the instruction. Bits in the `cmask` field, described in TABLE 7-14, specify additional constraints on the order of memory references and the processing of instructions. If `cmask` is zero, then MEMBAR enforces the partial ordering specified by the `mmask` field; if `cmask` is nonzero, then completion and partial order constraints are applied.

TABLE 7-14 MEMBAR `cmask` Encodings

Mask Bit	Function	Assembly Language Name	Description
<code>cmask{2}</code>	Synchronization barrier	<code>#Sync</code>	All operations (including nonmemory reference operations) appearing prior to the MEMBAR must have been performed and the effects of any exceptions must be visible before any instruction after the MEMBAR may be initiated.
<code>cmask{1}</code>	Memory issue barrier	<code>#MemIssue</code>	All memory reference operations appearing prior to the MEMBAR must have been performed before any memory operation after the MEMBAR may be initiated.
<code>cmask{0}</code>	Lookaside barrier	<code>#Lookaside^D</code>	(Deprecated - software should <i>not</i> use) A store appearing prior to the MEMBAR must complete before any load following the MEMBAR referencing the same address can be initiated. MEMBAR <code>#Lookaside</code> is deprecated and is supported <i>only</i> for legacy code; it should <i>not</i> be used in new software. A slightly more restrictive MEMBAR operation (such as MEMBAR <code>#StoreLoad</code>) should be used, instead. Implementation Note: Since <code>#Lookaside</code> is deprecated, implementations are not expected to perform address matching, but instead provide <code>#Lookaside</code> functionality using a more restrictive MEMBAR operation (such as <code>#StoreLoad</code>).

A MEMBAR instruction with both `mmask = 0` and `cmask = 0` is functionally a NOP.

For information on the use of MEMBAR, see . For additional information about the memory models themselves, see Chapter 9, *Memory*.

The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent (impl. dep. #120-V9).

V9 Compatibility Note | MEMBAR with `mmask = 816` and `cmask = 016` (MEMBAR `#StoreStore`) is identical in function to the SPARC V8 STBAR instruction, which is deprecated.

An attempt to execute a MEMBAR instruction when instruction bits 12:7 are nonzero causes an *illegal_instruction* exception.

Implementation Note | MEMBAR shares an opcode with RDasr #15; MEMBAR is distinguished by `rs1 = 15`, `rd = 0`, `i = 1`, and bit 12 = 0.

7.90.1 Memory Synchronization

The Oracle SPARC Architecture provides some level of software control over memory synchronization, through use of the MEMBAR and FLUSH instructions for explicit control of memory ordering in program execution.

IMPL. DEP. #412-S10: An Oracle SPARC Architecture implementation may define the operation of each MEMBAR variant in any manner that provides the required semantics.

MEMBAR

Implementation Note For an Oracle SPARC Architecture virtual processor that only provides TSO memory ordering semantics, three of the ordering MEMBARs would normally be implemented as NOPs. TABLE 7-15 shows an acceptable implementation of MEMBAR for a TSO-only Oracle SPARC Architecture implementation.

TABLE 7-15 MEMBAR Semantics for TSO-only implementation

MEMBAR variant	Preferred Implementation
#StoreStore	NOP
#LoadStore	NOP
#StoreLoad	#Sync
#LoadLoad	NOP
#Sync	#Sync
#MemIssue	#Sync
#Lookaside ^D	#Sync

If an Oracle SPARC Architecture implementation provides a less restrictive memory model than TSO (for example, RMO), the implementation of the MEMBAR variants may be different. See implementation-specific documentation for details.

7.90.2 Synchronization of the Virtual Processor

Synchronization of a virtual processor forces all outstanding instructions to be completed and any associated hardware errors to be detected and reported before any instruction after the synchronizing instruction is issued.

Synchronization can be explicitly caused by executing a synchronizing MEMBAR instruction (MEMBAR #Sync) or by executing an LDXA/STXA/LDDFA/STDFA instruction with an ASI that forces synchronization.

Programming Note Completion of a MEMBAR #Sync instruction does *not* guarantee that data previously stored has been written all the way out to external memory. Software cannot rely on that behavior. There is no mechanism in the Oracle SPARC Architecture that allows software to wait for all previous stores to be written to external memory.

7.90.3 TSO Ordering Rules affecting Use of MEMBAR

For detailed rules on use of MEMBAR to enable software to adhere to the ordering rules on a virtual processor running with the TSO memory model, refer to *TSO Ordering Rules* on page 417.

Exceptions *illegal_instruction*

7.91 MONTMUL Crypto

This instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax		Class
MONTMUL ^N	1 0100 1001	Montgomery Multiplication	montmul	<i>imm5</i>	N1



Description MONTMUL performs the Montgomery Multiplication shown below with length equal to *imm5*. The starting locations of the operands are constant in the integer register file (IRF), regardless of size. The starting location of the result is also constant. Below are the locations for N=31. For smaller MONTMULs, the remaining operand locations are not used and the remaining result locations will be unchanged.

FSR.fcc3 is set upon completion of MONTMUL to reflect whether or not a hardware error occurred during execution (see the Programming Note below).

The following pseudo-code specifies the operation of the MONTMUL instruction:

```

MontMul (l_uint A, l_uint B, l_uint N, l_uint *M, l_uint Nprime,
        char Length, l_uint *X) {
    // compute Montgomery Multiplication as described below
    // use M as temporary variable
    // return result in X
    // A,B,N,M,X 64*(Length+1) bit long
    // Nprime is 64 bits long

    ACCUM ← 0
    for I←0 to Length begin           // Length is one less than the number of words
        for j←0 to I-1 begin         //skipped on first I iteration
            ACCUM ← ACCUM + (A[j]×B[I-j])
            ACCUM ← ACCUM + (M[j]×N[I-j])
        end
        ACCUM ← ACCUM + (A[I]×B[0])
        M[I] ← ACCUM × Nprime        // 64 LSB of accum, store 64 LSB of product
        ACCUM ← ACCUM + (M[I]×N[0])
        ACCUM ← ACCUM >> 64
    end

    for I ← (Length+1) to ((2×Length)+1) begin
        for j ← (I-Length) to Length begin           // skip last I iteration
            ACCUM ← ACCUM + (A[j]×B[I-j])
            ACCUM ← ACCUM + (M[j]×N[I-j])
        end
        X[I-Length-1] ← ACCUM                    // 64 LSB of accum
        ACCUM ← ACCUM >> 64
    end

    ModReduction(ACCUM, X, N, Length)           // Reduce the final result
}
    
```


MONTMUL

```

ModReduction (bit ACCUM, l_uint *X, l_uint N, char Length) {
    // compute (ACCUM|X) mod N
    // return result in X
    // ACCUM 1 bit long
    // N,X 64×(Length+1) bit long

    if (ACCUM ≠ 0) begin
        for I ← 0 to Length begin // Length is one less than the number of words
            X[I] ← X[I] - N[I] // Subtraction with borrow
        end
    end
    else begin
        I ← Length
        while ((I≥0) and (X[I] = N[I])) begin
            I ← I-1
        end
        if ((I<0) or (X[I]>N[I])) begin
            for I ← 0 to Length begin
                X[I] ← X[I] - N[I] // Subtraction with borrow
            end
        end
    end
end
}

```

MONTMUL Operand Locations

Operand	Win- dow (CWP value)	Register(s)	
<i>Source operands</i>			
Nprime	—	F _D [60]	
M[7:0]	i-6	(F _D [2] :: F _D [0] :: R[29] :: R[28] :: R[27] :: R[26] :: R[25] :: R[24])	f2,f0,i5...i0
M[15:8]	i-6	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
M[23:16]	i-6	(F _D [6] :: F _D [4] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f6,f4,o5...o0
M[31:24]	—	(F _D [22] :: F _D [20] :: F _D [18] :: F _D [16] :: F _D [14] :: F _D [12] :: F _D [10] :: F _D [8])	
A[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
A[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24, o5...o0
A[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
A[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	
N[7:0]	i-4	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
N[13:8]	i-4	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[21:14]	i-3	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
N[27:22]	i-3	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[31:28]	i-2	(R[19] :: R[18] :: R[17] :: R[16])	13...10
B[5:0]	i-2	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
B[13:6]	i-1	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
B[19:14]	i-1	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
B[27:20]	i	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
B[31:28]	i	(R[11] :: R[10] :: R[9] :: R[8])	o3...o0

MONTMUL

MONTMUL Operand Locations

Operand	Win- dow (CWP value)	Register(s)	
<i>Destination operands</i>			
X[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
X[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24,o5...o0
X[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
X[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	

Programming Note | The MONTMUL instruction uses seven windows of the integer register file (IRF).

Programming Note | Due to the special nature of the MONTMUL instruction, the processor cannot protect the programmer from hardware errors in integer and floating-point register file locations during execution of MONTMUL. The MONTMUL instruction sets FSR.fcc3 to 00₂ if no hardware error occurred, or to 11₂ (unordered) if an error did occur. It is the responsibility of the programmer to check FSR.fcc3 when the instruction completes and to retry the instruction if it encountered a hardware error. The programmer should provide an error counter to allow for a limited number of attempts to retry the operation. The instruction may be retried after first reloading all the input data from a backing storage location; thus, the programmer must take care to preserve a clean copy of the input data.

Exceptions. If rs1 ≠ 0, an attempt to execute a MONTMUL instruction causes an *illegal_instruction* exception.

If CFR.montmul = 0, an attempt to execute a MONTMUL instruction causes a *compatibility_feature* exception.

Programming Note | Software *must* check that CFR.montmul = 1 before executing the MONTMUL instruction. If CFR.montmul = 0, then software should assume that an attempt to execute the MONTMUL instruction either
 (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
 (2) will execute but perform some other operation.
 Therefore, if CFR.montmul = 0, software should perform the MONTMUL operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a MONTMUL instruction causes an *fp_disabled* exception.

The MONTMUL instruction causes a *fill_n_normal* or *fill_n_other* exception if CANRESTORE is not equal to NWINDOWS-2. The fill trap handler is called with CWP set to point to the window to be filled, that is, old CWP-CANRESTORE-1. The trap vector for the fill trap is based on the values of OTHERWIN and WSTATE, as described in Trap Type for Spill/Fill Traps on page 411. The fill trap handler performs a RESTORED and a RETRY as part of filling the window. When the virtual processor reexecutes the MONTMUL (due to the handler ending in RETRY), another fill trap will result if more than one window needed to be filled.

MONTMUL

Implementation Note | MONTMUL and XMONTMUL share a basic opcode (op and opf). They are differentiated by the instruction rd field; rd = 0 0000₂ for MONTMUL and rd = 0 0001₂ for XMONTMUL.

Exceptions *fp_disabled*
 fill_n_normal (n = 0-7)
 fill_n_other (n = 0-7)

See Also MONTSQR on page 276
 MPMUL on page 286
 XMONTMUL on page 378
 XMONTSQR on page 381
 XMPMUL on page 384

MONTSQR

7.92 MONTSQR Crypto

This instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class
MONTSQR ^N	1 0100 1010	Montgomery Squaring	montsqr <i>imm5</i>	N1



Description MONTSQR performs the Montgomery Squaring shown below with length equal to imm5. The starting locations of the operands are constant in the integer register file (IRF) regardless of size. The starting location of the result is also constant. Below are the locations for N=31. For smaller MONTSQRs, the remaining operand locations are not used and the remaining result locations will be unchanged.

FSR.fcc3 is set upon completion of MONTSQR to reflect whether or not a hardware error occurred during execution (see the Programming Note below).

```

MontSqr (l_uint A, l_uint N, l_uint *M, l_uint Nprime, char Length, l_uint *X) {
    // compute Montgomery Squaring as described below
    // use M as temporary variable
    // return result in X
    // A,N,M,X 64×(Length+1) bit long
    // Nprime is 64 bits long

    ACCUM ← 0
    for I ← 0 to Length begin           // Length is one less than the number of
        for j ← 0 to ((I-1) >> 1) begin // words skipped on first I iteration
            ACCUM ← ACCUM + (2×A[j]×A[I-j])
        end
        if I is even begin
            ACCUM ← ACCUM + A[I/2]^2
        end
        for j ← 0 to I-1 begin
            ACCUM ← ACCUM + M[j] × N[I-j]
        end
        M[I] ← ACCUM × Nprime           // 64 LSB of accum, store 64 LSB of product
        ACCUM ← ACCUM + (M[I]×N[0])
        ACCUM ← (ACCUM >> 64)
    end

    for I ← (Length+1) to ((2×Length)+1) begin
        for j ← (I-Length) to ((I-1) >> 1) begin           // skip last I iteration
            ACCUM ← ACCUM + (2×A[j]×A[I-j])
        end
        if I is even begin
            ACCUM ← ACCUM + A[I/2]^2
        end
        for j = (I-Length) to Length begin
            ACCUM ← ACCUM + (M[j]×N[I-j])
        end
    end
}

```

MONTSQR

```

    end
    X[I-Length-1] ← ACCUM          // 64 LSB of accum
    ACCUM ← (ACCUM >> 64)
end

ModReduction(ACCUM, X, N, Length) // Reduce the final result
}

ModReduction (bit ACCUM, l_uint *X, l_uint N, char Length) {
    // compute (ACCUM|X) mod N
    // return result in X
    // ACCUM 1 bit long
    // N,X 64×(Length+1) bit long

    if (ACCUM ≠ 0) begin
        for I ← 0 to Length begin // Length is one less than the number of words
            X[I] ← X[I] - N[I] // Subtraction with borrow
        end
    end
    else begin
        I ← Length
        while ((I ≥ 0) and (X[I] = N[I])) begin
            I ← I-1
        end
        if ((I < 0) or (X[I] > N[I])) begin
            for I ← 0 to Length begin
                X[I] ← X[I] - N[I] // Subtraction with borrow
            end
        end
    end
end
}

```

MONTSQR Operand Location

Operand	Window (CWP value)	Register(s)	Notes
<i>Source operands</i>			
Nprime	—	F _D [60]	
M[7:0]	i-6	(F _D [2] :: F _D [0] :: R[29] :: R[28] :: R[27] :: R[26] :: R[25] :: R[24])	f2,f0,i5...i0
M[15:8]	i-6	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
M[23:16]	i-6	(F _D [6] :: F _D [4] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f6,f4,o5...o0
M[31:24]	—	(F _D [22] :: F _D [20] :: F _D [18] :: F _D [16] :: F _D [14] :: F _D [12] :: F _D [10] :: F _D [8])	
A[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
A[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24,o5...o0
A[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
A[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	
N[7:0]	i-4	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
N[13:8]	i-4	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[21:14]	i-3	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
N[27:22]	i-3	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[31:28]	i-2	(R[19] :: R[18] :: R[17] :: R[16])	13...10

MONTSQR

MONTSQR Operand Location

Operand	Win- dow (CWP value)	Register(s)	Notes
<i>Destination operands</i>			
X[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
X[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24,o5...o0
X[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
X[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	

Programming Note | The MONTSQR instruction uses seven windows of the integer register file (IRF).

Programming Note | Due to the special nature of the MONTSQR instruction, the processor cannot protect the programmer from hardware errors in integer and floating-point register file locations during execution of MONTSQR. The MONTSQR instruction sets FSR.fcc3 to 00₂ if no hardware error occurred, or to 11₂ (unordered) if an error did occur. It is the responsibility of the programmer to check FSR.fcc3 when the instruction completes and to retry the instruction if it encountered a hardware error. The programmer should provide an error counter to allow for a limited number of attempts to retry the operation. The instruction may be retried after first reloading all the input data from a backing storage location; thus, the programmer must take care to preserve a clean copy of the input data.

Exceptions. If rs1 ≠ 0, an attempt to execute a MONTSQR instruction causes an *illegal_instruction* exception.

If CFR.montsqr = 0, an attempt to execute a MONTSQR instruction causes a *compatibility_feature* exception.

Programming Note | Software *must* check that CFR.montsqr = 1 before executing the MONTSQR instruction. If CFR.montsqr = 0, then software should assume that an attempt to execute the MONTSQR instruction either
 (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
 (2) will execute but perform some other operation.
 Therefore, if CFR.montsqr = 0, software should perform the MONTSQR operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a MONTSQR instruction causes an *fp_disabled* exception.

The MONTSQR instruction causes a *fill_n_normal* or *fill_n_other* exception if CANRESTORE is not equal to NWINDOWS-2. The fill trap handler is called with CWP set to point to the window to be filled, that is, old CWP-CANRESTORE-1. The trap vector for the fill trap is based on the values of OTHERWIN and WSTATE, as described in Trap Type for Spill/Fill Traps on page 411. The fill trap handler performs a RESTORED and a RETRY as part of filling the window. When the virtual processor reexecutes the MONTSQR (due to the handler ending in RETRY), another fill trap will result if more than one window needed to be filled.

MONTSQR

Implementation | MONTSQR and XMONTSQR share a basic opcode (op and opf).
Note | They are differentiated by the instruction rd field; rd = 0 0000₂ for MONTSQR and rd = 0 0001₂ for XMONTSQR.

Exceptions *fp_disabled*
 fill_n_normal (n = 0-7)
 fill_n_other (n = 0-7)

See Also MONTMUL on page 272
 MPMUL on page 286
 XMONTMUL on page 378
 XMONTSQR on page 381
 XMPMUL on page 384

MOVcc

7.93 Move Integer Register on Condition (MOVcc)

For Integer Condition Codes

Instruction	op3	cond	Operation	icc / xcc Test	Assembly Language Syntax	Class
MOVA	10 1100	1000	Move Always	1	<code>mova i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVN	10 1100	0000	Move Never	0	<code>movn i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVNE	10 1100	1001	Move if Not Equal	not Z	<code>movne[†] i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVE	10 1100	0001	Move if Equal	Z	<code>move[‡] i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVG	10 1100	1010	Move if Greater	not (Z or N xor V)	<code>movg i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVLE	10 1100	0010	Move if Less or Equal	Z or (N xor V)	<code>movle i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVGE	10 1100	1011	Move if Greater or Equal	not (N xor V)	<code>movge i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVL	10 1100	0011	Move if Less	N xor V	<code>movl i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVGU	10 1100	1100	Move if Greater, Unsigned	not (C or Z)	<code>movgu i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVLEU	10 1100	0100	Move if Less or Equal, Unsigned	(C or Z)	<code>movleu i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVCC	10 1100	1101	Move if Carry Clear (Greater or Equal, Unsigned)	not C	<code>movcc[◇] i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVCS	10 1100	0101	Move if Carry Set (Less than, Unsigned)	C	<code>movcs[∇] i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVPOS	10 1100	1110	Move if Positive	not N	<code>movpos i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVNEG	10 1100	0110	Move if Negative	N	<code>movneg i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVVC	10 1100	1111	Move if Overflow Clear	not V	<code>movvc i_or_x_cc, reg_or_imm11, reg_rd</code>	A1
MOVVS	10 1100	0111	Move if Overflow Set	V	<code>movvs i_or_x_cc, reg_or_imm11, reg_rd</code>	A1

[†] *synonym: movnz*

[‡] *synonym: movz*

[◇] *synonym: movgeu*

[∇] *synonym: movlu*

Programming Note | In assembly language, to select the appropriate condition code, include `%icc` or `%xcc` before the `reg_or_imm11` field.


MOVcc

For Floating-Point Condition Codes

Instruction	op3	cond	Operation	fcc Test	Assembly Language Syntax	Class
MOVFA	10 1100	1000	Move Always	1	mova %fccn, reg_or_imm11, reg _{rd}	A1
MOVFN	10 1100	0000	Move Never	0	movn %fccn, reg_or_imm11, reg _{rd}	A1
MOVFU	10 1100	0111	Move if Unordered	U	movu %fccn, reg_or_imm11, reg _{rd}	A1
MOVFG	10 1100	0110	Move if Greater	G	movg %fccn, reg_or_imm11, reg _{rd}	A1
MOVFUG	10 1100	0101	Move if Unordered or Greater	G or U	movug %fccn, reg_or_imm11, reg _{rd}	A1
MOVFL	10 1100	0100	Move if Less	L	movl %fccn, reg_or_imm11, reg _{rd}	A1
MOVFUL	10 1100	0011	Move if Unordered or Less	L or U	movul %fccn, reg_or_imm11, reg _{rd}	A1
MOVFLG	10 1100	0010	Move if Less or Greater	L or G	movlg %fccn, reg_or_imm11, reg _{rd}	A1
MOVFNE	10 1100	0001	Move if Not Equal	L or G or U	movne [†] %fccn, reg_or_imm11, reg _{rd}	A1
MOVFE	10 1100	1001	Move if Equal	E	move [‡] %fccn, reg_or_imm11, reg _{rd}	A1
MOVFUE	10 1100	1010	Move if Unordered or Equal	E or U	movue %fccn, reg_or_imm11, reg _{rd}	A1
MOVFGE	10 1100	1011	Move if Greater or Equal	E or G	movge %fccn, reg_or_imm11, reg _{rd}	A1
MOVFUGE	10 1100	1100	Move if Unordered or Greater or Equal	E or G or U	movuge %fccn, reg_or_imm11, reg _{rd}	A1
MOVFLE	10 1100	1101	Move if Less or Equal	E or L	movle %fccn, reg_or_imm11, reg _{rd}	A1
MOVFULE	10 1100	1110	Move if Unordered or Less or Equal	E or L or U	movule %fccn, reg_or_imm11, reg _{rd}	A1
MOVFO	10 1100	1111	Move if Ordered	E or L or G	movo %fccn, reg_or_imm11, reg _{rd}	A1

[†] synonym: movnz [‡] synonym: movz

Programming Note In assembly language, to select the appropriate condition code, include %fcc0, %fcc1, %fcc2, or %fcc3 before the *reg_or_imm11* field.



cc2	cc1	cc0	Condition Code
0	0	0	fcc0
0	0	1	fcc1
0	1	0	fcc2
0	1	1	fcc3
1	0	0	icc
1	0	1	Reserved (illegal_instruction)
1	1	0	xcc
1	1	1	Reserved (illegal_instruction)

MOVcc

Description These instructions test to see if `cond` is `TRUE` for the selected condition codes. If so, they copy the value in `R[rs2]` if `i` field = 0, or “`sign_ext(simm11)`” if `i` = 1 into `R[rd]`. The condition code used is specified by the `cc2`, `cc1`, and `cc0` fields of the instruction. If the condition is `FALSE`, then `R[rd]` is not changed.

These instructions copy an integer register to another integer register if the condition is `TRUE`. The condition code that is used to determine whether the move will occur can be either integer condition code (`icc` or `xcc`) or any floating-point condition code (`fcc0`, `fcc1`, `fcc2`, or `fcc3`).

These instructions do not modify any condition codes.

Programming Note Branches cause the performance of many implementations to degrade significantly. Frequently, the `MOVcc` and `FMOVcc` instructions can be used to avoid branches. For example, the C language if-then-else statement

```
if (A > B) then X = 1; else X = 0;
```

can be coded as

```
    cmp    %i0,%i2
    bg,a   %xcc,label
    or     %g0,1,%i3! X = 1
    or     %g0,0,%i3! X = 0
label:...
```

The above sequence requires four instructions, including a branch. With `MOVcc` this could be coded as:

```
    cmp    %i0,%i2
    or     %g0,1,%i3! assume X = 1
    movle  %xcc,0,%i3! overwrite with X = 0
```

This approach takes only three instructions and no branches and may boost performance significantly. Use `MOVcc` and `FMOVcc` instead of branches wherever these instructions would increase performance.

An attempt to execute a `MOVcc` instruction when either instruction bits 10:5 are nonzero or $(cc2 :: cc1 :: cc0) = 101_2$ or 111_2 causes an *illegal_instruction* exception.

If `cc2` = 0 (that is, a floating-point condition code is being referenced in the `MOVcc` instructions) and either the FPU is not enabled (`FPRS.fef` = 0 or `PSTATE.pef` = 0) or if no FPU is present, an attempt to execute a `MOVcc` instruction causes an *fp_disabled* exception.


Exceptions *illegal_instruction*
fp_disabled

MOVr

7.94 Move Integer Register on Register Condition (MOVr)

Instruction	op3	rcond	Operation	Test	Assembly Language Syntax	Class
—	10 1111	000	<i>Reserved (illegal_instruction)</i>			—
MOVrZ	10 1111	001	Move if Register Zero	$R[rs1] = 0$	<code>movrZ[†] reg_rs1, reg_or_imm10, reg_rd</code>	A1
MOVrLEZ	10 1111	010	Move if Register Less Than or Equal to Zero	$R[rs1] \leq 0$	<code>movrLEZ reg_rs1, reg_or_imm10, reg_rd</code>	A1
MOVrLZ	10 1111	011	Move if Register Less Than Zero	$R[rs1] < 0$	<code>movrLZ reg_rs1, reg_or_imm10, reg_rd</code>	A1
—	10 1111	100	<i>Reserved (illegal_instruction)</i>			—
MOVrNZ	10 1111	101	Move if Register Not Zero	$R[rs1] \neq 0$	<code>movrNZ[‡] reg_rs1, reg_or_imm10, reg_rd</code>	A1
MOVrGZ	10 1111	110	Move if Register Greater Than Zero	$R[rs1] > 0$	<code>movrGZ reg_rs1, reg_or_imm10, reg_rd</code>	A1
MOVrGEZ	10 1111	111	Move if Register Greater Than or Equal to Zero	$R[rs1] \geq 0$	<code>movrGEZ reg_rs1, reg_or_imm10, reg_rd</code>	A1

[†] synonym: `movre` [‡] synonym: `movrne`



Description If the contents of integer register R[rs1] satisfy the condition specified in the rcond field, these instructions copy their second operand (if $i = 0$, R[rs2]; if $i = 1$, `sign_ext(simm10)`) into R[rd]. If the contents of R[rs1] do not satisfy the condition, then R[rd] is not modified.

These instructions treat the register contents as a signed integer value; they do not modify any condition codes.

Programming Note The MOVr instructions are “64-bit-only” instructions; there is no version of these instructions that operates on just the less-significant 32 bits of their source operands.

Implementation Note If this instruction is implemented by tagging each register value with an n (negative) and a z (zero) bit, use the table below to determine if rcond is TRUE.

Move	Test
MOVrNZ	not Z
MOVrZ	Z
MOVrGEZ	not N
MOVrLZ	N
MOVrLEZ	N or Z
MOVrGZ	N nor Z

An attempt to execute a MOVr instruction when either instruction bits 9:5 are nonzero or $rcond = 000_2$ or 100_2 causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*

MOVfTOi

7.95 Move Floating-Point Register to Integer Register

VIS 3B

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
MOVstOsw	1 0001 0011	Copy 32-bit (single-precision) floating-point register to 64-bit Integer register (with sign extension)	movstosw <i>freg_{rs2}, reg_{rd}</i>	C1	OSA 2011
MOVstOuw	1 0001 0001	Copy 32-bit (single-precision) floating-point register to 64-bit Integer register (no sign extension)	movstouw <i>freg_{rs2}, reg_{rd}</i>	C1	OSA 2011
MOVdTOx	1 0001 0000	Copy 64-bit (double-precision) floating-point register to 64-bit Integer register	movdtox <i>freg_{rs2}, reg_{rd}</i>	C1	OSA 2011



Description The MOVstOsw instruction copies 32 bits from floating-point register $F_S[rs2]$ to general-purpose register $R[rd]\{31:0\}$ (with no conversion). It places a copy of the sign bit in each bit of $R[rd]\{63:32\}$; that is, it sign-extends the destination result from 32 to 64 bits.

The MOVstOuw instruction copies 32 bits from floating-point register $F_S[rs2]$ to general-purpose register $R[rd]\{31:0\}$ (with no conversion). It sets $R[rd]\{63:32\}$ to zero; that is, it does not sign-extend the result.

The MOVdTOx instruction copies 64 bits from general-purpose register $F_D[rs2]$ to general-purpose register $R[rd]$. No conversion is performed.

An attempt to execute a MOVfTOi instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute a MOVfTOi instruction causes an *fp_disabled* exception.


Exceptions *illegal_instruction*
fp_disabled

■ **See Also** MOViTOf on page 285

7.96 Move Integer Register to Floating-Point Register

VIS 3B

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
MOVwTOs	1 0001 1001	Copy 32-bit Integer register to 32-bit (single-precision) floating-point register	movwtos <i>reg_{rs2}, freg_{rd}</i>	C1	OSA 2011
MOVxTOd	1 0001 1000	Copy 64-bit Integer register to 64-bit (double-precision) floating-point register	movxtod <i>reg_{rs2}, freg_{rd}</i>	C1	OSA 2011



Description The MOVwTOs instruction copies 32 bits from general-purpose register R[rs1][31:0] to floating-point register F_S[rd]. No conversion is performed on the copied bits.

The MOVxTOd instruction copies 64 bits from general-purpose register R[rs1] to floating-point register F_D[rd]. No conversion is performed on the copied bits.

An attempt to execute a MOViTOf instruction when instruction bits 18:14 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a MOViTOf instruction causes an *fp_disabled* exception.

Exceptions. *illegal_instruction*
fp_disabled

■ *See Also* MOVfTOi on page 284

MPMUL

7.97 MPMUL Crypto

This instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax		Class
MPMUL ^N	1 0100 1000	Multiple Precision Multiply	mpmul	imm5	N1



Description MPMUL multiplies two values, each of width $(N+1) \times 64$ bits, where N is specified in imm5 field of the instruction. The starting locations of the multiplier and multiplicand are constant relative to CWP in the integer register file (IRF), regardless of size. The starting location of the product is also constant. Below are the locations for $N = 31$. For smaller MPMULs ($N < 31$), the remaining multiplier and multiplicand locations are unused and the remaining product locations are unchanged.

MPMUL : product[], multiplier[], and multiplicand[] are arrays of 64-bit doublewords.
 $\text{product}[(2N+1):0] \leftarrow \text{multiplier}[N:0] \times \text{multiplicand}[N:0]$
 Let $i = \text{CWP}$ when MPMUL is executed.
 Note: doubleword 0 is the least significant doubleword.

The operands are located in registers as follows:

MPMUL Operand Locations			
Operand	Window (CWP value)	Register(s)	Notes
<i>Source operands</i>			
multiplier[7:0]	i-6	(F _D [2] :: F _D [0] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	
multiplier[15:8]	i-6	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
multiplier[23:16]	i-6	(F _D [6] :: F _D [4] :: R[29] :: R[28] :: R[27] :: R[26] :: R[25] :: R[24])	f6,f4, i5...i0
multiplier[31:24]	—	(F _D [22] :: F _D [20] :: F _D [18] :: F _D [16] :: F _D [14] :: F _D [12] :: F _D [10] :: F _D [8])	
multiplicand[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
multiplicand[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24, o5...o0
multiplicand[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
multiplicand[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	
<i>Destination operands</i>			
product[7:0]	i-4	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[13:8]	i-4	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...o0
product[21:14]	i-3	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[27:22]	i-3	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...o0
product[35:28]	i-2	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[41:36]	i-2	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...o0

MPMUL

MPMUL Operand Locations

Operand	Window (CWP value)	Register(s)	Notes
product[49:42]	i-1	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[55:50]	i-1	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...00
product[63:56]	i	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10

Programming Note The MPMUL instruction uses seven windows of the integer register file (IRF). Before using the code sequence below, the contents of the IRF must be saved. The code shown below assumes the following window register values prior to execution: CANSAVE=NWINDOWS-2, CANRESTORE=0, OTHERWIN=0.

The following code shows an example usage of MPMUL with length equal to 31.

```

setx    multiplier, %g1, %g4
setx    multiplicand, %g1, %g5

load_multiplier:
    ldd    [%g4 + 0x000], %f22        !# CWP = i-6
    ldd    [%g4 + 0x008], %f20
    ldd    [%g4 + 0x010], %f18
    ldd    [%g4 + 0x018], %f16
    ldd    [%g4 + 0x020], %f14
    ldd    [%g4 + 0x028], %f12
    ldd    [%g4 + 0x030], %f10
    ldd    [%g4 + 0x038], %f8

    ldd    [%g4 + 0x040], %f6
    ldd    [%g4 + 0x048], %f4
    ldx    [%g4 + 0x050], %i5
    ldx    [%g4 + 0x058], %i4
    ldx    [%g4 + 0x060], %i3
    ldx    [%g4 + 0x068], %i2
    ldx    [%g4 + 0x070], %i1
    ldx    [%g4 + 0x078], %i0

    ldx    [%g4 + 0x080], %l7
    ldx    [%g4 + 0x088], %l6
    ldx    [%g4 + 0x090], %l5
    ldx    [%g4 + 0x098], %l4
    ldx    [%g4 + 0x0a0], %l3
    ldx    [%g4 + 0x0a8], %l2
    ldx    [%g4 + 0x0b0], %l1
    ldx    [%g4 + 0x0b8], %l0

    ldd    [%g4 + 0x0c0], %f2
    ldd    [%g4 + 0x0c8], %f0
    ldx    [%g4 + 0x0d0], %o5
    ldx    [%g4 + 0x0d8], %o4
    ldx    [%g4 + 0x0e0], %o3
    ldx    [%g4 + 0x0e8], %o2
    ldx    [%g4 + 0x0f0], %o1
    ldx    [%g4 + 0x0f8], %o0

save                                     !# CWP = i-5

```

MPMUL

```
load_multiplicand:
    ldd    [%g5 + 0x000], %f58
    ldd    [%g5 + 0x008], %f56
    ldd    [%g5 + 0x010], %f54
    ldd    [%g5 + 0x018], %f52
    ldd    [%g5 + 0x020], %f50
    ldd    [%g5 + 0x028], %f48
    ldd    [%g5 + 0x030], %f46
    ldd    [%g5 + 0x038], %f44

    ldd    [%g5 + 0x040], %f42
    ldd    [%g5 + 0x048], %f40
    ldd    [%g5 + 0x050], %f38
    ldd    [%g5 + 0x058], %f36
    ldd    [%g5 + 0x060], %f34
    ldd    [%g5 + 0x068], %f32
    ldd    [%g5 + 0x070], %f30
    ldd    [%g5 + 0x078], %f28

    ldd    [%g5 + 0x080], %f26
    ldd    [%g5 + 0x088], %f24
    ldx    [%g5 + 0x090], %o5
    ldx    [%g5 + 0x098], %o4
    ldx    [%g5 + 0x0a0], %o3
    ldx    [%g5 + 0x0a8], %o2
    ldx    [%g5 + 0x0b0], %o1
    ldx    [%g5 + 0x0b8], %o0

    ldx    [%g5 + 0x0c0], %l7
    ldx    [%g5 + 0x0c8], %l6
    ldx    [%g5 + 0x0d0], %l5
    ldx    [%g5 + 0x0d8], %l4
    ldx    [%g5 + 0x0e0], %l3
    ldx    [%g5 + 0x0e8], %l2
    ldx    [%g5 + 0x0f0], %l1
    ldx    [%g5 + 0x0f8], %l0

    save                    !# CWP = i-4
    save                    !# CWP = i-3
    save                    !# CWP = i-2
    save                    !# CWP = i-1
    save                    !# CWP = i

run_mpmul:
    mpmul    0x1f                    !# CWP = i

store_result:
    setx    vt_result, %g1, %g4
    stx     %l7, [%g4 + 0x000]        !# CWP = i
    stx     %l6, [%g4 + 0x008]
    stx     %l5, [%g4 + 0x010]
    stx     %l4, [%g4 + 0x018]
    stx     %l3, [%g4 + 0x020]
    stx     %l2, [%g4 + 0x028]
    stx     %l1, [%g4 + 0x030]
    stx     %l0, [%g4 + 0x038]

    restore                    !# CWP = i-1

    stx     %o5, [%g4 + 0x040]
```


MPMUL

```
stx    %o4, [%g4 + 0x048]
stx    %o3, [%g4 + 0x050]
stx    %o2, [%g4 + 0x058]
stx    %o1, [%g4 + 0x060]
stx    %o0, [%g4 + 0x068]
```

```
stx    %l7, [%g4 + 0x070]
stx    %l6, [%g4 + 0x078]
stx    %l5, [%g4 + 0x080]
stx    %l4, [%g4 + 0x088]
stx    %l3, [%g4 + 0x090]
stx    %l2, [%g4 + 0x098]
stx    %l1, [%g4 + 0x0a0]
stx    %l0, [%g4 + 0x0a8]
```

```
restore                                !# CWP = i-2
```

```
stx    %o5, [%g4 + 0x0b0]
stx    %o4, [%g4 + 0x0b8]
stx    %o3, [%g4 + 0x0c0]
stx    %o2, [%g4 + 0x0c8]
stx    %o1, [%g4 + 0x0d0]
stx    %o0, [%g4 + 0x0d8]
```

```
stx    %l7, [%g4 + 0x0e0]
stx    %l6, [%g4 + 0x0e8]
stx    %l5, [%g4 + 0x0f0]
stx    %l4, [%g4 + 0x0f8]
stx    %l3, [%g4 + 0x100]
stx    %l2, [%g4 + 0x108]
stx    %l1, [%g4 + 0x110]
stx    %l0, [%g4 + 0x118]
```

```
restore                                !# CWP = i-3
```

```
stx    %o5, [%g4 + 0x120]
stx    %o4, [%g4 + 0x128]
stx    %o3, [%g4 + 0x130]
stx    %o2, [%g4 + 0x138]
stx    %o1, [%g4 + 0x140]
stx    %o0, [%g4 + 0x148]
```

```
stx    %l7, [%g4 + 0x150]
stx    %l6, [%g4 + 0x158]
stx    %l5, [%g4 + 0x160]
stx    %l4, [%g4 + 0x168]
stx    %l3, [%g4 + 0x170]
stx    %l2, [%g4 + 0x178]
stx    %l1, [%g4 + 0x180]
stx    %l0, [%g4 + 0x188]
```

```
restore                                !# CWP = i-4
```

```
stx    %o5, [%g4 + 0x190]
stx    %o4, [%g4 + 0x198]
stx    %o3, [%g4 + 0x1a0]
stx    %o2, [%g4 + 0x1a8]
stx    %o1, [%g4 + 0x1b0]
stx    %o0, [%g4 + 0x1b8]
```

MPMUL

```
stx    %17, [%g4 + 0x1c0]
stx    %16, [%g4 + 0x1c8]
stx    %15, [%g4 + 0x1d0]
stx    %14, [%g4 + 0x1d8]
stx    %13, [%g4 + 0x1e0]
stx    %12, [%g4 + 0x1e8]
stx    %11, [%g4 + 0x1f0]
stx    %10, [%g4 + 0x1f8]

restore                                !# CWP = i-5
restore                                !# CWP = i-6
```

Exceptions. If $rs1 \neq 0$, an attempt to execute an MPMUL instruction causes an *illegal_instruction* exception.

If $CFR.mpmul = 0$, an attempt to execute an MPMUL instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that $CFR.mpmul = 1$ before executing the MPMUL instruction. If $CFR.mpmul = 0$, then software should assume that an attempt to execute the MPMUL instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute but perform some other operation.

Therefore, if $CFR.mpmul = 0$, software should perform the MPMUL operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an MPMUL instruction causes an *fp_disabled* exception.

The MPMUL instruction causes a *fill_n_normal* or *fill_n_other* exception if $CANRESTORE$ is not equal to $NWINDOWS-2$. The fill trap handler is called with CWP set to point to the window to be filled, that is, $old\ CWP - CANRESTORE - 1$. The trap vector for the fill trap is based on the values of $OTHERWIN$ and $WSTATE$, as described in Trap Type for Spill/Fill Traps on page 458. The fill trap handler performs a $RESTORED$ and a $RETRY$ as part of filling the window. When the virtual processor reexecutes MPMUL (due to the handler ending in $RETRY$), another fill trap results if more than one window needed to be filled.


Implementation Note MPMUL and XMPMUL share a basic opcode (op and opf). They are differentiated by the instruction rd field; $rd = 0\ 0000_2$ for MPMUL and $rd = 0\ 0001_2$ for XMPMUL.

Exceptions *fp_disabled*
fill_n_normal ($n = 0-7$)
fill_n_other ($n = 0-7$)

See Also MONTMUL on page 272
MONTSQR on page 276
XMONTMUL on page 378
XMONTSQR on page 381
XMPMUL on page 384

7.98 Multiply and Divide (64-bit)

Instruction	op3	Operation	Assembly Language	Class
MULX	00 1001	Multiply (signed or unsigned)	<code>mulx</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
SDIVX	10 1101	Signed Divide	<code>sdivx</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
UDIVX	00 1101	Unsigned Divide	<code>udivx</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1



Description MULX computes “ $R[rs1] \times R[rs2]$ ” if $i = 0$ or “ $R[rs1] \times \text{sign_ext}(simm13)$ ” if $i = 1$, and writes the 64-bit product into $R[rd]$. MULX can be used to calculate the 64-bit product for signed or unsigned operands (the product is the same).

SDIVX and UDIVX compute “ $R[rs1] \div R[rs2]$ ” if $i = 0$ or “ $R[rs1] \div \text{sign_ext}(simm13)$ ” if $i = 1$, and write the 64-bit result into $R[rd]$. SDIVX operates on the operands as signed integers and produces a corresponding signed result. UDIVX operates on the operands as unsigned integers and produces a corresponding unsigned result.

For SDIVX, if the largest negative number is divided by -1 , the result should be the largest negative number. That is:

$$8000\ 0000\ 0000\ 0000_{16} \div \text{FFFF}\ \text{FFFF}\ \text{FFFF}\ \text{FFFF}_{16} = 8000\ 0000\ 0000\ 0000_{16}.$$

These instructions do not modify any condition codes.


An attempt to execute a MULX, SDIVX, or UDIVX instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
division_by_zero

MWAIT

7.99 MWait

Instruction	op3	Operation	Assembly Language Syntax	Class	Added
MWAIT	11 0000	Monitor-wait virtual processor (strand)	<code>mwait reg_or_imm</code>	C1	OSA 2015



Description

MWAIT attempts to voluntarily pause (temporarily suspend execution on) the virtual processor until a specified number of nanoseconds has elapsed, a block of data containing the memory location accessed by the most recently-executed Load-Monitor instruction is modified by a another virtual processor, or certain other events occur.

If $i = 0$, the virtual processor is requested to pause for a maximum of “R[rs2]” nanoseconds; if $i = 1$, the virtual processor is requested to pause for a maximum of “**sign_ext**(simm13)” nanoseconds.

An MWAIT instruction behaves as a NOP if any of the following is true (all events are assumed to be on the *same* virtual processor as the MWAIT, unless otherwise stated):

- no Load-Monitor instruction has been previously executed
- between execution of the most recent previously-executed Load-Monitor instruction and execution of this MWAIT instruction, any of the following has occurred:
 - a block of data containing the memory location accessed by the Load-Monitor instruction was modified by a virtual processor *other than* the virtual processor executing the MWAIT instruction
 - another MWAIT instruction was executed
 - a synchronous or asynchronous trap occurred
 - certain other synchronous or asynchronous events occurred

If an MWAIT does not behave as a NOP per any of the above conditions, then execution pauses (is temporarily suspended on) its virtual processor until the first of the following events occurs:

- the specified number of nanoseconds has elapsed
- a block of data containing the memory location accessed by the most recent previously-executed Load-Monitor instruction is modified (stored to) by a virtual processor *other than* the virtual processor executing the MWAIT instruction
- an asynchronous trap occurs
- certain other synchronous or asynchronous events occur

The MWAIT instruction operates by writing the specified value to the MWait Count (MWAIT) register. For details, see *Mwait Count (mwait) Register (ASR 28)* on page 61.

Programming Note

MWAIT is intended to be used as part of a progressive (exponential) backoff algorithm. An MWAIT may not affect a given implementation’s behavior if its duration is below some implementation-specific threshold. However, an MWAIT of sufficient duration (larger than the implementation-dependent threshold) and which is not otherwise rendered a NOP causes the virtual processor to suspend execution and retirement of instructions, freeing up any resources shared with other virtual processors for their use.

MWAIT

Programming Note MWAIT 0 can be used to nullify the effect of any earlier (in program order) Load-Monitor instruction. This may be desirable in certain context-switching code.

Implementation Note The MWAIT instruction is implemented as a degenerate case of a WRasr to ASR 28 (1C₁₆), the MWait Count register (MWAIT). A full WRasr to the MWAIT register (WRMWAIT, see page 373) may also be used to write a value to MWAIT, allowing the value to be constructed as the **xor** of two operands.


Exceptions (none)

See Also MWAIT register on page 61
Load-Monitor ASIs on page 437
PAUSE on page 298
WRasr/WRMWAIT on page 373

NOP

7.100 No Operation

Instruction	op2	Operation	Assembly Language Syntax	Class
NOP	100	No Operation	nop	A1



Description The NOP instruction changes no program-visible state (except that of the PC register).

NOP is a special case of the SETHI instruction, with `imm22 = 0` and `rd = 0`.


Programming Note There are many other opcodes that may execute as NOPs; however, this dedicated NOP instruction is the only one guaranteed to be implemented efficiently across all implementations.

Exceptions None

NORMALW

7.101 NORMALW

Instruction	Operation	Assembly Language Syntax	Class
NORMALW ^P	"Other" register windows become "normal" register windows	normalw	A1



Description NORMALW^P is a privileged instruction that copies the value of the OTHERWIN register to the CANRESTORE register, then sets the OTHERWIN register to zero.

Programming Notes The NORMALW instruction is used when changing address spaces. NORMALW indicates the current "other" windows are now "normal" windows and should use the *spill_n_normal* and *fill_n_normal* traps when they generate a trap due to window spill or fill exceptions. The window state may become inconsistent if NORMALW is used when CANRESTORE is nonzero.

An attempt to execute a NORMALW instruction when instruction bits 18:0 are nonzero causes an *illegal_instruction* exception.

An attempt to execute a NORMALW instruction in nonprivileged mode (PSTATE.priv = 0) causes a *privileged_opcode* exception.


Exceptions *illegal_instruction*
privileged_opcode

See Also ALLCLEAN on page 118
INVALW on page 237
OTHERW on page 297
RESTORED on page 317
SAVED on page 324

OR

7.102 OR Logical Operation

Instruction	op3	Operation	Assembly Language Syntax	Class
OR	00 0010	Inclusive or	<code>or</code> $reg_{rs1}, reg_{or_imm}, reg_{rd}$	A1
ORcc	01 0010	Inclusive or and modify cc's	<code>orcc</code> $reg_{rs1}, reg_{or_imm}, reg_{rd}$	A1
ORN	00 0110	Inclusive or not	<code>orn</code> $reg_{rs1}, reg_{or_imm}, reg_{rd}$	A1
ORNcc	01 0110	Inclusive or not and modify cc's	<code>orncc</code> $reg_{rs1}, reg_{or_imm}, reg_{rd}$	A1



Description These instructions implement bitwise logical **or** operations. They compute “R[rs1] **op** R[rs2]” if $i = 0$, or “R[rs1] **op** **sign_ext**(simm13)” if $i = 1$, and write the result into R[rd].

ORcc and ORNcc modify the integer condition codes (icc and xcc). They set the condition codes as follows:

- `icc.v`, `icc.c`, `xcc.v`, and `xcc.c` are set to 0
- `icc.n` is copied from bit 31 of the result
- `xcc.n` is copied from bit 63 of the result
- `icc.z` is set to 1 if bits 31:0 of the result are zero (otherwise to 0)
- `xcc.z` is set to 1 if all 64 bits of the result are zero (otherwise to 0)

ORN and ORNcc logically negate their second operand before applying the main (**or**) operation.


An attempt to execute an OR[N][cc] instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*

OTHERW

7.103 OTHERW

Instruction	Operation	Assembly Language Syntax	Class
OTHERW ^P	“Normal” register windows become “other” register windows	otherw	A1



Description OTHERW^P is a privileged instruction that copies the value of the CANRESTORE register to the OTHERWIN register, then sets the CANRESTORE register to zero.

Programming Notes The OTHERW instruction is used when changing address spaces. OTHERW indicates the current "normal" register windows are now "other" register windows and should use the *spill_n_other* and *fill_n_other* traps when they generate a trap due to window spill or fill exceptions. The window state may become inconsistent if OTHERW is used when OTHERWIN is nonzero.

An attempt to execute an OTHERW instruction when instruction bits 18:0 are nonzero causes an *illegal_instruction* exception.

An attempt to execute an OTHERW instruction in nonprivileged mode (PSTATE.priv = 0) causes a *privileged_opcode* exception.


Exceptions *illegal_instruction*
privileged_opcode

See Also ALLCLEAN on page 118
INVALW on page 237
NORMALW on page 295
RESTORED on page 317
SAVED on page 324

PAUSE

7.104 Pause

Instruction	op3	Operation	Assembly Language Syntax	Class	Added
PAUSE	11 0000	Pause virtual processor (strand)	<code>pause reg_or_imm</code>	C1	OSA 2011



Description Execution of the PAUSE instruction voluntarily pauses (temporarily suspends execution on) a virtual processor for up to the specified number of nanoseconds. If $i = 0$, the virtual processor is requested to pause for “R[rs2]” nanoseconds; if $i = 1$, the virtual processor is requested to pause for “**sign_ext**(simm13)” nanoseconds. If certain asynchronous events occur, the virtual processor may terminate the pause sooner than the requested length of time.

The PAUSE instruction operates by writing the specified value to the Pause Count (PAUSE) register. For details, see *Pause Count (pause) Register (ASR 27)* on page 60.

Programming Note PAUSE is intended to be used as part of a progressive (exponential) backoff algorithm. A pause may not affect a given implementation’s behavior if its duration is below some implementation-specific threshold. However, a pause of sufficient duration (larger than the implementation-dependent threshold) causes the virtual processor to suspend execution and retirement of instructions.

Implementation Note The PAUSE instruction is implemented as a degenerate case of a WRasr to ASR 27 (1B₁₆), the Pause Count register (PAUSE). A full WRasr to the PAUSE register (WRPAUSE) may also be used to write a value to PAUSE, allowing the value to be constructed as the **xor** of two operands.

Exceptions (none)

See Also PAUSE register on page 59
MWAIT on page 292
WRasr on page 373

PDIST

7.105 Pixel Component Distance (with Accumulation) VIS 1

The PDIST instruction is deprecated and should not be used in new software. The PDISTN instructions should be used, instead.

Instruction	opf	Operation	Assembly Language Syntax	Class
PDIST	0 0011 1110	Distance between eight 8-bit components, with accumulation	<code>pdist freg_{rs1}, freg_{rs2}, freg_{rd}</code>	D2



Description Eight unsigned 8-bit values are contained in the 64-bit floating-point source registers $F_D[rs1]$ and $F_D[rs2]$. The corresponding 8-bit values in the source registers are subtracted (that is, each byte in $F_D[rs2]$ is subtracted from the corresponding byte in $F_D[rs1]$). The sum of the absolute value of each difference is added to the integer in $F_D[rd]$ and the resulting integer sum is stored in the destination register, $F_D[rd]$.

Programming Notes PDIST is a “destructive” instruction, in that $F_D[rd]$ serves as both a source and a destination register (read-modify-write). Typically, PDIST is used for motion estimation in video compression algorithms. The (deprecated) PDIST instruction does not perform well on every Oracle SPARC Architecture implementation and its performance is expected to degrade further on future implementations. Use of PDISTN plus a subsequent ADD is recommended as a higher-performance replacement for PDIST.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute a PDIST instruction causes an *fp_disabled* exception.

Exceptions *fp_disabled*

See Also PDISTN on page 300

PDISTN

7.106 Pixel Component Distance (No Accumulation) VIS 3

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
PDISTN	0 0011 1111	Distance between eight 8-bit components (with no accumulation)	f64	f64	i64	<code>pdistn freg_{rs1}, freg_{rs2}, reg_{rd}</code>	C1	OSA 2011



Description Eight unsigned 8-bit values are contained in the 64-bit floating-point source registers, $F_D[rs1]$ and $F_D[rs2]$. The corresponding 8-bit values in $F_D[rs1]$ and $F_D[rs2]$ are subtracted (that is, $F_D[rs1]\{n+7:n\} - F_D[rs2]\{n+7:n\}$). The integer sum of the absolute values of each of the eight differences is stored in the destination register, $R[rd]$.

Note PDISTN is intended as a functional replacement for PDIST. As explained on the PDIST instruction page, new software should avoid using PDIST and use PDISTN, instead.

Note There are two differences between PDISTN and PDIST.
 (1) PDISTN is a two-source-operand instruction, unlike PDIST, which is a three-source-operand instruction that also performs accumulation.
 (2) PDISTN's destination register is an integer register, while PDIST's destination register is a floating-point register.

Programming Notes Typically, PDISTN is used for motion estimation in video compression algorithms and is followed by an `ADD[cc]` instruction, which performs accumulation.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute a PDISTN instruction causes an *fp_disabled* exception.


Exceptions *fp_disabled*

See Also PDIST on page 299

POPC

7.107 Population Count

Instruction	op3	Operation	Assembly Language Syntax	Class
POPC	10 1110	Population Count	popc <i>reg_or_imm</i> , <i>reg_rd</i>	C2



Description POPC counts the number of ‘1’ bits in R[rs2] if $i = 0$, or the number of ‘1’ bits in `sign_ext(simm13)` if $i = 1$, and stores the count in R[rd]. This instruction does not modify the condition codes.

V9 Compatibility Note Instruction bits 18 through 14 must be zero for POPC. Other encodings of this field (rs1) may be used in future versions of the SPARC architecture for other instructions.

Programming Note POPC can be used to “find first bit set” from the least-significant (right) end (or, equivalently, Trailing Zero Count) in an integer register. A ‘C’-language program illustrating how POPC can be used for this purpose follows:

```
int ffs(x) /* finds first 1 bit, counting from the LSB */
unsigned long long int x;
{
    return popc(in ^ (~(-x))); /* for nonzero x */
}
```

Inline assembly language code for `ffs()` is:

```
neg    %X, %NEG_X    ! -x(2's complement)
xnor   %X, %NEG_X, %TEMP ! ^ ~ -x (exclusive nor)
popc   %TEMP, %RESULT ! result = popc(x ^ ~ -x)
movrzs %X, %g0, %RESULT ! %RESULT should be 0 for %X=0
```

where `X`, `NEG_X`, `TEMP`, and `RESULT` are integer registers.

Example computation:

```

X = ...00101000 !1st '1' bit from right is
-X = ...11011000 ! bit 3 (4th bit)
~ -X = ...00100111
X ^ ~ -X = ...00001111
popc(X ^ ~ -X) = 4
```

Programming Note Another means of implementing “Trailing Zero Count” is:
or `TZCNT(x) = POPC((not(x) and (x-1)))`

Programming Note POPC can be used to “centrifuge” all the ‘1’ bits in a register to the least significant end of a destination register. Assembly-language code illustrating how POPC can be used for this purpose follows:

```
popc   %X, %DEST
cmp    %X, -1      ! Test for pattern of all 1's
mov    -1, %TEMP   ! Constant -1 -> temp register
sllx  %TEMP, %DEST, %DEST ! (shift count of 64 same as 0)
not    %DEST
movcc  %xcc, -1, %DEST ! If src was -1, result is -1
```

where `X`, `TEMP`, and `DEST` are integer registers.

POPC

Programming Note POPC is a “64-bit-only” instruction; there is no version of this instruction that operates on just the less-significant 32 bits of its source operand.

An attempt to execute a POPC instruction when either instruction bits 18:14 are nonzero, or *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*


■ *See Also* *Leading Zeroes Count on page 266*

PREFETCH

7.108 Prefetch

Instruction	op3	Operation	Assembly Language Syntax	Class
PREFETCH	10 1101	Prefetch Data	<code>prefetch [address], prefetch_fcn</code>	A1
PREFETCHA ^{PASI}	11 1101	Prefetch Data from Alternate Space	<code>prefetcha [regaddr] imm_asi, prefetch_fcn</code> <code>prefetcha [reg_plus_imm] %asi, prefetch_fcn</code>	A1

PREFETCH



PREFETCHA




TABLE 7-16 Prefetch Variants, by Function Code

fcn	Prefetch Variant	<i>prefetch_fcn</i>
0	(Weak) Prefetch for several reads	#n_reads
1	(Weak) Prefetch for one read	#one_read
2	(Weak) Prefetch for several writes and possibly reads	#n_writes
3	(Weak) Prefetch for one write	#one_write
4	Prefetch page	#page
5–15 (05 ₁₆ –0F ₁₆)	Reserved (<i>illegal_instruction</i>)	
16 (10 ₁₆)	Implementation dependent (NOP if not implemented)	
17 (11 ₁₆)	Prefetch to nearest unified cache	#unified
18–19 (12 ₁₆ –13 ₁₆)	Implementation dependent (NOP if not implemented)	
20 (14 ₁₆)	Strong Prefetch for several reads	#n_reads_strong
21 (15 ₁₆)	Strong Prefetch for one read	#one_read_strong
22 (16 ₁₆)	Strong Prefetch for several writes and possibly reads	#n_writes_stron
23 (17 ₁₆)	Strong Prefetch for one write	#one_write_strong
24–31 (18 ₁₆ –1F ₁₆)	Implementation dependent (NOP if not implemented)	
24 (18 ₁₆)	Implementation dependent (NOP if not implemented)	
25 (19 ₁₆)	Implementation dependent (NOP if not implemented)	
26–31 (1A ₁₆ –1F ₁₆)	Implementation dependent (NOP if not implemented)	

Description

A PREFETCH[A] instruction provides a hint to the virtual processor that software expects to access a particular address in memory in the near future, so that the virtual processor may take action to reduce the latency of accesses near that address. Typically, execution of a prefetch instruction initiates

PREFETCH

movement of a block of data containing the addressed byte from memory toward the virtual processor or creates an address mapping.

Implementation | A PREFETCH[A] instruction may be used by software to:

- Note**
- prefetch a cache line into a cache
 - prefetch a valid address translation into a TLB
 -

If $i = 0$, the effective address operand for the PREFETCH instruction is “ $R[rs1] + R[rs2]$ ”; if $i = 1$, it is “ $R[rs1] + \text{sign_ext}(\text{simm13})$ ”.

PREFETCH instructions access the primary address space (ASI_PRIMARY[_LITTLE]).

PREFETCHA instructions access an alternate address space. If $i = 0$, the address space identifier (ASI) to be used for the instruction is in the `imm_asi` field. If $i = 1$, the ASI is found in the ASI register.

A prefetch operates much the same as a regular load operation, but with certain important differences. In particular, a PREFETCH[A] instruction is non-blocking; subsequent instructions can continue to execute while the prefetch is in progress.

Implementation | A PREFETCH[A] instruction is “released” by hardware after the
Note | TLB access, allowing subsequent instructions to continue to execute while the virtual processor performs the hardware tablewalk (in the case of a TLB miss for a Strong prefetch) or the cache access in the background.

When executed in nonprivileged or privileged mode, PREFETCH[A] has the same observable effect as a NOP. A prefetch instruction will not cause a trap if applied to an illegal or nonexistent memory address. (impl. dep. #103-V9-Ms10(e))

IMPL. DEP. #103-V9-Ms10(a): The size and alignment in memory of the data block prefetched is implementation dependent; the minimum size is 64 bytes and the minimum alignment is a 64-byte boundary.

Programming | Software may prefetch 64 bytes beginning at an arbitrary address
Note | address by issuing the instructions

```
prefetch [address], prefetch_fcn
prefetch [address + 63], prefetch_fcn
```

Variants of the prefetch instruction can be used to prepare the memory system for different types of accesses.

IMPL. DEP. #103-V9-Ms10(b): An implementation may implement none, some, or all of the defined PREFETCH[A] variants. It is implementation-dependent whether each variant is (1) not implemented and executes as a NOP, (2) is implemented and supports the full semantics for that variant, or (3) is implemented and only supports the simple common-case prefetching semantics for that variant.

PREFETCH

7.108.1 Exceptions

Prefetch instructions PREFETCH and PREFETCHA generate exceptions under the conditions detailed in TABLE 7-17. Only the implementation-dependent prefetch variants (see TABLE 7-16) may generate an exception under conditions not listed in this table; the predefined variants only generate the exceptions listed here.

TABLE 7-17 Behavior of PREFETCH[A] Instructions Under Exceptional Conditions

fcn	Instruction	Condition	Result
any	PREFETCH	$i = 0$ and instruction bits 12:5 are nonzero	<i>illegal_instruction</i>
any	PREFETCHA	reference to an ASI in the range 0_{16} - $7F_{16}$, while in nonprivileged mode (<i>privileged_action</i> condition)	executes as NOP
any	PREFETCHA	reference to an ASI in range 30_{16} - $7F_{16}$, while in privileged mode (<i>privileged_action</i> condition)	executes as NOP
0-3 (weak)	PREFETCH[A]	condition detected for MMU miss	executes as NOP
0-4	PREFETCH[A]	variant unimplemented	executes as NOP
0-4	PREFETCHA	reference to an invalid ASI (ASI not listed in following table)	executes as NOP
0-4, 17, 20-23	PREFETCH[A]	condition detected for <i>DAE_invalid_asi</i> (see following table), <i>DAE_privilege_violation</i> , <i>DAE_nc_page</i> (TTE.cp = 0), <i>DAE_nfo_page</i> , or <i>DAE_side_effect_page</i> (TTE.e = 1)	executes as NOP
4, 20-23 (strong)	PREFETCH[A]	prefetching the requested data would be a very time-consuming operation	executes as NOP
5-15 (05_{16} - $0F_{16}$)	PREFETCH[A]	(always)	<i>illegal_instruction</i>
16 (10_{16})	PREFETCH[A]	variant unimplemented	executes as NOP
17-31 (11_{16} - $1F_{16}$)	PREFETCH[A]	variant unimplemented	executes as NOP
24 (18_{16})	PREFETCH[A]	variant unimplemented	executes as NOP
25 (19_{16})	PREFETCH[A]	variant unimplemented	executes as NOP

ASIs valid for PREFETCHA (all other ASIs are invalid)

04_{16} ASI_NUCLEUS	$0C_{16}$ ASI_NUCLEUS_LITTLE
10_{16} ASI_AS_IF_USER_PRIMARY	18_{16} ASI_AS_IF_USER_PRIMARY_LITTLE
11_{16} ASI_AS_IF_USER_SECONDARY	19_{16} ASI_AS_IF_USER_SECONDARY_LITTLE
14_{16} ASI_REAL	$1C_{16}$ ASI_REAL_LITTLE
80_{16} ASI_PRIMARY	88_{16} ASI_PRIMARY_LITTLE
81_{16} ASI_SECONDARY	89_{16} ASI_SECONDARY_LITTLE
82_{16} ASI_PRIMARY_NO_FAULT	$8A_{16}$ ASI_PRIMARY_NO_FAULT_LITTLE
83_{16} ASI_SECONDARY_NO_FAULT	$8B_{16}$ ASI_SECONDARY_NO_FAULT_LITTLE

PREFETCH

7.108.2 Weak versus Strong Prefetches

Some prefetch variants are available in two versions, “Weak” and “Strong”.

From software’s perspective, the difference between the two is the degree of certainty that the data being prefetched will subsequently be accessed. That, in turn, affects the amount of effort (time) the underlying hardware will invest to perform the prefetch. If the prefetch is speculative (software believes the data will probably be needed, but isn’t sure), a Weak prefetch will initiate data movement if the operation can be performed quickly, but abort the prefetch and behave like a NOP if it turns out that performing the full prefetch will be time-consuming. If software has very high confidence that data being prefetched will subsequently be accessed, then a Strong prefetch will ensure that the prefetch operation will continue, even if the prefetch operation does become time-consuming.

From the virtual processor’s perspective, the difference between a Weak and a Strong prefetch is whether the prefetch is allowed to perform a time-consuming operation in order to complete. If a time-consuming operation is required, a Weak prefetch will abandon the operation and behave like a NOP while a Strong prefetch will pay the cost of performing the time-consuming operation so it can finish initiating the requested data movement.

Behavioral differences among loads, strong prefetches, and weak prefetches are compared in TABLE 7-18.

TABLE 7-18 Comparative Behavior of Load, Weak Prefetch, and Strong Prefetch Operations

Condition	Behavior		
	Load	Weak Prefetch	Strong Prefetch
Upon detection of <i>privileged_action</i> , any of the <i>DAE_*</i> exceptions or <i>VA_watchpoint</i> exception...	Traps	NOP‡	NOP‡
If page table entry has <i>cp</i> = 0, and <i>e</i> = 1 for Prefetch for Several Reads	Traps	NOP‡	NOP‡
If page table entry has <i>nfo</i> = 1 for a non-NoFault access...	Traps	NOP‡	NOP‡
If page table entry has <i>w</i> = 0 for any prefetch for write access (<i>fcn</i> = 2, 3, 22, or 23)...	Traps	NOP‡	NOP‡
Instruction blocks until cache line filled?	Yes	No	No

‡ The PREFETCH[A] instruction aborts, appearing to behave like a NOP instruction.

7.108.3 Prefetch Variants

The prefetch variant is selected by the *fcn* field of the instruction. *fcn* values 5–15 are reserved for future extensions of the architecture, and PREFETCH *fcn* values of 16–19 and 24–31 are implementation dependent in Oracle SPARC Architecture 2015.

Each prefetch variant reflects an intent on the part of the compiler or programmer, a “hint” to the underlying virtual processor. This is different from other instructions (except BPN), all of which cause specific actions to occur. An Oracle SPARC Architecture implementation may implement a prefetch variant by any technique, as long as the intent of the variant is achieved (impl. dep. #103-V9-Ms10(b)).

The prefetch instruction is designed to treat common cases well. The variants are intended to provide scalability for future improvements in both hardware and compilers. If a variant is implemented, it should have the effects described below. In case some of the variants listed below are implemented and some are not, a recommended overloading of the unimplemented variants is provided in the SPARC V9 specification. An implementation must treat any unimplemented prefetch *fcn* values as NOPs (impl. dep. #103-V9-Ms10).

PREFETCH

7.108.3.1 Prefetch for Several Reads (f_{cn} = 0, 20(14₁₆))

The intent of these variants is to cause movement of data into the cache nearest the virtual processor.

There are Weak and Strong versions of this prefetch variant; f_{cn} = 0 is Weak and f_{cn} = 20 is Strong. The choice of Weak or Strong variant controls the degree of effort that the virtual processor may expend to obtain the data.

Programming Note	The intended use of this variant is for streaming relatively small amounts of data into the primary data cache of the virtual processor.
-------------------------	--

7.108.3.2 Prefetch for One Read (f_{cn} = 1, 21(15₁₆))

The data to be read from the given address are expected to be read once and not reused (read or written) soon after that. Use of this PREFETCH variant indicates that, if possible, the data cache should be minimally disturbed by the data read from the given address.

There are Weak and Strong versions of this prefetch variant; f_{cn} = 1 is Weak and f_{cn} = 21 is Strong. The choice of Weak or Strong variant controls the degree of effort that the virtual processor may expend to obtain the data.

Programming Note	The intended use of this variant is in streaming medium amounts of data into the virtual processor without disturbing the data in the primary data cache memory.
-------------------------	--

7.108.3.3 Prefetch for Several Writes (and Possibly Reads) (f_{cn} = 2, 22(16₁₆))

The intent of this variant is to cause movement of data in preparation for multiple writes.

There are Weak and Strong versions of this prefetch variant; f_{cn} = 2 is Weak and f_{cn} = 22 is Strong. The choice of Weak or Strong variant controls the degree of effort that the virtual processor may expend to obtain the data.

Programming Note	An example use of this variant is to initialize a cache line, in preparation for a partial write.
-------------------------	---

Implementation Note	On a multiprocessor system, this variant indicates that exclusive ownership of the addressed data is needed. Therefore, it may have the additional effect of obtaining exclusive ownership of the addressed cache line.
----------------------------	---

7.108.3.4 Prefetch for One Write (f_{cn} = 3, 23(17₁₆))

The intent of this variant is to initiate movement of data in preparation for a single write. This variant indicates that, if possible, the data cache should be minimally disturbed by the data written to this address, because those data are not expected to be reused (read or written) soon after they have been written once.

There are Weak and Strong versions of this prefetch variant; f_{cn} = 3 is Weak and f_{cn} = 23 is Strong. The choice of Weak or Strong variant controls the degree of effort that the virtual processor may expend to obtain the data.

PREFETCH

7.108.3.5 Prefetch Page (fcn = 4)

In a virtual memory system, the intended action of this variant is for hardware (or privileged or hyperprivileged software) to initiate asynchronous mapping of the referenced virtual address (assuming that it is legal to do so).

Programming Note | Prefetch Page is used is to avoid a later page fault for the given address, or at least to shorten the latency of a page fault.

In a non-virtual-memory system or if the addressed page is already mapped, this variant has no effect.

Implementation Note | The mapping required by Prefetch Page may be performed by privileged software, hyperprivileged software, or hardware.

7.108.3.6 Prefetch to Nearest Unified Cache (fcn = 17(11₁₆))

The intent of this variant is to cause movement of data from memory into the nearest unified (combined instruction and data) cache. At the successful completion of this variant, the selected line from memory will be in the unified cache in the shared state, and in caches (if any) below it in the cache hierarchy.

Prefetch to Nearest Unified Cache is a Strong prefetch variant.

7.108.4 Implementation-Dependent Prefetch Variants (fcn = 16, 18, 19, and 24–31)

IMPL. DEP. #103-V9-Ms10(c): Whether and how PREFETCH fcns 16, 18, 19 and 24-31 are implemented are implementation dependent. If a variant is not implemented, it must execute as a NOP.

7.108.5 Additional Notes

Programming Note | Prefetch instructions do have some “cost to execute”. As long as the cost of executing a prefetch instruction is well less than the cost of a cache miss, use of prefetching provides a net gain in performance.

It does not appear that prefetching causes a significant number of useless fetches from memory, though it may increase the rate of *useful* fetches (and hence the bandwidth), because it more efficiently overlaps computing with fetching.

Programming Note | A compiler that generates PREFETCH instructions should generate each of the variants where its use is most appropriate. That will help portable software be reasonably efficient across a range of hardware configurations.

Implementation Note | Any effects of a data prefetch operation in privileged code should be reasonable (for example, no page prefetching is allowed within code that handles page faults). The benefits of prefetching should be available to most privileged code.

PREFETCH

Implementation Note | A prefetch from a nonprefetchable location has no effect. It is up to memory management hardware to determine how locations are identified as not prefetchable.

Exceptions *illegal_instruction*



7.109 Read Ancillary State Register

Instruction	rs1	Operation	Assembly Language Syntax	Class	Added
RDY ^D	0	Read Y register (<i>deprecated</i>)	rd %y, reg _{rd}	D3	
—	1	<i>Reserved</i>			
RDCCR	2	Read Condition Codes register (CCR)	rd %ccr, reg _{rd}	A1	
RDASI	3	Read ASI register	rd %asi, reg _{rd}	A1	
RDTICK ^{P_{dis},H_{dis}}	4	Read TICK register	rd %tick, reg _{rd}	A1	
RDPC	5	Read Program Counter (PC)	rd %pc, reg _{rd}	A2	
RDFPRS	6	Read Floating-Point Registers Status (FPRS) register	rd %fprs, reg _{rd}	A1	
—	7–14 (7-0E ₁₆)	<i>Reserved</i>			
—	13 (0D ₁₆)	<i>Reserved</i>			
—	14 (0E ₁₆)	<i>Reserved</i>			
See text	15 (F ₁₆)	MEMBAR or <i>Reserved</i> ; see text			
—	16–18 (10 ₁₆ –12 ₁₆)	<i>Reserved</i> (impl. dep. #8-V8-Cs20, 9-V8-Cs20)			
RDGSR	19 (13 ₁₆)	Read General Status register (GSR)	rd %gsr, reg _{rd}	A1	
—	20–21 (14 ₁₆ –15 ₁₆)	<i>Reserved</i> (impl. dep. #8-V8-Cs20, #9-V8-Cs20)			
RDSOFTINT ^P	22 (16 ₁₆)	Read per-virtual processor Soft Interrupt register (SOFTINT)	rd %softint, reg _{rd}	A2	
—	23 (17 ₁₆)	<i>Reserved</i>			
RDSTICK ^{P_{dis},H_{dis}}	24 (18 ₁₆)	Read System Tick Register (STICK)	rd %stick†, reg _{rd}	A2	
RDSTICK_CMPR ^P	25 (19 ₁₆)	Read System Tick Compare register (STICK_CMPR)	rd %stick_cmpr†, reg _{rd}	A2	
RDCFR	26 (1A ₁₆)	Read Compatibility Features register (CFR)	rd %cfr, reg _{rd}	A2	OSA 2011
—	27 (1B ₁₆)	<i>Reserved</i> (used for PAUSE, via WRAsr)		A1	
—	28 (1C ₁₆)	<i>Reserved</i> (used for MWAIT, via WRAsr)			
—	29 (1D ₁₆)	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)			
—	30 (1E ₁₆)	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)			
—	31 (1F ₁₆)	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)			

† The original assembly language names for %stick and %stick_cmpr were, respectively, %sys_tick and %sys_tick_cmpr, which are now deprecated. Over time, assemblers will support the new %stick and %stick_cmpr names for these registers (which are consistent with %tick). In the meantime, some existing assemblers may only recognize the original names.



RDAsr

Description

The Read Ancillary State Register (RDAsr) instructions copy the contents of the state register specified by `rs1` into `R[rd]`.

An RDAsr instruction with `rs1 = 0` is a (deprecated) RDY instruction (which should not be used in new software).

The RDY instruction is deprecated. It is recommended that all instructions that reference the Y register be avoided.

RDPC copies the contents of the PC register into `R[rd]`. If `PSTATE.am = 0`, the full 64-bit address is copied into `R[rd]`. If `PSTATE.am = 1`, only a 32-bit address is saved; `PC{31:0}` is copied to `R[rd]{31:0}` and `R[rd]{63:32}` is set to 0. (closed impl. dep. #125-V9-Cs10)

RDFPRS waits for all pending FPOps and loads of floating-point registers to complete before reading the FPRS register.

The following values of `rs1` are reserved for future versions of the architecture: 1, 7–13, 16–18, 20–21, and 27.

IMPL. DEP. #47-V8-Cs20: RDAsr instructions with `rd` in the range 28–31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For an RDAsr instruction with `rs1` in the range 28–31, the following are implementation dependent:

- the interpretation of bits 13:0 and 29:25 in the instruction
- whether the instruction is nonprivileged or privileged (impl. dep. #9-V8-Cs20), and
- whether an attempt to execute the instruction causes an *illegal_instruction* exception.

Note | Ancillary state registers may include (for example) timer, counter, diagnostic, self-test, and trap-control registers.

SPARC V8 Compatibility Note | The SPARC V8 RDPSR, RDWIM, and RDTBR instructions do not exist in the Oracle SPARC Architecture, since the PSR, WIM, and TBR registers do not exist.

See *Ancillary State Registers* on page 48 for more detailed information regarding ASR registers.

Exceptions. An attempt to execute a RDAsr instruction when any of the following conditions are true causes an *illegal_instruction* exception:

- `rs1` contains a reserved value (that is not assigned for implementation-dependent use)
- instruction bits 13:0 are nonzero

An attempt to execute `aRDSTICK_CMPR`, or `RDSOFTINT` instruction in nonprivileged mode (`PSTATE.priv = 0`) causes a *privileged_opcode* exception (impl. dep. #250-U3-Cs10).

Nonprivileged software can read the TICK register by using the RDTICK instruction, but only when nonprivileged access to TICK is enabled. If nonprivileged access is disabled, an attempt by nonprivileged software to read the TICK register using the RDTICK instruction causes a *privileged_action* exception. See *Tick (tick) Register (ASR 4)* on page 52 for details.

Nonprivileged software can read the STICK register by using the RDSTICK instruction, but only when nonprivileged access to STICK is enabled. If nonprivileged access is disabled, an attempt by nonprivileged software to read the STICK register causes a *privileged_action* exception. See *System Tick (stick) Register (ASR 24)* on page 56 for details.

Privileged software can read the STICK register with the RDSTICK instruction, but only when privileged access to STICK is enabled by hyperprivileged software. An attempt by privileged software to read the STICK register when privileged access is disabled causes a *privileged_action* exception. See *System Tick (stick) Register (ASR 24)* on page 56 for details.

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute a RDGSR instruction causes an *fp_disabled* exception.

RDasr

In nonprivileged mode (PSTATE.priv = 0), the following cause a *privileged_action* exception:

- execution of RDTICK when nonprivileged access to TICK is disabled
- execution of RDSTICK when nonprivileged access to STICK is disabled

In privileged mode (PSTATE.priv = 1), the following cause a *privileged_action* exception:

- execution of RDTICK or RDSTICK when privileged access to TICK and STICK is disabled

Implementation | RDasr shares an opcode with MEMBAR; RDasr is distinguished

Note | by rs1 = 15 or rd = 0 or (i = 0, and bit 12 = 0).

Exceptions

illegal_instruction
privileged_opcode
fp_disabled
privileged_action


See Also

RDPR on page 313
WRasr on page 373

RDPR

7.110 Read Privileged Register

Instruction	op3	Operation	rs1	Assembly Language Syntax	Class
RDPR ^P	10 1010	Read Privileged register			A2?
		TPC	0	rdpr %tpc, regrd	A1?
		TNPC	1	rdpr %tnpc, regrd	
		TSTATE	2	rdpr %tstate, regrd	
		TT	3	rdpr %tt, regrd	
		TICK	4	rdpr %tick, regrd	
		TBA	5	rdpr %tba, regrd	
		PSTATE	6	rdpr %pstate, regrd	
		TL	7	rdpr %tl, regrd	
		PIL	8	rdpr %pil, regrd	
		CWP	9	rdpr %cwp, regrd	
		CANSAVE	10	rdpr %cansave, regrd	
		CANRESTORE	11	rdpr %canrestore, regrd	
		CLEANWIN	12	rdpr %cleanwin, regrd	
		OTHERWIN	13	rdpr %otherwin, regrd	
		WSTATE	14	rdpr %wstate, regrd	
		Reserved	15		
		GL	16	rdpr %gl, regrd	
		Reserved	17–22		
		Reserved	23		
		Reserved	24–31		



Description The rs1 field in the instruction determines the privileged register that is read. There are *MAXPTL* copies of the TPC, TNPC, TT, and TSTATE registers. A read from one of these registers returns the value in the register indexed by the current value in the trap level register (TL). A read of TPC, TNPC, TT, or TSTATE when the trap level is zero (TL = 0) causes an *illegal_instruction* exception.

An attempt to execute a RDPR instruction when any of the following conditions exist causes an *illegal_instruction* exception:

- instruction bits 13:0 are nonzero
- rs1 contains a reserved value (see above)
- $0 \leq rs1 \leq 3$ (attempt to read TPC, TNPC, TSTATE, or TT register) while TL = 0 (current trap level is zero) and the virtual processor is in privileged mode.

Implementation In nonprivileged mode, *illegal_instruction* exception due to
Note $0 \leq rs1 \leq 3$ and TL = 0 does not occur; the *privileged_opcode* exception occurs instead.

An attempt to execute a RDPR instruction in nonprivileged mode (PSTATE.priv = 0) causes a *privileged_opcode* exception.

RDPR

Historical Note | On some early SPARC implementations, floating-point exceptions could cause deferred traps. To ensure that execution could be correctly resumed after handling a deferred trap, hardware provided a floating-point queue (FQ), from which the address of the trapping instruction could be obtained by the trap handler. The front of the FQ was accessed by executing a RDPR instruction with `rs1 = 15`.

On Oracle SPARC Architecture implementations, all floating-point traps are precise. When one occurs, the address of a trapping instruction can be found by the trap handler in the TPC[TL], so no floating-point queue (FQ) is needed or implemented (impl. dep. #25-V8) and RDPR with `rs1 = 15` generates an *illegal_instruction* exception.


Exceptions *illegal_instruction*
 privileged_opcode

| *See Also* RDAsr on page 310
 WRPR on page 376

RESTORE

7.111 RESTORE

Instruction	op3	Operation	Assembly Language Syntax	Class
RESTORE	11 1101	Restore Caller's Window	<code>restore <i>reg_rs1</i>, <i>reg_or_imm</i>, <i>reg_rd</i></code>	A1



Description The RESTORE instruction restores the register window saved by the last SAVE instruction executed by the current process. The *in* registers of the old window become the *out* registers of the new window. The *in* and *local* registers in the new window contain the previous values.

Furthermore, if and only if a fill trap is not generated, RESTORE behaves like a normal ADD instruction, except that the source operands R[rs1] or R[rs2] are read from the *old* window (that is, the window addressed by the original CWP) and the sum is written into R[rd] of the *new* window (that is, the window addressed by the new CWP).

Note CWP arithmetic is performed modulo the number of implemented windows, *N_REG_WINDOWS*.

Programming Notes Typically, if a RESTORE instruction traps, the fill trap handler returns to the trapped instruction to reexecute it. So, although the ADD operation is not performed the first time (when the instruction traps), it is performed the second time the instruction executes. The same applies to changing the CWP.

There is a performance trade-off to consider between using SAVE/RESTORE and saving and restoring selected registers explicitly.

Description (Effect on Privileged State)

If a RESTORE instruction does not trap, it decrements the CWP (**mod** *N_REG_WINDOWS*) to restore the register window that was in use prior to the last SAVE instruction executed by the current process. It also updates the state of the register windows by decrementing CANRESTORE and incrementing CANSERVE.

If the register window to be restored has been spilled (CANRESTORE = 0), then a fill trap is generated. The trap vector for the fill trap is based on the values of OTHERWIN and WSTATE, as described in *Trap Type for Spill/Fill Traps* on page 458. The fill trap handler is invoked with CWP set to point to the window to be filled, that is, old CWP - 1.

Programming Note The vectoring of fill traps can be controlled by setting the value of the OTHERWIN and WSTATE registers appropriately.

The fill handler normally will end with a RESTORED instruction followed by a RETRY instruction.

An attempt to execute a RESTORE instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

RESTORE


Exceptions *illegal_instruction*
 fill_n_normal (*n* = 0–7)
 fill_n_other (*n* = 0–7)

■ *See Also* SAVE on page 322

RESTORED

7.112 RESTORED

Instruction	Operation	Assembly Language Syntax	Class
RESTORED ^P	Window has been restored	restored	A1



Description RESTORED adjusts the state of the register-windows control registers.

RESTORED increments CANRESTORE.

If CLEANWIN < (N_REG_WINDOWS-1), then RESTORED increments CLEANWIN.

If OTHERWIN = 0, RESTORED decrements CANSAVE. If OTHERWIN ≠ 0, it decrements OTHERWIN.

Programming Notes Trap handler software for register window fills use the RESTORED instruction to indicate that a window has been filled successfully.

Normal privileged software would probably not execute a RESTORED instruction from trap level zero (TL = 0). However, it is not illegal to do so and doing so does not cause a trap.

Executing a RESTORED instruction outside of a window fill trap handler is likely to create an inconsistent window state. Hardware will not signal an exception, however, since maintaining a consistent window state is the responsibility of privileged software.

If ((CANSAVE = 0) and (OTHERWIN = 0)) or (CANRESTORE ≥ (N_REG_WINDOWS - 2)) just prior to execution of a RESTORED instruction, the subsequent behavior of the processor is undefined. In neither of these cases can RESTORED generate a register window state that is both valid (see *Register Window State Definition* on page 66) and consistent with the state prior to the RESTORED.

An attempt to execute a RESTORED instruction when instruction bits 18:0 are nonzero causes an *illegal_instruction* exception.

An attempt to execute a RESTORED instruction in nonprivileged mode (PSTATE.priv = 0) causes a *privileged_opcode* exception.


Exceptions *illegal_instruction*
privileged_opcode

See Also ALLCLEAN on page 118
INVALW on page 237
NORMALW on page 295
OTHERW on page 297
SAVED on page 324

RETRY

7.113 RETRY

Instruction	op3	Operation	Assembly Language Syntax	Class
RETRY ^P	11 1110	Return from Trap (retry trapped instruction)	retry	A1



Description The RETRY instruction restores the saved state from TSTATE[TL] (GL, CCR, ASI, PSTATE, and CWP), sets PC and NPC, and decrements TL. RETRY sets $PC \leftarrow TPC[TL]$ and $NPC \leftarrow TNPC[TL]$ (normally, the values of PC and NPC saved at the time of the original trap).

Programming Note | The DONE and RETRY instructions are used to return from privileged trap handlers.

If the saved TPC[TL] and TNPC[TL] were not altered by trap handler software, RETRY causes execution to resume at the instruction that originally caused the trap (“retrying” it).

Execution of a RETRY instruction in the delay slot of a control-transfer instruction produces undefined results.

If software writes invalid or inconsistent state to TSTATE before executing RETRY, virtual processor behavior during and after execution of the RETRY instruction is undefined.

When PSTATE.am = 1, the more-significant 32 bits of the target instruction address are masked out (set to 0) before being sent to the memory system.

IMPL. DEP. #417-S10: If (1) TSTATE[TL].pstate.am = 1 and (2) a RETRY instruction is executed (which sets PSTATE.am to ‘1’ by restoring the value from TSTATE[TL].pstate.am to PSTATE.am), it is implementation dependent whether the RETRY instruction masks (zeroes) the more-significant 32 bits of the values it places into PC and NPC.

Exceptions. An attempt to execute the RETRY instruction when either of the following conditions is true causes an *illegal_instruction* exception:

- instruction bits 18:0 are nonzero
- TL = 0 and the virtual processor is in privileged mode (PSTATE.priv = 1)

An attempt to execute a RETRY instruction in nonprivileged mode (PSTATE.priv = 0) causes a *privileged_opcode* exception.

Implementation Note | In nonprivileged mode, *illegal_instruction* exception due to TL = 0 does not occur. The *privileged_opcode* exception occurs instead, regardless of the current trap level (TL).

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20) and PSTATE.tct = 1, then RETRY generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the RETRY instruction) is stored in TPC[TL] and the value of NPC from before the RETRY was executed is stored in TNPC[TL]. The full 64-bit (nonmasked) PC and NPC values are stored in TPC[TL] and TNPC[TL], regardless of the value of PSTATE.am.

Note that since PSTATE.tct is automatically set to 0 during entry to a trap handler, the execution of a RETRY instruction at the end of a trap handler will not cause a *control_transfer_instruction* exception unless trap handler software has explicitly set PSTATE.tct to 1. During execution of the RETRY instruction, the value of PSTATE.tct is restored from TSTATE.

RETRY

Programming Note | RETRY should *not* normally be used to return from the trap handler for the *control_transfer_instruction* exception itself. See the DONE instruction on page 147 and *Trap on Control Transfer (tct)* on page 70.

Exceptions *illegal_instruction*
 privileged_opcode
 control_transfer_instruction (impl. dep. #450-S20)

■ *See Also* DONE on page 147

RETURN

7.114 RETURN

Instruction	op3	Operation	Assembly Language Syntax	Class
RETURN	11 1001	Return	return <i>address</i>	A1



Description The RETURN instruction causes a register-indirect delayed transfer of control to the target address and has the window semantics of a RESTORE instruction; that is, it restores the register window prior to the last SAVE instruction. The target address is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + sign_ext(simm13)” if $i = 1$. Registers R[rs1] and R[rs2] come from the *old* window.

Like other DCTIs, all effects of RETURN (including modification of CWP) are visible prior to execution of the delay slot instruction.

Programming Note To reexecute the trapped instruction when returning from a user trap handler, use the RETURN instruction in the delay slot of a JMPL instruction, for example:

```

    jmp1    %l6,%g0    ! Trapped PC supplied to user trap handler
    return %l7        ! Trapped NPC supplied to user trap handler

```

Programming Note A routine that uses a register window may be structured either as:

```

    save    %sp,-framesize,%sp
    . . .
    ret     ! "ret" is shorthand for "jmp1 %i7 + 8,%g0"
    restore ! A useful instruction in the delay slot, such as
           ! "restore %o2,%l2,%o0"

```

or as:

```

    save    %sp,-framesize,%sp
    . . .
    return %i7 + 8
    nop    ! Instead of "nop", could do some useful work in the
           ! caller's window, for example, "or %o1,%o2,%o0"

```

An attempt to execute a RETURN instruction when bits 29:25 are nonzero causes an *illegal_instruction* exception.

An attempt to execute a RETURN instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

A RETURN instruction may cause a *window_fill* exception as part of its RESTORE semantics.

When PSTATE.am = 1, the more-significant 32 bits of the target instruction address are masked out (set to 0) before being sent to the memory system. However, if a *control_transfer_instruction* trap occurs, the full 64-bit (nonmasked) address of the RETURN instruction is written into TPC[TL].

A RETURN instruction causes a *mem_address_not_aligned* exception if either of the two least-significant bits of the target address is nonzero.

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20) and PSTATE.tct = 1, then RETURN generates a *control_transfer_instruction* exception instead of causing a control transfer.

RETURN


Exceptions

- illegal_instruction*
- fill_n_normal* ($n = 0-7$)
- fill_n_other* ($n = 0-7$)
- mem_address_not_aligned*
- control_transfer_instruction* (impl. dep. #450-S20)

SAVE

7.115 SAVE

Instruction	op3	Operation	Assembly Language Syntax	Class
SAVE	11 1100	Save Caller's Window	save <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1



Description The SAVE instruction provides the routine executing it with a new register window. The *out* registers from the old window become the *in* registers of the new window. The contents of the *out* and the *local* registers in the new window are zero or contain values from the executing process; that is, the process sees a clean window.

Furthermore, if and only if a spill trap is not generated, SAVE behaves like a normal ADD instruction, except that the source operands R[rs1] or R[rs2] are read from the *old* window (that is, the window addressed by the original CWP) and the sum is written into R[rd] of the *new* window (that is, the window addressed by the new CWP).

Note CWP arithmetic is performed modulo the number of implemented windows, *N_REG_WINDOWS*.

Programming Notes Typically, if a SAVE instruction traps, the spill trap handler returns to the trapped instruction to reexecute it. So, although the ADD operation is not performed the first time (when the instruction traps), it is performed the second time the instruction executes. The same applies to changing the CWP.

The SAVE instruction can be used to atomically allocate a new window in the register file and a new software stack frame in memory.

There is a performance trade-off to consider between using SAVE/RESTORE and saving and restoring selected registers explicitly.

Description (Effect on Privileged State)

If a SAVE instruction does not trap, it increments the CWP (**mod** *N_REG_WINDOWS*) to provide a new register window and updates the state of the register windows by decrementing CANSAVE and incrementing CANRESTORE.

If the new register window is occupied (that is, CANSAVE = 0), a spill trap is generated. The trap vector for the spill trap is based on the value of OTHERWIN and WSTATE. The spill trap handler is invoked with the CWP set to point to the window to be spilled (that is, old CWP + 2).

An attempt to execute a SAVE instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

SAVE

If `CANSAVE` \neq 0, the `SAVE` instruction checks whether the new window needs to be cleaned. It causes a `clean_window` trap if the number of unused clean windows is zero, that is, $(\text{CLEANWIN} - \text{CANRESTORE}) = 0$. The `clean_window` trap handler is invoked with the `CWP` set to point to the window to be cleaned (that is, $\text{old CWP} + 1$).

Programming Note	The vectoring of spill traps can be controlled by setting the value of the <code>OTHERWIN</code> and <code>WSTATE</code> registers appropriately. The spill handler normally will end with a <code>SAVED</code> instruction followed by a <code>RETRY</code> instruction.
-------------------------	--

Exceptions


- illegal_instruction*
- spill_n_normal* ($n = 0-7$)
- spill_n_other* ($n = 0-7$)
- clean_window*

■ *See Also* `RESTORE` on page 315

SAVED

7.116 SAVED

Instruction	Operation	Assembly Language Syntax	Class
SAVED ^P	Window has been saved	saved	A1



Description SAVED adjusts the state of the register-windows control registers.

SAVED increments CANSERVE. If OTHERWIN = 0, SAVED decrements CANRESTORE. If OTHERWIN ≠ 0, it decrements OTHERWIN.

Programming Notes Trap handler software for register window spills uses the SAVED instruction to indicate that a window has been spilled successfully.

Normal privileged software would probably not execute a SAVED instruction from trap level zero (TL = 0). However, it is not illegal to do so and doing so does not cause a trap.

Executing a SAVED instruction outside of a window spill trap handler is likely to create an inconsistent window state. Hardware will not signal an exception, however, since maintaining a consistent window state is the responsibility of privileged software.

If $(CANSERVE \geq (N_REG_WINDOWS - 2))$ or $((CANRESTORE = 0) \text{ and } (OTHERWIN = 0))$ just prior to execution of a SAVED instruction, the subsequent behavior of the processor is undefined. In neither of these cases can SAVED generate a register window state that is both valid (see *Register Window State Definition* on page 66) and consistent with the state prior to the SAVED.

An attempt to execute a SAVED instruction when instruction bits 18:0 are nonzero causes an *illegal_instruction* exception.

An attempt to execute a SAVED instruction in nonprivileged mode (PSTATE.priv = 0) causes a *privileged_opcode* exception.

Exceptions *illegal_instruction*
privileged_opcode


See Also ALLCLEAN on page 118
INVALW on page 237
NORMALW on page 295
OTHERW on page 297
RESTORED on page 317

SDIV, SDIVcc (Deprecated)

7.117 Signed Divide (64-bit ÷ 32-bit)

The SDIV and SDIVcc instructions are deprecated and should not be used in new software. The SDIVX instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
SDIV ^D	00 1111	Signed Integer Divide	sdiv <i>reg_rs1, reg_or_imm, reg_rd</i>	D3
SDIVcc ^D	01 1111	Signed Integer Divide and modify cc's	sdivcc <i>reg_rs1, reg_or_imm, reg_rd</i>	D3



Description The signed divide instructions perform 64-bit by 32-bit division, producing a 32-bit result. If $i = 0$, they compute “ $(Y :: R[rs1]\{31:0\}) \div R[rs2]\{31:0\}$ ”. Otherwise (that is, if $i = 1$), the divide instructions compute “ $(Y :: R[rs1]\{31:0\}) \div (\text{sign_ext}(\text{simm13})\{31:0\})$ ”. In either case, if overflow does not occur, the less significant 32 bits of the integer quotient are sign- or zero-extended to 64 bits and are written into $R[rd]$.

The contents of the Y register are undefined after any 64-bit by 32-bit integer divide operation.

Signed Divide Signed divide (SDIV, SDIVcc) assumes a signed integer doubleword dividend ($Y ::$ lower 32 bits of $R[rs1]$) and a signed integer word divisor (lower 32 bits of $R[rs2]$ or lower 32 bits of $\text{sign_ext}(\text{simm13})$) and computes a signed integer word quotient ($R[rd]$).

Signed division rounds an inexact quotient toward zero. For example, $-7 \div 4$ equals the rational quotient of -1.75 , which rounds to -1 (not -2) when rounding toward zero.

The result of a signed divide can overflow the low-order 32 bits of the destination register $R[rd]$ under certain conditions. When overflow occurs, the largest appropriate signed integer is returned as the quotient in $R[rd]$. The conditions under which overflow occurs and the value returned in $R[rd]$ under those conditions are specified in TABLE 7-19.

TABLE 7-19 SDIV / SDIVcc Overflow Detection and Value Returned

Condition Under Which Overflow Occurs	Value Returned in $R[rd]$
Rational quotient $\geq 2^{31}$	$2^{31} - 1$ (0000 0000 7FFF FFFF ₁₆)
Rational quotient $\leq -2^{31} - 1$	-2^{31} (FFFF FFFF 8000 0000 ₁₆)

When no overflow occurs, the 32-bit result is sign-extended to 64 bits and written into register $R[rd]$.

SDIV, SDIVcc (Deprecated)

SDIV does not affect the condition code bits. SDIVcc writes the integer condition code bits as shown in the following table. Note that negative (N) and zero (Z) are set according to the value of R[rd] after it has been set to reflect overflow, if any.

Bit	Effect on bit of SDIVcc instruction
icc.n	Set to 1 if R[rd]{31} = 1; otherwise, set to 0
icc.z	Set to 1 if R[rd]{31:0} = 0; otherwise, set to 0
icc.v	Set to 1 if overflow (<i>per</i> TABLE 7-12); otherwise set to 0
icc.c	Set to 0
xcc.n	Set to 1 if R[rd]{63} = 1; otherwise, set to 0
xcc.z	Set to 1 if R[rd]{63:0} = 0; otherwise, set to 0
xcc.v	Set to 0
xcc.c	Set to 0

An attempt to execute an SDIV or SDIVcc instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
 division_by_zero

See Also RDY on page 310
 UDIV[cc] on page 369

SETHI

7.118 SETHI

Instruction	op2	Operation	Assembly Language Syntax	Class
SETHI	100	Set High 22 Bits of Low Word	<code>sethi const22, reg_{rd}</code> <code>sethi %hi(value), reg_{rd}</code>	A1



Description SETHI zeroes the least significant 10 bits and the most significant 32 bits of R[rd] and replaces bits 31 through 10 of R[rd] with the value from its imm22 field.

SETHI does not affect the condition codes.

Some SETHI instructions with rd = 0 have special uses:

- rd = 0 and imm22 = 0: defined to be a NOP instruction (described in *No Operation*)
- rd = 0 and imm22 ≠ 0 may be used to trigger hardware performance counters in some Oracle SPARC Architecture implementations (for details, see implementation-specific documentation).

Programming Note

The most common form of 64-bit constant generation is creating stack offsets whose magnitude is less than 2³². The code below can be used to create the constant 0000 0000 ABCD 1234₁₆:

```
sethi %hi(0xabcd1234),%o0
or    %o0, 0x234, %o0
```

The following code shows how to create a negative constant. **Note:** The immediate field of the xor instruction is sign extended and can be used to place 1's in all of the upper 32 bits. For example, to set the negative constant FFFF FFFF ABCD 1234₁₆:

```
sethi %hi(0x5432edcb),%o0! note 0x5432EDCB, not 0xABCD1234
xor   %o0, 0x1e34, %o0! part of imm. overlaps upper bits
```

Exceptions None

SHA1, SHA256, SHA512

Programming Note Software *must* check that the corresponding capability bit in the CFR register (sha1, sha256, or sha512) = 1 before executing one of these instructions. If the corresponding capability bit in CFR = 0, then software should assume that an attempt to execute the hash instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.


Therefore, if the corresponding capability bit in CFR = 0, software should perform the hash operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

Exceptions *fp_disabled*

SIAM

7.120 Set Interval Arithmetic Mode VIS 2

Instruction	opf	Operation	Assembly Language Syntax	Class
SIAM	0 1000 0001	Set the interval arithmetic mode fields in the GSR	<code>siam</code> <code>siam_mode</code>	A1



Description The SIAM instruction sets the GSR.im and GSR.irnd fields as follows:

GSR.im ← mode{2}

GSR.irnd ← mode{1:0}

Note When GSR.im is set to 1, all subsequent floating-point instructions requiring round mode settings derive rounding-mode information from the General Status Register (GSR.irnd) instead of the Floating-Point State Register (FSR.rd).

Note When GSR.im = 1, the processor operates in standard floating-point mode regardless of the setting of FSR.ns.

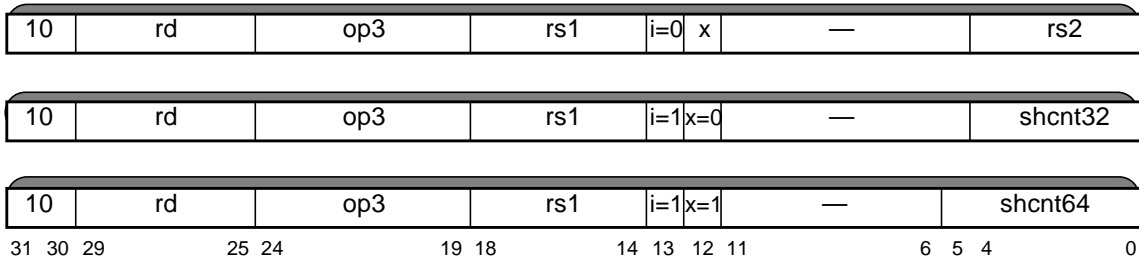
An attempt to execute a SIAM instruction when instruction bits 29:25, 18:14, or 4:3 are nonzero causes an *illegal_instruction* exception.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a SIAM instruction causes an *fp_disabled* exception.

Exceptions *illegal_instruction*
fp_disabled

7.121 Shift

Instruction	op3	x	Operation	Assembly Language Syntax	Class
SLL	10 0101	0	Shift Left Logical – 32 bits	<code>sll</code> <i>reg_{rs1}</i> , <i>reg_or_shcnt</i> , <i>reg_{rd}</i>	A1
SRL	10 0110	0	Shift Right Logical – 32 bits	<code>srl</code> <i>reg_{rs1}</i> , <i>reg_or_shcnt</i> , <i>reg_{rd}</i>	A1
SRA	10 0111	0	Shift Right Arithmetic – 32 bits	<code>sra</code> <i>reg_{rs1}</i> , <i>reg_or_shcnt</i> , <i>reg_{rd}</i>	A1
SLLX	10 0101	1	Shift Left Logical – 64 bits	<code>sllx</code> <i>reg_{rs1}</i> , <i>reg_or_shcnt</i> , <i>reg_{rd}</i>	A1
SRLX	10 0110	1	Shift Right Logical – 64 bits	<code>srlx</code> <i>reg_{rs1}</i> , <i>reg_or_shcnt</i> , <i>reg_{rd}</i>	A1
SRAX	10 0111	1	Shift Right Arithmetic – 64 bits	<code>srax</code> <i>reg_{rs1}</i> , <i>reg_or_shcnt</i> , <i>reg_{rd}</i>	A1



Description These instructions perform logical or arithmetic shift operations.

When $i = 0$ and $x = 0$, the shift count is the least significant five bits of $R[rs2]$.

When $i = 0$ and $x = 1$, the shift count is the least significant six bits of $R[rs2]$. When $i = 1$ and $x = 0$, the shift count is the immediate value specified in bits 0 through 4 of the instruction.

When $i = 1$ and $x = 1$, the shift count is the immediate value specified in bits 0 through 5 of the instruction.

TABLE 7-20 shows the shift count encodings for all values of i and x .

TABLE 7-20 Shift Count Encodings

i	x	Shift Count
0	0	bits 4–0 of $R[rs2]$
0	1	bits 5–0 of $R[rs2]$
1	0	bits 4–0 of instruction
1	1	bits 5–0 of instruction

SLL and SLLX shift all 64 bits of the value in $R[rs1]$ left by the number of bits specified by the shift count, replacing the vacated positions with zeroes, and write the shifted result to $R[rd]$.

SRL shifts the low 32 bits of the value in $R[rs1]$ right by the number of bits specified by the shift count. Zeroes are shifted into bit 31. The upper 32 bits are set to zero, and the result is written to $R[rd]$.

SRLX shifts all 64 bits of the value in $R[rs1]$ right by the number of bits specified by the shift count. Zeroes are shifted into the vacated high-order bit positions, and the shifted result is written to $R[rd]$.

SRA shifts the low 32 bits of the value in $R[rs1]$ right by the number of bits specified by the shift count and replaces the vacated positions with bit 31 of $R[rs1]$. The high-order 32 bits of the result are all set with bit 31 of $R[rs1]$, and the result is written to $R[rd]$.

SRAX shifts all 64 bits of the value in $R[rs1]$ right by the number of bits specified by the shift count and replaces the vacated positions with bit 63 of $R[rs1]$. The shifted result is written to $R[rd]$.

SLL / SRL / SRA

No shift occurs when the shift count is 0, but the high-order bits are affected by the 32-bit shifts as noted above.

These instructions do not modify the condition codes.

Programming Notes	“Arithmetic left shift by 1 (and calculate overflow)” can be effected with the ADDcc instruction. The instruction “sra <i>reg_{rs1}</i> , 0, <i>reg_{rd}</i> ” can be used to convert a 32-bit value to 64 bits, with sign extension into the upper word. “srl <i>reg_{rs1}</i> , 0, <i>reg_{rd}</i> ” can be used to clear the upper 32 bits of R[rd].
--------------------------	--

An attempt to execute a SLL, SRL, or SRA instruction when instruction bits 11:5 are nonzero causes an *illegal_instruction* exception.

An attempt to execute a SLLX, SRLX, or SRAX instruction when either of the following conditions exist causes an *illegal_instruction* exception:

- *i* = 0 or *x* = 0 and instruction bits 11:5 are nonzero
- *x* = 1 and instruction bits 11:6 are nonzero

Exceptions *illegal_instruction*

SMUL, SMULcc (Deprecated)

7.122 Signed Multiply (32-bit)

The SMUL and SMULcc instructions are deprecated and should not be used in new software. The MULX instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
SMUL ^D	00 1011	Signed Integer Multiply	smul <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	D3
SMULcc ^D	01 1011	Signed Integer Multiply and modify cc's	smulcc <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	D3



Description The signed multiply instructions perform 32-bit by 32-bit multiplications, producing 64-bit results. They compute “R[rs1]{31:0} × R[rs2]{31:0}” if *i* = 0, or “R[rs1]{31:0} × sign_ext(simm13){31:0}” if *i* = 1. They write the 32 most significant bits of the product into the Y register and all 64 bits of the product into R[rd].

Signed multiply instructions (SMUL, SMULcc) operate on signed integer word operands and compute a signed integer doubleword product.

SMUL does not affect the condition code bits. SMULcc writes the integer condition code bits, *icc* and *xcc*, as shown below.

Bit	Effect on bit by execution of SMULcc
icc.n	Set to 1 if product{31} = 1; otherwise, set to 0
icc.z	Set to 1 if product{31:0} = 0; otherwise, set to 0
icc.v	Set to 0
icc.c	Set to 0
xcc.n	Set to 1 if product{63} = 1; otherwise, set to 0
xcc.z	Set to 1 if product{63:0} = 0; otherwise, set to 0
xcc.v	Set to 0
xcc.c	Set to 0

Note 32-bit negative (*icc.n*) and zero (*icc.z*) condition codes are set according to the *less* significant word of the product, not according to the full 64-bit result.

Programming Notes 32-bit overflow after SMUL or SMULcc is indicated by $Y \neq (R[rd] \gg 31)$, where “ \gg ” indicates 32-bit arithmetic right-shift.

An attempt to execute a SMUL or SMULcc instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*

See Also UMUL[cc] on page 371

7.123 Store Integer

Instruction	op3	Operation	Assembly Language Syntax	Class
STB	00 0101	Store Byte	<code>stb[†] reg_{rd}, [address]</code>	A1
STH	00 0110	Store Halfword	<code>sth[‡] reg_{rd}, [address]</code>	A1
STW	00 0100	Store Word	<code>stw[◇] reg_{rd}, [address]</code>	A1
STX	00 1110	Store Extended Word	<code>stx reg_{rd}, [address]</code>	A1

[†] synonyms: `stwb, stsb` [‡] synonyms: `stuh, stsh` [◇] synonyms: `st, stuw, stsw`



Description The store integer instructions (except store doubleword) copy the whole extended (64-bit) integer, the less significant word, the least significant halfword, or the least significant byte of R[rd] into memory.

These instructions access memory using the implicit ASI (see page 83). The effective address for these instructions is “R[rs1] + R[rs2]” if *i* = 0, or “R[rs1] + **sign_ext**(simm13)” if *i* = 1.

A successful store (notably, STX) integer instruction operates atomically.

An attempt to execute a store integer instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Any of the following conditions cause a *mem_address_not_aligned* exception:

- an attempt to execute STH when the effective address is not halfword-aligned
- an attempt to execute STW when the effective address is not word-aligned
- an attempt to execute STX when the effective address is not doubleword-aligned


Exceptions *illegal_instruction*
mem_address_not_aligned
VA_watchpoint

See Also STTW on page 352

7.124 Store Integer into Alternate Space

Instruction	op3	Operation	Assembly Language Syntax	Class
STBA ^{PASI}	01 0101	Store Byte into Alternate Space	stba [†] <i>reg_{rd}</i> [<i>regaddr</i>] <i>imm_asi</i> stba <i>reg_{rd}</i> [<i>reg_plus_imm</i>] %asi	A1
STHA ^{PASI}	01 0110	Store Halfword into Alternate Space	stha [‡] <i>reg_{rd}</i> [<i>regaddr</i>] <i>imm_asi</i> stha <i>reg_{rd}</i> [<i>reg_plus_imm</i>] %asi	A1
STWA ^{PASI}	01 0100	Store Word into Alternate Space	stwa [◇] <i>reg_{rd}</i> [<i>regaddr</i>] <i>imm_asi</i> stwa <i>reg_{rd}</i> [<i>reg_plus_imm</i>] %asi	A1
STXA ^{PASI}	01 1110	Store Extended Word into Alternate Space	stxa <i>reg_{rd}</i> [<i>regaddr</i>] <i>imm_asi</i> stxa <i>reg_{rd}</i> [<i>reg_plus_imm</i>] %asi	A1

[†] *synonyms*: stuba, stsba [‡] *synonyms*: stuha, stsha [◇] *synonyms*: sta, stuwa, stswa



Description The store integer into alternate space instructions copy the whole extended (64-bit) integer, the less significant word, the least significant halfword, or the least significant byte of R[rd] into memory.

Store integer to alternate space instructions contain the address space identifier (ASI) to be used for the store in the *imm_asi* field if *i* = 0, or in the ASI register if *i* = 1. The effective address for these instructions is “R[rs1] + R[rs2]” if *i* = 0, or “R[rs1]+**sign_ext**(simm13)” if *i* = 1.

A successful store (notably, STXA) instruction operates atomically.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, these instructions cause a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, these instructions cause a *privileged_action* exception.

Any of the following conditions cause a *mem_address_not_aligned* exception:

- an attempt to execute STHA when the effective address is not halfword-aligned
- an attempt to execute STWA when the effective address is not word-aligned
- an attempt to execute STXA when the effective address is not doubleword-aligned

STBA, STHA, and STWA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with these instructions causes a *DAE_invalid_asi* exception.

ASIs valid for STBA, STHA, and STWA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
15 ₁₆ ASI_REAL_IO	1D ₁₆ ASI_REAL_IO_LITTLE
22 ₁₆ ASI_STBI_AIUP	2A ₁₆ ASI_STBI_AIUP_L
23 ₁₆ ASI_STBI_AIUS	2B ₁₆ ASI_STBI_AIUS_L
27 ₁₆ ASI_STBI_N	2F ₁₆ ASI_STBI_N_L
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE

STBA / STHA / STWA / STXA

ASIs valid for STBA, STHA, and STWA

E2 ₁₆ ASI_STBI_P	EA ₁₆ ASI_STBI_PL
E3 ₁₆ ASI_STBI_S	EB ₁₆ ASI_STBI_SL

STXA can be used with any ASI (including, but not limited to, the above two lists), unless it either (a) violates the privilege mode rules described for the *privileged_action* exception above or (b) is used with any of the following ASIs, which causes a *DAE_invalid_asi* exception.

ASIs invalid for STXA (cause *DAE_invalid_asi* exception)

16 ₁₆ ASI_BLOCK_AS_IF_USER_PRIMARY	1E ₁₆ ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE
17 ₁₆ ASI_BLOCK_AS_IF_USER_SECONDARY	1F ₁₆ ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE
22 ₁₆ (ASI_TWIX_AIUP)	2A ₁₆ (ASI_TWIX_AIUP_L)
23 ₁₆ (ASI_TWIX_AIUS)	2B ₁₆ (ASI_TWIX_AIUS_L)
26 ₁₆ (ASI_TWIX_REAL)	2E ₁₆ (ASI_TWIX_REAL_L)
27 ₁₆ (ASI_TWIX_N)	2F ₁₆ (ASI_TWIX_NL)
16 ₁₆ ASI_BLOCK_AS_IF_USER_PRIMARY	1E ₁₆ ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE
17 ₁₆ ASI_BLOCK_AS_IF_USER_SECONDARY	1F ₁₆ ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE
82 ₁₆ ASI_PRIMARY_NO_FAULT	8A ₁₆ ASI_PRIMARY_NO_FAULT_LITTLE
83 ₁₆ ASI_SECONDARY_NO_FAULT	8B ₁₆ ASI_SECONDARY_NO_FAULT_LITTLE
C0 ₁₆ ASI_PST8_PRIMARY	C8 ₁₆ ASI_PST8_PRIMARY_LITTLE
C1 ₁₆ ASI_PST8_SECONDARY	C9 ₁₆ ASI_PST8_SECONDARY_LITTLE
C2 ₁₆ ASI_PST16_PRIMARY	CA ₁₆ ASI_PST16_PRIMARY_LITTLE
C3 ₁₆ ASI_PST16_SECONDARY	CB ₁₆ ASI_PST16_SECONDARY_LITTLE
C4 ₁₆ ASI_PST32_PRIMARY	CC ₁₆ ASI_PST32_PRIMARY_LITTLE
C5 ₁₆ ASI_PST32_SECONDARY	CD ₁₆ ASI_PST32_SECONDARY_LITTLE
D0 ₁₆ ASI_FL8_PRIMARY	D8 ₁₆ ASI_FL8_PRIMARY_LITTLE
D1 ₁₆ ASI_FL8_SECONDARY	D9 ₁₆ ASI_FL8_SECONDARY_LITTLE
D2 ₁₆ ASI_FL16_PRIMARY	DA ₁₆ ASI_FL16_PRIMARY_LITTLE
D3 ₁₆ ASI_FL16_SECONDARY	DB ₁₆ ASI_FL16_SECONDARY_LITTLE
E0 ₁₆ ASI_BLOCK_COMMIT_PRIMARY	E1 ₁₆ ASI_BLOCK_COMMIT_SECONDARY
F0 ₁₆ ASI_BLOCK_PRIMARY	F8 ₁₆ ASI_BLOCK_PRIMARY_LITTLE
F1 ₁₆ ASI_BLOCK_SECONDARY	F9 ₁₆ ASI_BLOCK_SECONDARY_LITTLE

V8 Compatibility | The SPARC V8 STA instruction was renamed STWA in the
Note | SPARC V9 architecture.

Exceptions | *mem_address_not_aligned* (all except STBA)
privileged_action
VA_watchpoint
DAE_invalid_asi
DAE_privilege_violation
DAE_nfo_page

See Also | LDA on page 240
STWA on page 354


STBLOCKF (deprecated)

7.125 Block Store VIS 1

The STBLOCKF^D instructions are deprecated and should not be used in new software. A sequence of STDF instructions should be used instead.

The STBLOCKF^D instruction is intended to be a processor-specific instruction, which may or may not be implemented in future Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	ASI Value	Operation	Assembly Language Syntax	Class
STBLOCKF ^D 16 ₁₆		64-byte block store to primary address space, user privilege	<i>stda freg_{rd}, [regaddr] #ASI_BLK_AIUP</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D 17 ₁₆		64-byte block store to secondary address space, user privilege	<i>stda freg_{rd}, [regaddr] #ASI_BLK_AIUS</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D 1E ₁₆		64-byte block store to primary address space, little-endian, user privilege	<i>stda freg_{rd}, [regaddr] #ASI_BLK_AIUPL</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D 1F ₁₆		64-byte block store to secondary address space, little-endian, user privilege	<i>stda freg_{rd}, [regaddr] #ASI_BLK_AIUSL</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D F0 ₁₆		64-byte block store to primary address space	<i>stda freg_{rd}, [regaddr] #ASI_BLK_P</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D F1 ₁₆		64-byte block store to secondary address space	<i>stda freg_{rd}, [regaddr] #ASI_BLK_S</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D F8 ₁₆		64-byte block store to primary address space, little-endian	<i>stda freg_{rd}, [regaddr] #ASI_BLK_PL</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D F9 ₁₆		64-byte block store to secondary address space, little-endian	<i>stda freg_{rd}, [regaddr] #ASI_BLK_SL</i>	A1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D2
STBLOCKF ^D E0 ₁₆		64-byte block commit store to primary address space	<i>stda freg_{rd}, [regaddr] #ASI_BLK_COMMIT_P</i>	B1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D3
STBLOCKF ^D E1 ₁₆		64-byte block commit store to secondary address space	<i>stda freg_{rd}, [regaddr] #ASI_BLK_COMMIT_S</i>	B1
			<i>stda freg_{rd}, [reg_plus_imm] %asi</i>	D3



Description A block store instruction references one of several special block-transfer ASIs. Block-transfer ASIs allow block stores to be performed accessing the same address space as normal stores. Little-endian ASIs (those with an 'L' suffix) access data in little-endian format; otherwise, the access is assumed to be big-endian. Byte swapping is performed separately for each of the eight double-precision registers accessed by the instruction.

Programming Note The block store instruction, STBLOCKF^D, and its companion, LDBLOCKF^D, were originally defined to provide a fast mechanism for block-copy operations. However, in modern implementations they are rarely much faster than a sequence of regular loads and stores, so are now deprecated.

STBLOCKF (deprecated)

STBLOCKF^D stores data from the eight double-precision floating-point registers specified by *rd* to a 64-byte-aligned memory area. The lowest-addressed eight bytes in memory are stored from the lowest-numbered double-precision *rd*.

While a STBLOCKF^D operation is in progress, any of the following values may be observed in a destination doubleword memory locations: (1) the old data value, (2) zero, or (3) the new data value. When the operation is complete, only the new data values will be seen.

Compatibility Note	Software written for older UltraSPARC implementations that reads data being written by STBLOCKF ^D instructions may or may not allow for case (2) above. Such software should be checked to verify that either it always waits for STBLOCKF ^D to complete before reading the values written, or that it will operate correctly if an intermediate value of zero (not the “old” or “new” data values) is observed while the STBLOCKF ^D operation is in progress.
---------------------------	---

A Block Store only guarantees atomicity for each 64-bit (8-byte) portion of the 64 bytes that it stores.

A Block Store with Commit forces the data to be written to memory and invalidates copies in all caches present. As a result, a Block Store with Commit maintains coherency with the I-cache¹. It does not, however, flush instructions that have already been fetched into the pipeline before executing the modified code. If a Block Store with Commit is used to write modified instructions, a FLUSH instruction must still be executed to guarantee that the instruction pipeline is flushed. (See *Synchronizing Instruction and Data Memory* on page 420 for more information.)

ASIs E0₁₆ and E1₁₆ are only used for block store-with-commit operations; they are not available for use by block load operations. See *Block Load and Store ASIs* on page 436 for more information.

Software should assume the following (where “load operation” includes load, load-store, and LDBLOCKF^D instructions and “store operation” includes store, load-store, and STBLOCKF^D instructions):

- A STBLOCKF^D does not follow memory ordering with respect to earlier or later load operations. If there is overlap between the addresses of destination memory locations of a STBLOCKF^D and the source address of a later load operation, the load operation may receive incorrect data. Therefore, if ordering with respect to later load operations is important, a MEMBAR #StoreLoad instruction must be executed between the STBLOCKF^D and subsequent load operations.
- A STBLOCKF^D does not follow memory ordering with respect to earlier or later store operations. Those instructions’ data may commit to memory in a different order from the one in which those instructions were issued. Therefore, if ordering with respect to later store operations is important, a MEMBAR #StoreStore instruction must be executed between the STBLOCKF^D and subsequent store operations.
- STBLOCKFs do not follow register dependency interlocks, as do ordinary stores.

Programming Note	STBLOCKF ^D is intended to be a processor-specific instruction (see the warning at the top of page 337). If STBLOCKF ^D <i>must</i> be used in software intended to be portable across current and previous processor implementations, then it must be coded to work in the face of any implementation variation that is permitted by implementation dependency #411-S10, described below.
-------------------------	--

IMPL. DEP. #411-S10: The following aspects of the behavior of the block store (STBLOCKF^D) instruction are implementation dependent:

- The memory ordering model that STBLOCKF^D follows (other than as constrained by the rules outlined above).
- Whether *VA_watchpoint* exceptions are recognized on accesses to all 64 bytes of the STBLOCKF^D (the recommended behavior), or only on accesses to the first eight bytes.

¹: Even if all data stores on a given implementation coherently update the instruction cache (see page 389), stores (other than Block Store with Commit) on SPARC V9 implementations in general do *not* maintain coherency between instruction and data caches.

STBLOCKF (deprecated)

- Whether STBLOCKFs to non-cacheable (TTE.cp = 0) pages execute in strict program order or not. If not, a STBLOCKF^D to a non-cacheable page causes an *illegal_instruction* exception.
- Whether STBLOCKF^D follows register dependency interlocks (as ordinary stores do).
- Whether a non-Commit STBLOCKF^D forces the data to be written to memory and invalidates copies in all caches present (as the Commit variants of STBLOCKF^D do).
- Any other restrictions on the behavior of STBLOCKF^D, as described in implementation-specific documentation.

Exceptions. An *illegal_instruction* exception occurs if the source floating-point registers are not aligned on an eight-register boundary.

If the FPU is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if no FPU is present, an attempt to execute a STBLOCKF^D instruction causes an *fp_disabled* exception.

If the least significant 6 bits of the memory address are not all zero, a *mem_address_not_aligned* exception occurs.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0 (ASIs 16₁₆, 17₁₆, 1E₁₆, and 1F₁₆), STBLOCKF^D causes a *privileged_action* exception.

An access caused by STBLOCKF^D may trigger a *VA_watchpoint* exception (impl. dep. #411-S10).

Implementation Note	STBLOCKF ^D shares an opcode with the STDFA, STPARTIALF, and STSHORTF instructions; they are distinguished by the ASI used.
----------------------------	---

Exceptions

- illegal_instruction*
- fp_disabled*
- mem_address_not_aligned*
- privileged_action*
- VA_watchpoint* (impl. dep. #411-S10)
- DAE_privilege_violation*
- DAE_nfo_page*

See Also

- LDBLOCKF^D on page 243
- STDF on page 340

7.126 Store Floating-Point

Instruction	op3	rd	Operation	Assembly Language		Class
STF	10 0100	0–31	Store Floating-Point register	st	<i>freg_{rd}</i> , [<i>address</i>]	A1
STDF	10 0111	†	Store Double Floating-Point register	std	<i>freg_{rd}</i> , [<i>address</i>]	A1
STQF	10 0110	†	Store Quad Floating-Point register	stq	<i>freg_{rd}</i> , [<i>address</i>]	C3

† Encoded floating-point register value, as described on page 51.



Description The store single floating-point instruction (STF) copies the contents of the 32-bit floating-point register $F_S[rd]$ into memory.

The store double floating-point instruction (STDF) copies the contents of 64-bit floating-point register $F_D[rd]$ into a word-aligned doubleword in memory. The unit of atomicity for STDF is 4 bytes (one word).

The store quad floating-point instruction (STQF) copies the contents of 128-bit floating-point register $F_Q[rd]$ into a word-aligned quadword in memory. The unit of atomicity for STQF is 4 bytes (one word).

These instructions access memory using the implicit ASI (see page 83). The effective address for these instructions is “ $R[rs1] + R[rs2]$ ” if $i = 0$, or “ $R[rs1] + \text{sign_ext}(\text{simm13})$ ” if $i = 1$.

Exceptions. An attempt to execute a STF or STDF instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the floating-point unit is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if the FPU is not present, then an attempt to execute a STF or STDF instruction causes an *fp_disabled* exception.

Any of the following conditions cause an exception:

- an attempt to execute STF, STDF, or STQF when the effective address is not word-aligned causes a *mem_address_not_aligned* exception
- an attempt to execute STDF when the effective address is word-aligned but not doubleword-aligned causes an *STDF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the STDF instruction and return (impl. dep. #110-V9-Cs10(a)).
- an attempt to execute STQF when the effective address is word-aligned but not quadword-aligned causes an *STQF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the STQF instruction and return (impl. dep. #112-V9-Cs10(a)).

Programming Note Some compilers issued sequences of single-precision stores for SPARC V8 processor targets when the compiler could not determine whether doubleword or quadword operands were properly aligned. For SPARC V9, since emulation of misaligned stores is expected to be fast, compilers should issue sets of single-precision stores only when they can determine that double- or quadword operands are *not* properly aligned.

STF / STDF / STQF

An attempt to execute an STQF instruction when $rd\{1\} \neq 0$ causes an *fp_exception_other* (FSR.ftt = *invalid_fp_register*) exception.

Implementation Note	Since Oracle SPARC Architecture 2015 processors do not implement in hardware instructions (including STQF) that refer to quad-precision floating-point registers, the <i>STQF_mem_address_not_aligned</i> and <i>fp_exception_other</i> (with FSR.ftt = <i>invalid_fp_register</i>) exceptions do not occur in hardware. However, their effects must be emulated by software when the instruction causes an <i>illegal_instruction</i> exception and subsequent trap.
----------------------------	--

Exceptions

- illegal_instruction*
- fp_disabled*
- STDF_mem_address_not_aligned*
- STQF_mem_address_not_aligned* (not used in Oracle SPARC Architecture 2015)
- mem_address_not_aligned*
- fp_exception_other* (FSR.ftt = *invalid_fp_register* (STQF only))
- VA_watchpoint*
- DAE_privilege_violation*
- DAE_nfo_page*


See Also

- Load Floating-Point Register on page 246*
- Block Store on page 337*
- Store Floating-Point into Alternate Space on page 342*
- Store Floating-Point State Register (Lower) on page 345*
- Store Short Floating-Point on page 350*
- DAE_invalid_asid eDAE_invalid_asid eStore Partial Floating-Point on page 347*
- Store Floating-Point State Register on page 356*

7.127 Store Floating-Point into Alternate Space

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
STFA ^{PASI}	11 0100	0–31	Store Floating-Point Register to Alternate Space	sta <i>freq_{rd}</i> , [<i>regaddr</i>] <i>imm_asi</i> sta <i>freq_{rd}</i> , [<i>reg_plus_imm</i>] %asi	A1
STDFA ^{PASI}	11 0111	†	Store Double Floating-Point Register to Alternate Space	stda <i>freq_{rd}</i> , [<i>regaddr</i>] <i>imm_asi</i> stda <i>freq_{rd}</i> , [<i>reg_plus_imm</i>] %asi	A1
STQFA ^{PASI}	11 0110	†	Store Quad Floating-Point Register to Alternate Space	stqa <i>freq_{rd}</i> , [<i>regaddr</i>] <i>imm_asi</i> stqa <i>freq_{rd}</i> , [<i>reg_plus_imm</i>] %asi	C3

† Encoded floating-point register value, as described on page 51.



Description The store single floating-point into alternate space instruction (STFA) copies the contents of the 32-bit floating-point register $F_S[rd]$ into memory.

The store double floating-point into alternate space instruction (STDFA) copies the contents of 64-bit floating-point register $F_D[rd]$ into a word-aligned doubleword in memory. The unit of atomicity for STDFA is 4 bytes (one word).

The store quad floating-point into alternate space instruction (STQFA) copies the contents of 128-bit floating-point register $F_Q[rd]$ into a word-aligned quadword in memory. The unit of atomicity for STQFA is 4 bytes (one word).

Store floating-point into alternate space instructions contain the address space identifier (ASI) to be used for the load in the *imm_asi* field if $i = 0$ or in the ASI register if $i = 1$. The effective address for these instructions is “ $R[rs1] + R[rs2]$ ” if $i = 0$, or “ $R[rs1] + \text{sign_ext}(\text{simm13})$ ” if $i = 1$.

Programming Note Some compilers issued sequences of single-precision stores for SPARC V8 processor targets when the compiler could not determine whether doubleword or quadword operands were properly aligned. For SPARC V9, since emulation of misaligned stores is expected to be fast, compilers should issue sets of single-precision stores only when they can determine that double- or quadword operands are *not* properly aligned.

Exceptions. If the floating-point unit is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if the FPU is not present, then an attempt to execute a STFA or STDFA instruction causes an *fp_disabled* exception.

Any of the following conditions cause an exception:

- an attempt to execute STFA, STDFA, or STQFA when the effective address is not word-aligned causes a *mem_address_not_aligned* exception
- an attempt to execute STDFA when the effective address is word-aligned but not doubleword-aligned causes an *STDF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the STDFA instruction and return (impl. dep. #110-V9-Cs10(a)).
- an attempt to execute STQFA when the effective address is word-aligned but not quadword-aligned causes an *STQF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the STQFA instruction and return (impl. dep. #112-V9-Cs10(a)).

STFA / STDFA / STQFA

Implementation Note STDFA shares an opcode with the STBLOCKF^D, STPARTIALF, and STSHORTF instructions; they are distinguished by the ASI used.

An attempt to execute an STQFA instruction when rd{1} ≠ 0 causes an *fp_exception_other* (FSR.ftt = invalid_fp_register) exception.

Implementation Note Since Oracle SPARC Architecture 2015 processors do not implement in hardware instructions (including STQFA) that refer to quad-precision floating-point registers, the *STQF_mem_address_not_aligned* and *fp_exception_other* (with FSR.ftt = invalid_fp_register) exceptions do not occur in hardware. However, their effects must be emulated by software when the instruction causes an *illegal_instruction* exception and subsequent trap.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, this instruction causes a *privileged_action* exception.

STFA and STQFA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with these instructions causes a *DAE_invalid_asi* exception.

ASIs valid for STFA and STQFA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
15 ₁₆ ASI_REAL_IO	1D ₁₆ ASI_REAL_IO_LITTLE
22 ₁₆ ASI_STBI_AIUP	2A ₁₆ ASI_STBI_AIUP_L
23 ₁₆ ASI_STBI_AIUS	2B ₁₆ ASI_STBI_AIUS_L
27 ₁₆ ASI_STBI_N	2F ₁₆ ASI_STBI_N_L
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE
E2 ₁₆ ASI_STBI_P	EA ₁₆ ASI_STBI_PL
E3 ₁₆ ASI_STBI_S	EB ₁₆ ASI_STBI_SL

STDFA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with the STDFA instruction causes a *DAE_invalid_asi* exception.

ASIs valid for STDFA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
15 ₁₆ ASI_REAL_IO	1D ₁₆ ASI_REAL_IO_LITTLE
16 ₁₆ ASI_BLOCK_AS_IF_USER_PRIMARY †	1E ₁₆ ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE †
17 ₁₆ ASI_BLOCK_AS_IF_USER_SECONDARY †	1F ₁₆ ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE †
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE
C0 ₁₆ ASI_PST8_PRIMARY *	C8 ₁₆ ASI_PST8_PRIMARY_LITTLE *
C1 ₁₆ ASI_PST8_SECONDARY *	C9 ₁₆ ASI_PST8_SECONDARY_LITTLE *
C2 ₁₆ ASI_PST16_PRIMARY *	CA ₁₆ ASI_PST16_PRIMARY_LITTLE *
C3 ₁₆ ASI_PST16_SECONDARY *	CB ₁₆ ASI_PST16_SECONDARY_LITTLE *

STFA / STDFA / STQFA

ASIs valid for STDFA

C4 ₁₆ ASI_PST32_PRIMARY *	CC ₁₆ ASI_PST32_PRIMARY_LITTLE *
C5 ₁₆ ASI_PST32_SECONDARY *	CD ₁₆ ASI_PST32_SECONDARY_LITTLE *
D0 ₁₆ ASI_FL8_PRIMARY ‡	D8 ₁₆ ASI_FL8_PRIMARY_LITTLE ‡
D1 ₁₆ ASI_FL8_SECONDARY ‡	D9 ₁₆ ASI_FL8_SECONDARY_LITTLE ‡
D2 ₁₆ ASI_FL16_PRIMARY ‡	DA ₁₆ ASI_FL16_PRIMARY_LITTLE ‡
D3 ₁₆ ASI_FL16_SECONDARY ‡	DB ₁₆ ASI_FL16_SECONDARY_LITTLE ‡
E0 ₁₆ ASI_BLOCK_COMMIT_PRIMARY †	E1 ₁₆ ASI_BLOCK_COMMIT_SECONDARY †
F0 ₁₆ ASI_BLOCK_PRIMARY †	F8 ₁₆ ASI_BLOCK_PRIMARY_LITTLE †
F1 ₁₆ ASI_BLOCK_SECONDARY †	F9 ₁₆ ASI_BLOCK_SECONDARY_LITTLE †

† If this ASI is used with the opcode for STDFA, the STBLOCKF^D instruction is executed instead of STFA. For behavior of STBLOCKF^D, see *Block Store* on page 337.

‡ If this ASI is used with the opcode for STDFA, the STSHORTF instruction is executed instead of STDFA. For behavior of STSHORTF, see *Store Short Floating-Point* on page 350.

* If this ASI is used with the opcode for STDFA, the STPARTIALF instruction is executed instead of STDFA. For behavior of STPARTIALF, see *DAE_invalid_asi eDAE_invalid_asi eStore Partial Floating-Point* on page 347.

Exceptions

fp_disabled
STDF_mem_address_not_aligned
STQF_mem_address_not_aligned (STQFA only) (not used in UA-2015)
mem_address_not_aligned
fp_exception_other (FSR.ftt = invalid_fp_register (STQFA only))
privileged_action
VA_watchpoint
DAE_invalid_asi
DAE_privilege_violation
DAE_nfo_page

See Also


Load Floating-Point from Alternate Space on page 248
Block Store on page 337
Store Floating-Point on page 340
Store Short Floating-Point on page 350
DAE_invalid_asi eDAE_invalid_asi eStore Partial Floating-Point on page 347

STFSR (Deprecated)

7.128 Store Floating-Point State Register (Lower)

The STFSR instruction is deprecated and should not be used in new software. The STXFSR instruction should be used instead.

Opcode	op3	rd	Operation	Assembly Language Syntax	Class
STFSR ^D	10 0101	0	Store Floating-Point State Register (Lower)	st %f _{sr} , [address]	D2
	10 0101	1-31	(see page 356)		



Description The Store Floating-point State Register (Lower) instruction (STFSR) waits for any currently executing FPop instructions to complete, and then it writes the less-significant 32 bits of FSR into memory.

After writing the FSR to memory, STFSR zeroes FSR.ftt

V9 Compatibility Note FSR.ftt should not be zeroed until it is known that the store will not cause a precise trap.

STFSR accesses memory using the implicit ASI (see page 83). The effective address for this instruction is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + sign_ext(sim13)” if $i = 1$.

An attempt to execute a STFSR instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the floating-point unit is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if the FPU is not present, then an attempt to execute a STFSR instruction causes an *fp_disabled* exception.

STFSR causes a *mem_address_not_aligned* exception if the effective memory address is not word-aligned.

V9 Compatibility Note Although STFSR is deprecated, Oracle SPARC Architecture implementations continue to support it for compatibility with existing SPARC V8 software. The STFSR instruction is defined to store only the less-significant 32 bits of the FSR into memory, while STXFSR allows SPARC V9 software to store all 64 bits of the FSR.

Implementation Note STFSR shares an opcode with the STXFSR instruction (and possibly with other implementation-dependent instructions); they are differentiated by the instruction rd field. An attempt to execute the $op = 10_2$, $op3 = 10\ 0101_2$ opcode with an invalid rd value causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
fp_disabled
mem_address_not_aligned
VA_watchpoint
DAE_privilege_violation
DAE_nfo_page

STFSR (Deprecated)

See Also

Store Floating-Point on page 340

Store Floating-Point State Register on page 356


STPARTIALF

7.129

DAE_invalid_asi DAE_invalid_asi Store Partial Floating-Point vis 1

Instruction	ASI Value	Operation	Assembly Language Syntax †	Class
STPARTIALF	C0 ₁₆	Eight 8-bit conditional stores to primary address space	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST8_P	B1
STPARTIALF	C1 ₁₆	Eight 8-bit conditional stores to secondary address space	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST8_S	B1
STPARTIALF	C8 ₁₆	Eight 8-bit conditional stores to primary address space, little-endian	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST8_PL	B1
STPARTIALF	C9 ₁₆	Eight 8-bit conditional stores to secondary address space, little-endian	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST8_SL	B1
STPARTIALF	C2 ₁₆	Four 16-bit conditional stores to primary address space	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST16_P	B1
STPARTIALF	C3 ₁₆	Four 16-bit conditional stores to secondary address space	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST16_S	B1
STPARTIALF	CA ₁₆	Four 16-bit conditional stores to primary address space, little-endian	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST16_PL	B1
STPARTIALF	CB ₁₆	Four 16-bit conditional stores to secondary address space, little-endian	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST16_SL	B1
STPARTIALF	C4 ₁₆	Two 32-bit conditional stores to primary address space	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST32_P	B1
STPARTIALF	C5 ₁₆	Two 32-bit conditional stores to secondary address space	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST32_S	B1
STPARTIALF	CC ₁₆	Two 32-bit conditional stores to primary address space, little-endian	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST32_PL	B1
STPARTIALF	CD ₁₆	Two 32-bit conditional stores to secondary address space, little-endian	stda freg _{rd} , reg _{rs2} , [reg _{rs1}] #ASI_PST32_SL	B1

† The original assembly language syntax for a Partial Store instruction (“stda freg_{rd}, [reg_{rs1}] reg_{rs2}, imm_asi”) has been deprecated because of inconsistency with the rest of the SPARC assembly language. Over time, assemblers will support the new syntax for this instruction. In the meantime, some existing assemblers may only recognize the original syntax.



Description The partial store instructions are selected by one of the partial store ASIs with the STDFA instruction.

Two 32-bit, four 16-bit, or eight 8-bit values from the 64-bit floating-point register $F_D[rd]$ are conditionally stored at the address specified by $R[rs1]$, using the mask specified in $R[rs2]$. STPARTIALF has the effect of merging selected data from its source register, $F_D[rd]$, into the existing data at the corresponding destination memory locations.

The mask value in $R[rs2]$ has the same format as the result specified by the edge-handling and pixel compare instructions (see *Edge Handling Instructions (no CC)* on page 150, *Partitioned Signed Compare* on page 207, and *Partitioned Unsigned Compare* on page 210). The most significant bit of the mask (not

STPARTIALF

of the entire register) corresponds to the most significant part of $F_D[rd]$. The data is stored in little-endian form in memory if the ASI name has an "L" (or "_LITTLE") suffix; otherwise, it is stored in big-endian format.

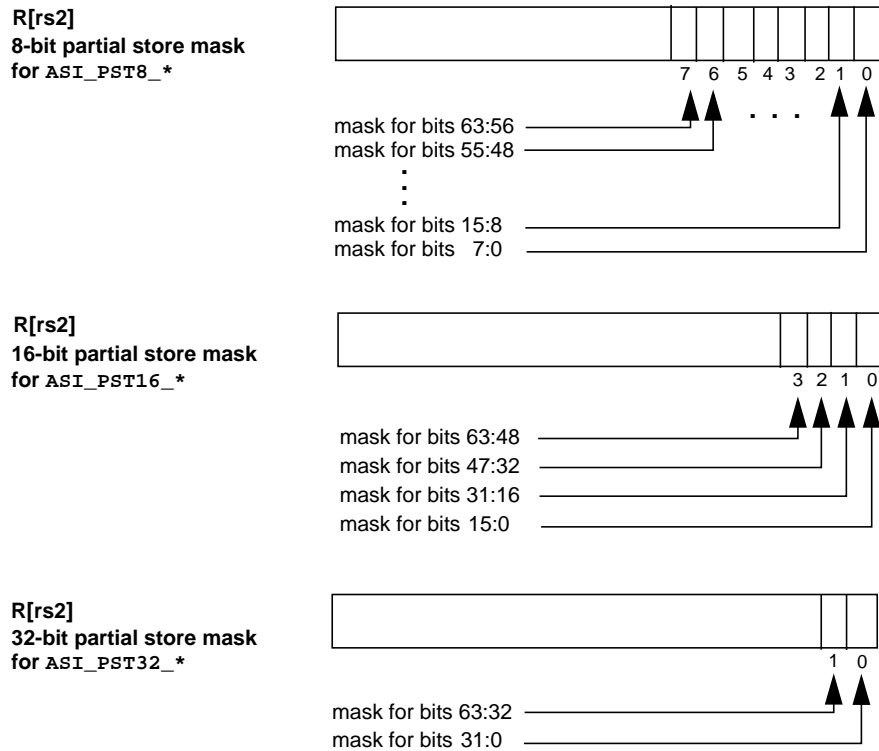


FIGURE 7-59 Mask Format for Partial Store

Exceptions.

IMPL. DEP. #_: It is implementation-dependent whether any of the following exceptions are suppressed when the store-mask in $R[rs2] = 0$. In particular, a store mask with value 0 may or may not prevent a *mem_address_not_aligned* exception or data access exception from occurring, if the conditions for triggering such an exception are otherwise met.

A Partial Store instruction can cause a virtual watchpoint exception when the following conditions are met:

- The virtual address in $R[rs1]$ matches the address in the VA Data Watchpoint Register.
- The byte store mask in $R[rs2]$ indicates that a byte, halfword or word is to be stored.
- The Virtual (Physical) Data Watchpoint Mask in `ASI_DCU_WATCHPOINT_CONTROL_REG` indicates that one or more of the bytes to be stored at the watched address is being watched.

For data watchpoints of partial stores in Oracle SPARC Architecture 2015, the byte store mask ($R[rs2]$) in the Partial Store instruction is ignored, and a watchpoint exception can occur even if the mask is zero (that is, no data will be written to memory). The `ASI_DCU_WATCHPOINT_CONTROL_REG` Data Watchpoint masks are only checked for nonzero value (watchpoint enabled) (impl. dep. #249).

An attempt to execute a STPARTIALF instruction when $i = 1$ causes an *illegal_instruction* exception.

If the floating-point unit is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if the FPU is not present, then an attempt to execute a STPARTIALF instruction causes an *fp_disabled* exception.

Any of the following conditions cause an exception:

- an attempt to execute STPARTIALF when the effective address is not word-aligned causes a *mem_address_not_aligned* exception

STPARTIALF

- an attempt to execute STPARTIALF when the effective address is word-aligned but not doubleword-aligned causes an *STDF_mem_address_not_aligned* exception. In this case, trap handler software must emulate the STPARTIALF instruction and return.

IMPL. DEP. #249-U3-Cs10: For an STPARTIAL instruction, the following aspects of data watchpoints are implementation dependent: (a) whether data watchpoint logic examines the byte store mask in R[rs2] or it conservatively behaves as if every Partial Store always stores all 8 bytes, and (b) whether data watchpoint logic examines individual bits in the Virtual (Physical) Data Watchpoint Mask in the LSU Control register DCUCR to determine which bytes are being watched or (when the Watchpoint Mask is nonzero) it conservatively behaves as if all 8 bytes are being watched.

ASIs C0₁₆–C5₁₆ and C8₁₆–CD₁₆ are only used for partial store operations. In particular, they should not be used with the LDDFA instruction; however, if any of them *is* used, the resulting behavior is specified in the LDDFA instruction description on page 250.

Implementation Note	STPARTIALF shares an opcode with the STBLOCKF ^D , STDFA, and STSHORTF instructions; they are distinguished by the ASI used.
----------------------------	--


Exceptions

- illegal_instruction*
- fp_disabled*
- mem_address_not_aligned*
- VA_watchpoint* (see text)
- DAE_privilege_violation*
- DAE_nc_page*
- DAE_nfo_page*

STSHORTF

7.130 Store Short Floating-Point VIS 1

Instruction	ASI Value	Operation	Assembly Language Syntax	Class
STSHORTF	D0 ₁₆	8-bit store to primary address space	<code>stda freg_{rd}, [regaddr] #ASI_FL8_P</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	D1 ₁₆	8-bit store to secondary address space	<code>stda freg_{rd}, [regaddr] #ASI_FL8_S</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	D8 ₁₆	8-bit store to primary address space, little-endian	<code>stda freg_{rd}, [regaddr] #ASI_FL8_PL</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	D9 ₁₆	8-bit store to secondary address space, little-endian	<code>stda freg_{rd}, [regaddr] #ASI_FL8_SL</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	D2 ₁₆	16-bit store to primary address space	<code>stda freg_{rd}, [regaddr] #ASI_FL16_P</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	D3 ₁₆	16-bit store to secondary address space	<code>stda freg_{rd}, [regaddr] #ASI_FL16_S</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	DA ₁₆	16-bit store to primary address space, little-endian	<code>stda freg_{rd}, [regaddr] #ASI_FL16_PL</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2
STSHORTF	DB ₁₆	16-bit store to secondary address space, little-endian	<code>stda freg_{rd}, [regaddr] #ASI_FL16_SL</code>	B1
			<code>stda freg_{rd}, [reg_plus_imm] %asi</code>	D2



Description The short floating-point store instruction allows 8- and 16-bit stores to be performed from the floating-point registers. Short stores access the low-order 8 or 16 bits of the register.

Little-endian ASIs transfer data in little-endian format from memory; otherwise, memory is assumed to be big-endian. Short floating-point stores are typically used with the FALIGNDATAg or FALIGNDATAi instruction to assemble or store 64 bits on noncontiguous components.

Implementation Note STSHORTF shares an opcode with the STBLOCKF^D, STDFA, and STPARTIALF instructions; they are distinguished by the ASI used.

If the floating-point unit is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if the FPU is not present, then an attempt to execute a STSHORTF instruction causes an *fp_disabled* exception.

An 8-bit STSHORTF (using ASI D0₁₆, D1₁₆, D8₁₆, or D9₁₆) can be performed to an arbitrary memory address (no alignment requirement).

An attempt to execute a 16-bit STSHORTF (using ASI D2₁₆, D3₁₆, DA₁₆, or DB₁₆) when the effective address is not halfword-aligned causes a *mem_address_not_aligned* exception.

Exceptions

- fp_disabled*
- mem_address_not_aligned*
- VA_watchpoint*
- DAE_privilege_violation*
- DAE_nfo_page*

STSHORTF

See Also

LDSHORTF on page 253

Align Data (using gsr.align) on page 154

Align Data (using Integer register) on page 155

STTW (Deprecated)

7.131 Store Integer Twin Word

The STTW instruction is deprecated and should not be used in new software. The STX instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax †	Class
STTW ^D	00 0111	Store Integer Twin Word	<i>sttw</i> <i>reg_{rd}</i> , [<i>address</i>]	D2

† The original assembly language syntax for this instruction used an “std” instruction mnemonic, which is now deprecated. Over time, assemblers will support the new “sttw” mnemonic for this instruction. In the meantime, some existing assemblers may only recognize the original “std” mnemonic.



Description The store integer twin word instruction (STTW) copies two words from an R register pair into memory. The least significant 32 bits of the even-numbered R register are written into memory at the effective address, and the least significant 32 bits of the following odd-numbered R register are written into memory at the “effective address + 4”.

The least significant bit of the *rd* field of a store twin word instruction is unused and should always be set to 0 by software.

STTW accesses memory using the implicit ASI (see page 83). The effective address for this instruction is “R[rs1] + R[rs2]” if *i* = 0, or “R[rs1] + **sign_ext**(*simm13*)” if *i* = 1.

A successful store twin word instruction operates atomically.

IMPL. DEP. #108-V9a: It is implementation dependent whether STTW is implemented in hardware. If not, an attempt to execute it will cause an *unimplemented_STTW* exception. (STTW is implemented in hardware in all Oracle SPARC Architecture 2015 implementations.)

An attempt to execute an STTW instruction when either of the following conditions exist causes an *illegal_instruction* exception:

- destination register number *rd* is an odd number (is misaligned)
- *i* = 0 and instruction bits 12:5 are nonzero

STTW causes a *mem_address_not_aligned* exception if the effective address is not doubleword-aligned.

With respect to little-endian memory, an STTW instruction behaves as if it is composed of two 32-bit stores, each of which is byte-swapped independently before being written into its respective destination memory word.

STTW (Deprecated)

Programming Notes	<p>STTW is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties. Therefore, software should avoid using STTW.</p> <p>If STTW is emulated in software, an STX instruction should be used for the memory access in the emulation code to preserve atomicity. Emulation software should examine TSTATE[TL].pstate.cle (and, if appropriate, TTE.ie) to determine the endianness of the emulated memory access.</p> <p>Note that the value of TTE.ie is not saved during a trap. Therefore, if it is examined in the emulation trap handler, that should be done as quickly as possible, to minimize the window of time during which the value of TTE.ie could possibly be changed from the value it had at the time of the attempted execution of STTW.</p>
--------------------------	--

Exceptions

- unimplemented_STTW* (not used in Oracle SPARC Architecture 2015)
- illegal_instruction*
- mem_address_not_aligned*
- VA_watchpoint*
- DAE_privilege_violation*
- DAE_nfo_page*

See Also

- STW/STX on page 334
- STTWA on page 354

STTWA (Deprecated)


7.132 Store Integer Twin Word into Alternate Space

The STTWA instruction is deprecated and should not be used in new software. The STXA instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
STTWA ^{D, P, ASI}	01 0111	Store Twin Word into Alternate Space	$sttwa\ reg_{rd}[regaddr]\ imm_asi$ $sttwa\ reg_{rd}[reg_plus_imm]\ \%asi$	D2, Y3 ‡

† The original assembly language syntax for this instruction used an “stda” instruction mnemonic, which is now deprecated. Over time, assemblers will support the new “sttwa” mnemonic for this instruction. In the meantime, some existing assemblers may only recognize the original “stda” mnemonic.

‡ **Y3** for restricted ASIs (00₁₆-7F₁₆); **D2** for unrestricted ASIs (80₁₆-FF₁₆)



Description The store twin word integer into alternate space instruction (STTWA) copies two words from an R register pair into memory. The least significant 32 bits of the even-numbered R register are written into memory at the effective address, and the least significant 32 bits of the following odd-numbered R register are written into memory at the “effective address + 4”.

The least significant bit of the rd field of an STTWA instruction is unused and should always be set to 0 by software.

Store integer twin word to alternate space instructions contain the address space identifier (ASI) to be used for the store in the imm_asi field if i = 0, or in the ASI register if i = 1. The effective address for these instructions is “R[rs1] + R[rs2]” if i = 0, or “R[rs1]+sign_ext(simm13)” if i = 1.

A successful store twin word instruction operates atomically.

With respect to little-endian memory, an STTWA instruction behaves as if it is composed of two 32-bit stores, each of which is byte-swapped independently before being written into its respective destination memory word.

IMPL. DEP. #108-V9b: It is implementation dependent whether STTWA is implemented in hardware. If not, an attempt to execute it will cause an *unimplemented_STTW* exception. (STTWA is implemented in hardware in all Oracle SPARC Architecture 2015 implementations.)

An attempt to execute an STTWA instruction with a misaligned (odd) destination register number rd causes an *illegal_instruction* exception.

STTWA causes a *mem_address_not_aligned* exception if the effective address is not doubleword-aligned.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, this instruction causes a *privileged_action* exception.

STTWA (Deprecated)

STTWA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with this instruction causes a *DAE_invalid_asi* exception (impl. dep. #300-U4-Cs10).

ASIs valid for STTWA	
04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
15 ₁₆ ASI_REAL_IO	1D ₁₆ ASI_REAL_IO_LITTLE
22 ₁₆ ASI_STBI_AIUP	2A ₁₆ ASI_STBI_AIUP_L
23 ₁₆ ASI_STBI_AIUS	2B ₁₆ ASI_STBI_AIUS_L
27 ₁₆ ASI_STBI_N	2F ₁₆ ASI_STBI_N_L
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE
E2 ₁₆ ASI_STBI_P	EA ₁₆ ASI_STBI_PL
E3 ₁₆ ASI_STBI_S	EB ₁₆ ASI_STBI_SL

Programming Note Nontranslating ASIs (see page 423) may only be accessed using STXA (not STTWA) instructions. If an STTWA referencing a nontranslating ASI is executed, per the above table, it generates a *DAE_invalid_asi* exception (impl. dep. #300-U4-Cs10).

Programming Notes STTWA is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties. Therefore, software should avoid using STTWA.

If STTWA is emulated in software, an STXA instruction should be used for Emulation software should examine TSTATE[TL].pstate.cle (and, if appropriate, TTE.ie) to determine the endianness of the emulated memory access.

Note that the value of TTE.ie is not saved during a trap. Therefore, if it is examined in the emulation trap handler, that should be done as quickly as possible, to minimize the window of time during which the value of TTE.ie could possibly be changed from the value it had at the time of the attempted execution of STTWA.

Exceptions

- unimplemented_STTW*
- illegal_instruction*
- mem_address_not_aligned*
- privileged_action*
- VA_watchpoint*
- DAE_invalid_asi*
- DAE_privilege_violation*
- DAE_nfo_page*


See Also

- STWA/STXA on page 335
- STTW on page 352

STXFSR

7.133 Store Floating-Point State Register

Instruction	op3	rd	Operation	Assembly Language	Class
	10 0101	0	(see page 345)		
STXFSR	10 0101	1	Store Floating-Point State register	stx %fsr, [address]	A1
—	10 0101	2–31	Reserved		



Description The store floating-point state register instruction (STXFSR) waits for any currently executing FPop instructions to complete, and then it writes all 64 bits of the FSR into memory.

STXFSR zeroes FSR.ftt after writing the FSR to memory.

Implementation Note FSR.ftt should not be zeroed by STXFSR until it is known that the store will not cause a precise trap.

STXFSR accesses memory using the implicit ASI (see page 83). The effective address for this instruction is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + sign_ext(simm13)” if $i = 1$.

Exceptions. An attempt to execute a STXFSR instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the floating-point unit is not enabled (FPRS.fef = 0 or PSTATE.pef = 0) or if the FPU is not present, then an attempt to execute a STXFSR instruction causes an *fp_disabled* exception.

If the effective address is not doubleword-aligned, an attempt to execute an STXFSR instruction causes a *mem_address_not_aligned* exception.

Implementation Note STXFSR shares an opcode with the (deprecated) STFSR instruction (and possibly with other implementation-dependent instructions); they are differentiated by the instruction rd field. An attempt to execute the $op = 10_2$, $op3 = 10\ 0101_2$ opcode with an invalid rd value causes an *illegal_instruction* exception.


Exceptions *illegal_instruction*
fp_disabled
mem_address_not_aligned
VA_watchpoint
DAE_privilege_violation
DAE_nfo_page

See Also *Load Floating-Point State Register* on page 264
Store Floating-Point on page 340
Store Floating-Point State Register (Lower) on page 345

SUB

7.134 Subtract

Instruction	op3	Operation	Assembly Language Syntax	Class
SUB	00 0100	Subtract	sub <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
SUBcc	01 0100	Subtract and modify cc's	subcc <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
SUBC	00 1100	Subtract with Carry	subc <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
SUBCcc	01 1100	Subtract with Carry and modify cc's	subccc <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1



Description These instructions compute “R[rs1] – R[rs2]” if *i* = 0, or “R[rs1] – **sign_ext**(simm13)” if *i* = 1, and write the difference into R[rd].

SUBC and SUBCcc (“SUBtract with carry”) also subtract the CCR register’s 32-bit carry (icc.c) bit; that is, they compute “R[rs1] – R[rs2] – icc.c” or “R[rs1] – **sign_ext**(simm13) – icc.c” and write the difference into R[rd].

SUBcc and SUBCcc modify the integer condition codes (CCR.icc and CCR.xcc). A 32-bit overflow (CCR.icc.v) occurs on subtraction if bit 31 (the sign) of the operands differs and bit 31 (the sign) of the difference differs from R[rs1]{31}. A 64-bit overflow (CCR.xcc.v) occurs on subtraction if bit 63 (the sign) of the operands differs and bit 63 (the sign) of the difference differs from R[rs1]{63}.

Programming Notes A SUBcc instruction with rd = 0 can be used to effect a signed or unsigned integer comparison. See the *cmp* synthetic instruction in Appendix C, *Assembly Language Syntax*.
SUBC and SUBCcc read the 32-bit condition codes’ carry bit (CCR.icc.c), not the 64-bit condition codes’ carry bit (CCR.xcc.c).

An attempt to execute a SUB instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*

SUBXC

7.135 Subtract Extended with 64-bit Carry

The SUBXC instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
SUBXC ^N	0 0100 0001	Subtract Extended with 64-bit Carry	subxc <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2015
SUBXCcc ^N	0 0100 0011	Subtract Extended with 64-bit Carry and modify cc's	subxccc <i>reg_{rs1}, reg_{rs2}, reg_{rd}</i>	C1	OSA 2015



Description SUBXC and SUBXCcc (“SUB extended with carry”) both compute “R[rs1] – R[rs2] – xcc.c” and write the difference into R[rd].

In addition, SUBXCcc modifies the integer condition codes (CCR.icc and CCR.xcc). Overflow occurs on subtraction if bit 63 (the sign bit) of the two operands differ and bit 63 (the sign bit) of the difference differs from bit 63 of the first operand (R[rs1]).

Exceptions None


See Also Add Extended with 64-bit Carry on page 111
Subtract on page 357

SWAP (Deprecated)

7.136 Swap Register with Memory

The SWAP instruction is deprecated and should not be used in new software. The CASA or CASXA instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
SWAP ^D	00 1111	Swap Register with Memory	swap [address], reg _{rd}	D2



Description SWAP exchanges the less significant 32 bits of R[rd] with the contents of the word at the addressed memory location. The upper 32 bits of R[rd] are set to 0. The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more virtual processors executing CASA, CASXA, SWAP, SWAPA, LDSTUB, or LDSTUBA instructions addressing any or all of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial, order.

SWAP accesses memory using the implicit ASI (see page 83). The effective address for these instructions is “R[rs1] + R[rs2]” if i = 0, or “R[rs1] + sign_ext(simm13)” if i = 1.

An attempt to execute a SWAP instruction when i = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If the effective address is not word-aligned, an attempt to execute a SWAP instruction causes a *mem_address_not_aligned* exception.

The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent (impl. dep. #120-V9).

Exceptions *illegal_instruction*
mem_address_not_aligned
VA_watchpoint
DAE_privilege_violation
DAE_nfo_page


SWAPA (Deprecated)

7.137 Swap Register with Alternate Space Memory

The SWAPA instruction is deprecated and should not be used in new software. The CASXA instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax		Class
SWAPA ^{D, P_{ASI}}	01 1111	Swap register with Alternate Space Memory	swapa	[regaddr] imm_asi, reg_rd	D2, Y3 [‡]
			swapa	[reg_plus_imm] %asi, reg_rd	

[‡] **Y3** for restricted ASIs (00₁₆-7F₁₆); **D2** for unrestricted ASIs (80₁₆-FF₁₆)



Description

SWAPA exchanges the less significant 32 bits of R[rd] with the contents of the word at the addressed memory location. The upper 32 bits of R[rd] are set to 0. The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more virtual processors executing CASA, CASXA, SWAP, SWAPA, LDSTUB, or LDSTUBA instructions addressing any or all of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial, order.

The SWAPA instruction contains the address space identifier (ASI) to be used for the load in the imm_asi field if $i = 0$, or in the ASI register if $i = 1$. The effective address for this instruction is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + sign_ext(simm13)” if $i = 1$.

This instruction causes a *mem_address_not_aligned* exception if the effective address is not word-aligned. It causes a *privileged_action* exception if PSTATE.priv = 0 and bit 7 of the ASI is 0.

The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent (impl. dep #120-V9).

If the effective address is not word-aligned, an attempt to execute a SWAPA instruction causes a *mem_address_not_aligned* exception.

In nonprivileged mode (PSTATE.priv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged_action* exception. In privileged mode (PSTATE.priv = 1), if the ASI is in the range 30₁₆ to 7F₁₆, this instruction causes a *privileged_action* exception.

SWAPA can be used with any of the following ASIs, subject to the privilege mode rules described for the *privileged_action* exception above. Use of any other ASI with this instruction causes a *DAE_invalid_asi* exception.

ASIs valid for SWAPA

04 ₁₆ ASI_NUCLEUS	0C ₁₆ ASI_NUCLEUS_LITTLE
10 ₁₆ ASI_AS_IF_USER_PRIMARY	18 ₁₆ ASI_AS_IF_USER_PRIMARY_LITTLE
11 ₁₆ ASI_AS_IF_USER_SECONDARY	19 ₁₆ ASI_AS_IF_USER_SECONDARY_LITTLE
14 ₁₆ ASI_REAL	1C ₁₆ ASI_REAL_LITTLE
80 ₁₆ ASI_PRIMARY	88 ₁₆ ASI_PRIMARY_LITTLE
81 ₁₆ ASI_SECONDARY	89 ₁₆ ASI_SECONDARY_LITTLE

SWAPA (Deprecated)

Exceptions


- mem_address_not_aligned*
- privileged_action*
- VA_watchpoint*
- DAE_invalid_asl*
- DAE_privilege_violation*
- DAE_nc_page*
- DAE_nfo_page*

█

TADDcc

7.138 Tagged Add

Instruction	op3	Operation	Assembly Language Syntax	Class
TADDcc	10 0000	Tagged Add and modify cc's	taddcc <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1



Description This instruction computes a sum that is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + sign_ext(simm13)” if $i = 1$.

TADDcc modifies the integer condition codes (*icc* and *xcc*).

A tag overflow condition occurs if bit 1 or bit 0 of either operand is nonzero or if the addition generates 32-bit arithmetic overflow (that is, both operands have the same value in bit 31 and bit 31 of the sum is different).

If a TADDcc causes a tag overflow, the 32-bit overflow bit (CCR.icc.v) is set to 1; if TADDcc does not cause a tag overflow, CCR.icc.v is set to 0.

In either case, the remaining integer condition codes (both the other CCR.icc bits and all the CCR.xcc bits) are also updated as they would be for a normal ADD instruction. In particular, the setting of the CCR.xcc.v bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). CCR.xcc.v is set based on the 64-bit arithmetic overflow condition, like a normal 64-bit add.

An attempt to execute a TADDcc instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*

See Also TADDccTV^D on page 363
TSUBcc on page 367

TADDccTV (Deprecated)

7.139 Tagged Add and Trap on Overflow

The TADDccTV instruction is deprecated and should not be used in new software. The TADDcc instruction followed by the BPVS instruction (with instructions to save the pre-TADDcc integer condition codes if necessary) should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
TADDccTV ^D	10 0010	Tagged Add and modify cc's or Trap on Overflow	taddoctv <i>reg_rs1, reg_or_imm, reg_rd</i>	D2



Description This instruction computes a sum that is “R[rs1] + R[rs2]” if $i = 0$, or “R[rs1] + **sign_ext**(simm13)” if $i = 1$.

TADDccTV modifies the integer condition codes if it does not trap.

An attempt to execute a TADDccTV instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

A tag overflow condition occurs if bit 1 or bit 0 of either operand is nonzero or if the addition generates 32-bit arithmetic overflow (that is, both operands have the same value in bit 31 and bit 31 of the sum is different).

If TADDccTV causes a tag overflow, a *tag_overflow* exception is generated and R[rd] and the integer condition codes remain unchanged. If a TADDccTV does not cause a tag overflow, the sum is written into R[rd] and the integer condition codes are updated. CCR.icc.v is set to 0 to indicate no 32-bit overflow.

In either case, the remaining integer condition codes (both the other CCR.icc bits and all the CCR.xcc bits) are also updated as they would be for a normal ADD instruction. In particular, the setting of the CCR.xcc.v bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). CCR.xcc.v is set only on the basis of the normal 64-bit arithmetic overflow condition, like a normal 64-bit add.

SPARC V8 Compatibility Note TADDccTV traps based on the 32-bit overflow condition, just as in the SPARC V8 architecture. Although the tagged add instructions set the 64-bit condition codes CCR.xcc, there is no form of the instruction that traps on the 64-bit overflow condition.

Exceptions *illegal_instruction*
tag_overflow


See Also TADDcc on page 362
TSUBccTV^D on page 368

Tcc

7.140 Trap on Integer Condition Codes (Tcc)

Instruction	op3	cond	Operation	cc	Test	Assembly Language Syntax	Class
TA	11 1010	1000	Trap Always	1	ta	<i>i_or_x_cc, software_trap_number</i>	A1
TN	11 1010	0000	Trap Never	0	tn	<i>i_or_x_cc, software_trap_number</i>	A1
TNE	11 1010	1001	Trap on Not Equal	not Z	tne [†]	<i>i_or_x_cc, software_trap_number</i>	A1
TE	11 1010	0001	Trap on Equal	Z	te [‡]	<i>i_or_x_cc, software_trap_number</i>	A1
TG	11 1010	1010	Trap on Greater	not (Z or (N xor V))	tg	<i>i_or_x_cc, software_trap_number</i>	A1
TLE	11 1010	0010	Trap on Less or Equal	Z or (N xor V)	tle	<i>i_or_x_cc, software_trap_number</i>	A1
TGE	11 1010	1011	Trap on Greater or Equal	not (N xor V)	tge	<i>i_or_x_cc, software_trap_number</i>	A1
TL	11 1010	0011	Trap on Less	N xor V	tl	<i>i_or_x_cc, software_trap_number</i>	A1
TGU	11 1010	1100	Trap on Greater, Unsigned	not (C or Z)	tgu	<i>i_or_x_cc, software_trap_number</i>	A1
TLEU	11 1010	0100	Trap on Less or Equal, Unsigned	(C or Z)	tleu	<i>i_or_x_cc, software_trap_number</i>	A1
TCC	11 1010	1101	Trap on Carry Clear (Greater than or Equal, Unsigned)	not C	tcc [◇]	<i>i_or_x_cc, software_trap_number</i>	A1
TCS	11 1010	0101	Trap on Carry Set (Less Than, Unsigned)	C	tcs [∇]	<i>i_or_x_cc, software_trap_number</i>	A1
TPOS	11 1010	1110	Trap on Positive or zero	not N	tpos	<i>i_or_x_cc, software_trap_number</i>	A1
TNEG	11 1010	0110	Trap on Negative	N	tneg	<i>i_or_x_cc, software_trap_number</i>	A1
TVC	11 1010	1111	Trap on Overflow Clear	not V	tvc	<i>i_or_x_cc, software_trap_number</i>	A1
TVS	11 1010	0111	Trap on Overflow Set	V	tvs	<i>i_or_x_cc, software_trap_number</i>	A1

[†] synonym: tnz [‡] synonym: tz [◇] synonym: tgeu [∇] synonym: tlu



cc1 :: cc0	Condition Codes Evaluated
00	CCR.icc
01	— (<i>illegal_instruction</i>)
10	CCR.xcc
11	— (<i>illegal_instruction</i>)

Tcc

Description

The Tcc instruction evaluates the selected integer condition codes (icc or xcc) according to the cond field of the instruction, producing either a TRUE or FALSE result. If TRUE and no higher-priority exceptions or interrupt requests are pending, then a *trap_instruction* or *htrap_instruction* exception is generated. If FALSE, the *trap_instruction* (or *htrap_instruction*) exception does not occur and the instruction behaves like a NOP.

For brevity, in the remainder of this section the value of the “software trap number” used by Tcc will be referred to as “SWTN”.

In nonprivileged mode, if $i = 0$ the SWTN is specified by the least significant seven bits of “R[rs1] + R[rs2]”. If $i = 1$, the SWTN is provided by the least significant seven bits of “R[rs1] + imm_trap_#”. Therefore, the valid range of values for SWTN in nonprivileged mode is 0 to 127. The most significant 57 bits of SWTN are unused and should be supplied as zeroes by software.

In privileged mode, if $i = 0$ the SWTN is specified by the least significant eight bits of “R[rs1] + R[rs2]”. If $i = 1$, the SWTN is provided by the least significant eight bits of “R[rs1] + imm_trap_#”. Therefore, the valid range of values for SWTN in privileged mode is 0 to 255. The most significant 56 bits of SWTN are unused and should be supplied as zeroes by software.

Generally, values of $0 \leq \text{SWTN} \leq 127$ are used to trap to privileged-mode software and values of $128 \leq \text{SWTN} \leq 255$ are used to trap to hyperprivileged-mode software. The behavior of Tcc, based on the privilege mode in effect when it is executed and the value of the supplied SWTN, is as follows:

Privilege Mode in effect when Tcc is executed	Behavior of Tcc instruction	
	$0 \leq \text{SWTN} \leq 127$	$128 \leq \text{SWTN} \leq 255$
Nonprivileged (PSTATE.priv = 0)	<i>trap_instruction</i> exception (to privileged mode) ($256 \leq \text{TT} \leq 383$)	— (not possible, because SWTN is a 7-bit value in nonprivileged mode)
Privileged (PSTATE.priv = 1)	<i>trap_instruction</i> exception (to privileged mode) ($256 \leq \text{TT} \leq 383$)	<i>htrap_instruction</i> exception (to hyperprivileged mode) ($384 \leq \text{TT} \leq 511$)

Programming Note Tcc can be used to implement breakpointing, tracing, and calls to privileged and hyperprivileged software. It can also be used for runtime checks, such as for out-of-range array indexes and integer overflow.

Exceptions. An attempt to execute a Tcc instruction when any of the following conditions exist causes an *illegal_instruction* exception:

- instruction bit 29 is nonzero
- $i = 0$ and instruction bits 10:5 are nonzero
- $i = 1$ and instruction bits 10:8 are nonzero
- cc0 = 1

If the Trap on Control Transfer feature is implemented (impl. dep. #450-S20) and PSTATE.tct = 1, then Tcc generates a *control_transfer_instruction* exception instead of causing a control transfer. When a *control_transfer_instruction* trap occurs, PC (the address of the Tcc instruction) is stored in TPC[TL] and the value of NPC from before the Tcc was executed is stored in TNPC[TL]. The full 64-bit (nonmasked) PC and NPC values are stored in TPC[TL] and TNPC[TL], regardless of the value of PSTATE.am.

If a Tcc instruction causes a *trap_instruction* trap, 256 plus the SWTN value is written into TT[TL]. Then the trap is taken and the virtual processor performs the normal trap entry procedure, as described in *Trap Processing* on page 458.

Tcc


Exceptions

- illegal_instruction*
- control_transfer_instruction* (impl. dep. #450-S20)
- trap_instruction* ($0 \leq \text{SWTN} \leq 127$)
- htrap_instruction* ($128 \leq \text{SWTN} \leq 255$)

TSUBcc

7.141 Tagged Subtract

Instruction	op3	Operation	Assembly Language Syntax	Class
TSUBcc	10 0001	Tagged Subtract and modify cc's	<code>tsubcc reg_rs1, reg_or_imm, reg_rd</code>	A1



Description This instruction computes “R[rs1] – R[rs2]” if $i = 0$, or “R[rs1] – `sign_ext(simm13)`” if $i = 1$.

TSUBcc modifies the integer condition codes (icc and xcc).

A tag overflow condition occurs if bit 1 or bit 0 of either operand is nonzero or if the subtraction generates 32-bit arithmetic overflow; that is, the operands have different values in bit 31 (the 32-bit sign bit) and the sign of the 32-bit difference in bit 31 differs from bit 31 of R[rs1].

If a TSUBcc causes a tag overflow, the 32-bit overflow bit (CCR.icc.v) is set to 1; if TSUBcc does not cause a tag overflow, CCR.icc.v is set to 0.

In either case, the remaining integer condition codes (both the other CCR.icc bits and all the CCR.xcc bits) are also updated as they would be for a normal subtract instruction. In particular, the setting of the CCR.xcc.v bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). `ccr.xcc.v` is set based on the 64-bit arithmetic overflow condition, like a normal 64-bit subtract.

An attempt to execute a TSUBcc instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*


See Also TADDcc on page 362
TSUBccTV^D on page 368

TSUBccTV (Deprecated)

7.142 Tagged Subtract and Trap on Overflow

The TSUBccTV instruction is deprecated and should not be used in new software. The TSUBcc instruction followed by BPVS instead (with instructions to save the pre-TSUBcc integer condition codes if necessary) should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
TSUBccTV ^D	10 0011	Tagged Subtract and modify cc's or Trap on Overflow	<code>tsubcc tv reg_rs1, reg_or_imm, reg_rd</code>	D2



Description This instruction computes “R[rs1] – R[rs2]” if $i = 0$, or “R[rs1] – **sign_ext**(simm13)” if $i = 1$.

TSUBccTV modifies the integer condition codes (icc and xcc) if it does not trap.

A tag overflow condition occurs if bit 1 or bit 0 of either operand is nonzero or if the subtraction generates 32-bit arithmetic overflow; that is, the operands have different values in bit 31 (the 32-bit sign bit) and the sign of the 32-bit difference in bit 31 differs from bit 31 of R[rs1].

An attempt to execute a TSUBccTV instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

If TSUBccTV causes a tag overflow, then a *tag_overflow* exception is generated and R[rd] and the integer condition codes remain unchanged. If a TSUBccTV does not cause a tag overflow condition, the difference is written into R[rd] and the integer condition codes are updated. CCR.icc.v is set to 0 to indicate no 32-bit overflow.

In either case, the remaining integer condition codes (both the other CCR.icc bits and all the CCR.xcc bits) are also updated as they would be for a normal subtract instruction. In particular, the setting of the CCR.xcc.v bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). CCR.xcc.v is set only on the basis of the normal 64-bit arithmetic overflow condition, like a normal 64-bit subtract.

SPARC V8 Compatibility Note | TSUBccTV traps based on the 32-bit overflow condition, just as in the SPARC V8 architecture. Although the tagged add instructions set the 64-bit condition codes CCR.xcc, there is no form of the instruction that traps on the 64-bit overflow condition.

Exceptions *illegal_instruction*
tag_overflow


See Also TADDccTV^D on page 363
TSUBcc on page 367

UDIV, UDIVcc (Deprecated)

7.143 Unsigned Divide (64-bit ÷ 32-bit)

The UDIV and UDIVcc instructions are deprecated and should not be used in new software. The UDIVX instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
UDIV ^D	00 1110	Unsigned Integer Divide	udiv <i>reg_rs1, reg_or_imm, reg_rd</i>	D3
UDIVcc ^D	01 1110	Unsigned Integer Divide and modify cc's	udivcc <i>reg_rs1, reg_or_imm, reg_rd</i>	D3



Description The unsigned divide instructions perform 64-bit by 32-bit division, producing a 32-bit result. If $i = 0$, they compute “ $(Y :: R[rs1]\{31:0\}) \div R[rs2]\{31:0\}$ ”. Otherwise (that is, if $i = 1$), the divide instructions compute “ $(Y :: R[rs1]\{31:0\}) \div (\text{sign_ext}(\text{simm13})\{31:0\})$ ”. In either case, if overflow does not occur, the less significant 32 bits of the integer quotient are sign- or zero-extended to 64 bits and are written into R[rd].

The contents of the Y register are undefined after any 64-bit by 32-bit integer divide operation.

Unsigned Divide

Unsigned divide (UDIV, UDIVcc) assumes an unsigned integer doubleword dividend ($Y :: R[rs1]\{31:0\}$) and an unsigned integer word divisor $R[rs2]\{31:0\}$ or $(\text{sign_ext}(\text{simm13})\{31:0\})$ and computes an unsigned integer word quotient (R[rd]). Immediate values in simm13 are in the ranges 0 to $2^{12} - 1$ and $2^{32} - 2^{12}$ to $2^{32} - 1$ for unsigned divide instructions.

Unsigned division rounds an inexact rational quotient toward zero.

Programming Note The *rational quotient* is the infinitely precise result quotient. It includes both the integer part and the fractional part of the result. For example, the rational quotient of $11/4 = 2.75$ (integer part = 2, fractional part = .75).

The result of an unsigned divide instruction can overflow the less significant 32 bits of the destination register R[rd] under certain conditions. When overflow occurs, the largest appropriate unsigned integer is returned as the quotient in R[rd]. The condition under which overflow occurs and the value returned in R[rd] under this condition are specified in TABLE 7-21.

TABLE 7-21 UDIV / UDIVcc Overflow Detection and Value Returned

Condition Under Which Overflow Occurs	Value Returned in R[rd]
Rational quotient $\geq 2^{32}$	$2^{32} - 1$ (0000 0000 FFFF FFFF ₁₆)

When no overflow occurs, the 32-bit result is zero-extended to 64 bits and written into register R[rd].

UDIV, UDIVcc (Deprecated)

UDIV does not affect the condition code bits. UDIVcc writes the integer condition code bits as shown in the following table. Note that negative (N) and zero (Z) are set according to the value of R[rd] after it has been set to reflect overflow, if any.

Bit	Effect on bit of UDIVcc instruction
icc.n	Set if R[rd][31] = 1
icc.z	Set if R[rd][31:0] = 0
icc.v	Set if overflow (<i>per</i> TABLE 7-21)
icc.c	Zero
xcc.n	Set if R[rd][63] = 1
xcc.z	Set if R[rd][63:0] = 0
xcc.v	Zero
xcc.c	Zero

An attempt to execute a UDIV or UDIVcc instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*
 division_by_zero

See Also RDY on page 310
 SDIV[cc] on page 325,
 UMUL[cc] on page 371

UMUL, UMULcc (Deprecated)

7.144 Unsigned Multiply (32-bit)

The UMUL and UMULcc instructions are deprecated and should not be used in new software. The MULX instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
UMUL ^D	00 1010	Unsigned Integer Multiply	umul <i>reg_rs1</i> , <i>reg_or_imm</i> , <i>reg_rd</i>	D3
UMULcc ^D	01 1010	Unsigned Integer Multiply and modify cc's	umulcc <i>reg_rs1</i> , <i>reg_or_imm</i> , <i>reg_rd</i>	D3



Description The unsigned multiply instructions perform 32-bit by 32-bit multiplications, producing 64-bit results. They compute “R[rs1]{31:0} × R[rs2]{31:0}” if *i* = 0, or “R[rs1]{31:0} × **sign_ext**(simm13){31:0}” if *i* = 1. They write the 32 most significant bits of the product into the Y register and all 64 bits of the product into R[rd].

Unsigned multiply instructions (UMUL, UMULcc) operate on unsigned integer word operands and compute an unsigned integer doubleword product.

UMUL does not affect the condition code bits. UMULcc writes the integer condition code bits, *icc* and *xcc*, as shown below.

Bit	Effect on bit by execution of UMULcc
icc.n	Set to 1 if product{31} = 1; otherwise, set to 0
icc.z	Set to 1 if product{31:0} = 0; otherwise, set to 0
icc.v	Set to 0
icc.c	Set to 0
xcc.n	Set to 1 if product{63} = 1; otherwise, set to 0
xcc.z	Set to 1 if product{63:0} = 0; otherwise, set to 0
xcc.v	Set to 0
xcc.c	Set to 0

Note 32-bit negative (*icc.n*) and zero (*icc.z*) condition codes are set according to the *less* significant word of the product, not according to the full 64-bit result.

Programming Notes 32-bit overflow after UMUL or UMULcc is indicated by *Y* ≠ 0.

An attempt to execute a UMUL or UMULcc instruction when *i* = 0 and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

Exceptions *illegal_instruction*


See Also RDY on page 310
SMUL[cc] on page 333,
UDIV[cc] on page 369

UMULXHI

7.145 Integer Multiply High (64-bit) VIS 3

The UMULXHI instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
UMULXHI	0 0001 0110	Unsigned integer multiply yielding upper 64 bits of 128-bit product	i64	i64	i64	umulxhi <i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	C1	OSA 2011



Description UMULXHI computes “R[rs1] × R[rs2]” and writes the upper 64-bits of the 128-bit product into R[rd]. The two 64-bit operands are treated as unsigned integer values.

UMULXHI does not modify any condition codes.




FIGURE 7-60 UMULXHI

UMULXHI, in conjunction with MULX, can be used to multiply very large numbers (e.g., 1024b x 1024b) together. Most of the required multiplications are unsigned. However, the few signed 64b X unsigned 64b and signed 64b X signed 64b “MULXHI” products can be calculated via UMULXHI and other standard SPARC V9 instructions:

Signed 64 bit × signed 64 bit {127:64} =
 UMULXHI product – (R[rs1]{63} × R[rs2] + R[rs2]{63} × R[rs1])


Unsigned 64 bit × signed 64 bit{127:64} =
 UMULXHI product – (R[rs2]{63} × R[rs1])

Exceptions None

7.146 Write Ancillary State Register

Instruction	rd	Operation	Assembly Language Syntax	Class
WRY ^D	0	Write Y register (<i>deprecated</i>)	<code>wr reg_rs1, reg_or_imm, %y</code>	D3
—	1	<i>Reserved</i>		
WRCCR	2	Write Condition Codes register	<code>wr reg_rs1, reg_or_imm, %ccr</code>	A1
WRASI	3	Write ASI register	<code>wr reg_rs1, reg_or_imm, %asi</code>	A1
—	4	<i>Reserved</i> (read-only ASR (TICK))		
—	5	<i>Reserved</i> (read-only ASR (PC))		
WRFPSR	6	Write Floating-Point Registers Status register	<code>wr reg_rs1, reg_or_imm, %fprs</code>	A1
—	7–13	<i>Reserved</i> (7-0D ₁₆)		
—	14	(0E ₁₆) <i>Reserved</i>		
—	15	(0F ₁₆) <i>used at higher privilege level</i>		
—	16–18	<i>Reserved</i> (impl. dep. #8-V8-Cs20, #9- (10-12 ₁₆)V8-Cs20)		
WRGSR	19 (13 ₁₆)	Write General Status register (GSR)	<code>wr reg_rs1, reg_or_imm, %gsr</code>	A1
WRSOFTINT_SET ^P	20 (14 ₁₆)	Set bits of per-virtual processor Soft Interrupt register	<code>wr reg_rs1, reg_or_imm, %softint_set</code>	N-
WRSOFTINT_CLR ^P	21 (15 ₁₆)	Clear bits of per-virtual processor Soft Interrupt register	<code>wr reg_rs1, reg_or_imm, %softint_clr</code>	N-
WRSOFTINT ^P	22 (16 ₁₆)	Write per-virtual processor Soft Interrupt register	<code>wr reg_rs1, reg_or_imm, %softint</code>	N-
—	23	(17 ₁₆) <i>Reserved</i>		
—	24	(18 ₁₆) <i>used at higher privilege level</i>		
WRSTICK_CMPR ^P	25 (19 ₁₆)	Write System Tick Compare register	<code>wr reg_rs1, reg_or_imm, %stick_cmpr†</code>	N-
—	26	(1A ₁₆) <i>Reserved</i> (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		
—	26	(1A ₁₆) <i>used at higher privilege level</i>		
WRPAUSE	27 (1B ₁₆)	Pause virtual processor	<code>wr reg_rs1, reg_or_imm, %pause</code>	
WRMWAIT	28 (1C ₁₆)	Monitor-Wait virtual processor	<code>wr reg_rs1, reg_or_imm, %mwait</code>	
—	29	(1D ₁₆) Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		
—	31	(1F ₁₆) Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		

† The original assembly language names for `%stick` and `%stick_cmpr` were, respectively, `%sys_tick` and `%sys_tick_cmpr`, which are now deprecated. Over time, assemblers will support the new `%stick` and `%stick_cmpr` names for these registers (which are consistent with `%tick`). In the meantime, some existing assemblers may only recognize the original names.



Description The WRAsr instructions each store a value to the writable fields of the ancillary state register (ASR) specified by rd.

WRAsr

The value stored by these instructions (other than the implementation-dependent variants) is as follows: if $i = 0$, store the value “R[rs1] xor R[rs2]”; if $i = 1$, store “R[rs1] xor sign_ext(simm13)”.

Note | The operation is **exclusive-or**.

The WRAsr instruction with $rd = 0$ is a (deprecated) WRY instruction (which should not be used in new software). WRY is *not* a delayed-write instruction; the instruction immediately following a WRY observes the new value of the Y register.

The WRY instruction is deprecated. It is recommended that all instructions that reference the Y register be avoided.

WRCCR, WRFPRS, and WRASI are *not* delayed-write instructions. The instruction immediately following a WRCCR, WRFPRS, or WRASI observes the new value of the CCR, FPRS, or ASI register.

WRFPRS waits for any pending floating-point operations to complete before writing the FPRS register.

WRPAUSE writes a value to the PAUSE register, requesting a temporary suspension of instruction execution on the virtual processor. The value written to PAUSE indicates the requested number of processor nanoseconds to pause. See *Pause* on page 298 and *Pause Count (pause) Register (ASR 27)* on page 60 for details.

WRMWAIT writes a value to the MWAIT register, requesting a temporary suspension of instruction execution on the virtual processor while waiting for modification of a memory location loaded by an older Load-Monitor instruction. The value written to MWAIT indicates the requested number of nanoseconds to pause. See *MWait* on page 292 and *Mwait Count (mwait) Register (ASR 28)* on page 61 for details.

IMPL. DEP. #48-V8-Cs20: WRAsr instructions with rd of 16-18, 28, 29, or 31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For a WRAsr instruction using one of those rd values, the following are implementation dependent:

- the interpretation of bits 18:0 in the instruction
- the operation(s) performed (for example, **xor**) to generate the value written to the ASR
- whether the instruction is nonprivileged or privileged (impl. dep. #9-V8-Cs20), and
- whether an attempt to execute the instruction causes an *illegal_instruction* exception.

V9 Compatibility Notes | Ancillary state registers may include (for example) timer, counter, diagnostic, self-test, and trap-control registers.
The SPARC V8 WRIER, WRPSR, WRWIM, and WRTBR instructions do not exist in the Oracle SPARC Architecture because the IER, PSR, TBR, and WIM registers do not exist in the Oracle SPARC Architecture.

See *Ancillary State Registers* on page 48 for more detailed information regarding ASR registers.

Exceptions. An attempt to execute a WRAsr instruction when any of the following conditions exist causes an *illegal_instruction* exception:

- $i = 0$ and instruction bits 12:5 are nonzero
- rd contains a reserved value (that is not assigned for implementation-dependent use)
- $rd = 15$ and $((rs1 \neq 0) \text{ or } (i = 0))$

An attempt to execute a WRSOFTINT_SET, WRSOFTINT_CLR, WRSOFTINT, or WRSTICK_CMPR instruction in nonprivileged mode ($PSTATE.priv = 0$) causes a *privileged_opcode* exception.

If the floating-point unit is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if the FPU is not present, then an attempt to execute a WRGSR instruction causes an *fp_disabled* exception.

Implementation Note | WRAsr to ASR 27 ($1B_{16}$) is a PAUSE instruction. See *Pause* on page 298 for details.


WRasr

Exceptions *illegal_instruction*
 privileged_opcode
 fp_disabled

| *See Also* MWAIT on page 292
 PAUSE on page 298
 RDasr on page 310
 WRPR on page 376

7.147 Write Privileged Register

Instruction	op3	Operation	rd	Assembly Language Syntax	Class
WRPR ^P	11 0010	Write Privileged register			A1
		TPC	0	<code>wrpr reg_rs1, reg_or_imm, %tpc</code>	
		TNPC	1	<code>wrpr reg_rs1, reg_or_imm, %tnpc</code>	
		TSTATE	2	<code>wrpr reg_rs1, reg_or_imm, %tstate</code>	
		TT	3	<code>wrpr reg_rs1, reg_or_imm, %tt</code>	
		(illegal_instruction)	4		
		TBA	5	<code>wrpr reg_rs1, reg_or_imm, %tba</code>	
		PSTATE	6	<code>wrpr reg_rs1, reg_or_imm, %pstate</code>	
		TL	7	<code>wrpr reg_rs1, reg_or_imm, %tl</code>	
		PIL	8	<code>wrpr reg_rs1, reg_or_imm, %pil</code>	
		CWP	9	<code>wrpr reg_rs1, reg_or_imm, %cwp</code>	
		CANSAVE	10	<code>wrpr reg_rs1, reg_or_imm, %cansave</code>	
		CANRESTORE	11	<code>wrpr reg_rs1, reg_or_imm, %canrestore</code>	
		CLEANWIN	12	<code>wrpr reg_rs1, reg_or_imm, %cleanwin</code>	
		OTHERWIN	13	<code>wrpr reg_rs1, reg_or_imm, %otherwin</code>	
		WSTATE	14	<code>wrpr reg_rs1, reg_or_imm, %wstate</code>	
		Reserved	15		
		GL	16	<code>wrpr reg_rs1, reg_or_imm, %gl</code>	
		Reserved	16–22		
		Reserved	23		
		Reserved	24–31		



Description

This instruction generates a source value of “R[rs1] xor R[rs2]” if *i* = 0, or “R[rs1] xor sign_ext(simm13)” if *i* = 1. It stores that source value to the writable fields of the specified privileged state register.

Note | The operation is **exclusive-or**.

The *rd* field in the instruction determines the privileged register that is written. There are *MAXPTL* copies of the TPC, TNPC, TT, and TSTATE registers, one for each trap level. A write to one of these registers sets the register, indexed by the current value in the trap-level register (TL).

The WRPR instruction is a *non*-delayed-write instruction. The instruction immediately following the WRPR observes any changes made to virtual processor state made by the WRPR.

A WRPR to TL only stores a value to TL; it does not cause a trap, cause a return from a trap, or alter any machine state other than TL and state (such as PC, NPC, TICK, etc.) that is indirectly modified by every instruction.

Programming Note | A WRPR of TL can be used to read the values of TPC, TNPC, and TSTATE for any trap level; however, software must take care that traps do not occur while the TL register is modified.

MAXPTL is the maximum value that may be written by a WRPR to TL; an attempt to write a larger value results in *MAXPTL* being written to TL. For details, see TABLE 5-21 on page 74.

WRPR

MAXPGL is the maximum value that may be written by a WRPR to *GL*; an attempt to write a larger value results in *MAXPGL* being written to *GL*. For details, see TABLE 5-22 on page 75.

An attempt to use a WRPR instruction to write a value greater than *N_REG_WINDOWS* – 1 to *CANSAVE*, *CANRESTORE*, *OTHERWIN*, or *CLEANWIN* causes an implementation-dependent value in the range 0 .. *N_REG_WINDOWS* – 1 to be written to the register (impl. dep. #126-V9-Ms10Cs40).

Exceptions. An attempt to execute a WRPR instruction in nonprivileged mode (*PSTATE.priv* = 0) causes a *privileged_opcode* exception.

An attempt to execute a WRPR instruction when any of the following conditions exist causes an *illegal_instruction* exception:

- *i* = 0 and instruction bits 12:5 are nonzero
- *rd* = 4
- *rd* contains a reserved value (see above)
- $0 \leq rd \leq 3$ (attempt to write *TPC*, *TNPC*, *TSTATE*, or *TT* register) while *TL* = 0 (current trap level is zero) and the virtual processor is in privileged mode.

Implementation Note	In nonprivileged mode, <i>illegal_instruction</i> exception due to $0 \leq rd \leq 3$ and <i>TL</i> = 0 does not occur; the <i>privileged_opcode</i> exception occurs instead.
----------------------------	--

Exceptions *privileged_opcode*
 illegal_instruction

See Also *RDPR* on page 313
 WRAsr on page 373

XMONTMUL

7.148 XMONTMUL Crypto

This instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax		Class	Added
XMONTMUL ^N	1 0100 1001	Montgomery GF(2 ^m) Multiplication	xmontmul	imm5	N1	OSA 2015



Description XMONTMUL performs the Montgomery GF(2^m) Multiplication shown below with length equal to imm5. Multiplies are bitwise where an XOR operation replaces the ADD operation when combining partial products. The starting locations of the operands are constant in the integer register file (IRF), regardless of size. The starting location of the result is also constant. Below are the locations for N=31. For smaller XMONTMULs, the remaining operand locations are not used and the remaining result locations will be unchanged.

FSR.fcc3 is set upon completion of XMONTMUL to reflect whether or not a hardware error occurred during execution (see the Programming Note below).

The following pseudo-code specifies the operation of the XMONTMUL instruction:

```

XMontMul (l_uint A, l_uint B, l_uint N, l_uint *M, l_uint Nprime,
          char Length, l_uint *X) {
    // compute Montgomery GF(2m) Multiplication as described below
    // use M as temporary variable
    // return result in X
    // A,B,N,M,X 64×(Length+1) bit long
    // Nprime is 64 bits long

    ACCUM ← 0
    for I←0 to Length begin // Length is one less than the number of words
        for j←0 to I-1 begin //skipped on first I iteration
            ACCUM ← ACCUM ^ (A[j]×^B[I-j])
            ACCUM ← ACCUM ^ (M[j]×^N[I-j])
        end
        ACCUM ← ACCUM ^ (A[I]×^B[0])
        M[I] ← ACCUM ×^ Nprime // 64 LSB of accum, store 64 LSB of product
        ACCUM ← ACCUM ^ (M[I]×^N[0])
        ACCUM ← ACCUM >> 64
    end

    for I ← (Length+1) to ((2×Length)+1) begin // skip last I iteration
        for j ← (I-Length) to Length begin
            ACCUM ← ACCUM ^ (A[j]×^B[I-j])
            ACCUM ← ACCUM ^ (M[j]×^N[I-j])
        end
        X[I-Length-1] ← ACCUM // 64 LSB of accum
        ACCUM ← ACCUM >> 64
    end
end

```

XMONTMUL

XMONTMUL Operand Locations

Operand	Win- dow (CWP value)	Register(s)	
<i>Source operands</i>			
Nprime	—	F _D [60]	
M[7:0]	i-6	(F _D [2] :: F _D [0] :: R[29] :: R[28] :: R[27] :: R[26] :: R[25] :: R[24])	f2,f0,i5...i0
M[15:8]	i-6	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
M[23:16]	i-6	(F _D [6] :: F _D [4] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f6,f4,o5...o0
M[31:24]	—	(F _D [22] :: F _D [20] :: F _D [18] :: F _D [16] :: F _D [14] :: F _D [12] :: F _D [10] :: F _D [8])	
A[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
A[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24, o5...o0
A[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
A[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	
N[7:0]	i-4	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
N[13:8]	i-4	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[21:14]	i-3	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
N[27:22]	i-3	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[31:28]	i-2	(R[19] :: R[18] :: R[17] :: R[16])	13...10
B[5:0]	i-2	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
B[13:6]	i-1	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
B[19:14]	i-1	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
B[27:20]	i	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
B[31:28]	i	(R[11] :: R[10] :: R[9] :: R[8])	o3...o0
<i>Destination operands</i>			
X[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
X[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24,o5...o0
X[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
X[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	

Programming Note | The XMONTMUL instruction uses seven windows of the integer register file (IRF).

XMONTMUL

Programming Note | Due to the special nature of the XMONTMUL instruction, the processor cannot protect the programmer from hardware errors in integer and floating-point register file locations during execution of XMONTMUL. The XMONTMUL instruction sets FSR.fcc3 to 00₂ if no hardware error occurred, or to 11₂ (unordered) if an error did occur. It is the responsibility of the programmer to check FSR.fcc3 when the instruction completes and to retry the instruction if it encountered a hardware error. The programmer should provide an error counter to allow for a limited number of attempts to retry the operation. The instruction may be retried after first reloading all the input data from a backing storage location; thus, the programmer must take care to preserve a clean copy of the input data.

Exceptions. If $rd \neq 1$ or $rs1 \neq 0$, an attempt to execute a XMONTMUL instruction causes an *illegal_instruction* exception.

If $CFR.xmontmul = 0$, an attempt to execute a XMONTMUL instruction causes a *compatibility_feature* exception.

Programming Note | Software *must* check that $CFR.xmontmul = 1$ before executing the XMONTMUL instruction. If $CFR.xmontmul = 0$, then software should assume that an attempt to execute the XMONTMUL instruction either
(1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
(2) will execute, but perform some other operation.
Therefore, if $CFR.xmontmul = 0$, software should perform the XMONTMUL operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute a MONTMUL instruction causes an *fp_disabled* exception.

The XMONTMUL instruction causes a *fill_n_normal* or *fill_n_other* exception if CANRESTORE is not equal to NWINDOWS-2. The fill trap handler is called with CWP set to point to the window to be filled, that is, old CWP-CANRESTORE-1. The trap vector for the fill trap is based on the values of OTHERWIN and WSTATE, as described in Trap Type for Spill/Fill Traps on page 411. The fill trap handler performs a RESTORED and a RETRY as part of filling the window. When the virtual processor reexecutes the XMONTMUL (due to the handler ending in RETRY), another fill trap will result if more than one window needed to be filled.

Implementation Note | MONTMUL and XMONTMUL share a basic opcode (op and opf). They are differentiated by the instruction rd field; $rd = 0\ 0000_2$ for MONTMUL and $rd = 0\ 0001_2$ for XMONTMUL.

Exceptions *fp_disabled*
 fill_n_normal ($n = 0-7$)
 fill_n_other ($n = 0-7$)


See Also MONTMUL on page 272
 MONTSQR on page 276
 MPMUL on page 286
 XMONTSQR on page 381
 XMPMUL on page 384

XMONTSQR

7.149 XMONTSQR Crypto

This instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax	Class	Added
XMONTSQR ^N	1 0100 1010	Montgomery GF(2 ^m) Squaring	xmontsqr imm5	N1	OSA 2015



Description XMONTSQR performs the Montgomery GF(2^m) Squaring shown below with length equal to imm5. Multiplies are bitwise where an XOR operation replaces the ADD operation when combining partial products. The starting locations of the operands are constant in the integer register file (IRF) regardless of size. The starting location of the result is also constant. Below are the locations for N=31. For smaller XMONTSQRs, the remaining operand locations are not used and the remaining result locations will be unchanged.

FSR.fcc3 is set upon completion of XMONTSQR to reflect whether or not a hardware error occurred during execution (see the Programming Note below).

```

XMontSqr (l_uint A, l_uint N, l_uint *M, l_uint Nprime, char Length, l_uint *X) {
    // compute Montgomery GF(2m)Squaring as described below
    // use M as temporary variable
    // return result in X
    // A,N,M,X 64×(Length+1) bit long
    // Nprime is 64 bits long

    ACCUM ← 0
    for I ← 0 to Length begin           // Length is one less than the number of
        if I is even begin
            ACCUM ← ACCUM ^ A[I/2]^2    // ^2 denotes squaring
        end
        for j ← 0 to I-1 begin
            ACCUM ← ACCUM ^ M[j] ^ N[I-j]
        end
        M[I] ← ACCUM ^ Nprime          // 64 LSB of accum, store 64 LSB of product
        ACCUM ← ACCUM ^ (M[I]^N[0])
        ACCUM ← (ACCUM >> 64)
    end

    for I ← (Length+1) to ((2×Length)+1) begin
        if I is even begin
            ACCUM ← ACCUM ^ A[I/2]^2    // ^2 denotes squaring
        end
        for j = (I-Length) to Length begin
            ACCUM ← ACCUM ^ (M[j]^N[I-j])
        end
        X[I-Length-1] ← ACCUM          // 64 LSB of accum
        ACCUM ← (ACCUM >> 64)
    end
end

```

XMONTSQR

XMONTSQR Operand Location

Operand	Win- dow (CWP value)	Register(s)	Notes
<i>Source operands</i>			
N _{prime}	—	F _D [60]	
M[7:0]	i-6	(F _D [2] :: F _D [0] :: R[29] :: R[28] :: R[27] :: R[26] :: R[25] :: R[24])	f2,f0,i5...i0
M[15:8]	i-6	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	l7...l0
M[23:16]	i-6	(F _D [6] :: F _D [4] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f6,f4,o5...o0
M[31:24]	—	(F _D [22] :: F _D [20] :: F _D [18] :: F _D [16] :: F _D [14] :: F _D [12] :: F _D [10] :: F _D [8])	
A[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	l7...l0
A[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24,o5...o0
A[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
A[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	
N[7:0]	i-4	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	l7...l0
N[13:8]	i-4	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[21:14]	i-3	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	l7...l0
N[27:22]	i-3	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	o5...o0
N[31:28]	i-2	(R[19] :: R[18] :: R[17] :: R[16])	l3...l0
<i>Destination operands</i>			
X[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	l7...l0
X[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24,o5...o0
X[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
X[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	

Programming Note | The XMONTSQR instruction uses seven windows of the integer register file (IRF).

Programming Note | Due to the special nature of the XMONTSQR instruction, the processor cannot protect the programmer from hardware errors in integer and floating-point register file locations during execution of XMONTSQR. The XMONTSQR instruction sets FSR.fcc3 to 00₂ if no hardware error occurred, or to 11₂ (unordered) if an error did occur. It is the responsibility of the programmer to check FSR.fcc3 when the instruction completes and to retry the instruction if it encountered a hardware error. The programmer should provide an error counter to allow for a limited number of attempts to retry the operation. The instruction may be retried after first reloading all the input data from a backing storage location; thus, the programmer must take care to preserve a clean copy of the input data.

Exceptions. If rd ≠ 1 or rs1 ≠ 0, an attempt to execute a XMONTSQR instruction causes an *illegal_instruction* exception.

XMONTSQR

If `CFR.xmongsqr = 0`, an attempt to execute an XMONTSQR instruction causes a *compatibility_feature* exception.

Programming Note	Software <i>must</i> check that <code>CFR.xmongsqr = 1</code> before executing the XMONTSQR instruction. If <code>CFR.xmongsqr = 0</code> , then software should assume that an attempt to execute the XMONTSQR instruction either (1) will generate an <i>illegal_instruction</i> exception because it is not implemented in hardware, or (2) will execute, but perform some other operation. Therefore, if <code>CFR.xmongsqr = 0</code> , software should perform the XMONTSQR operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.
-------------------------	--

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute a XMONTSQR instruction causes an *fp_disabled* exception.

The XMONTSQR instruction causes a *fill_n_normal* or *fill_n_other* exception if `CANRESTORE` is not equal to `NWINDOWS-2`. The fill trap handler is called with `CWP` set to point to the window to be filled, that is, old `CWP-CANRESTORE-1`. The trap vector for the fill trap is based on the values of `OTHERWIN` and `WSTATE`, as described in Trap Type for Spill/Fill Traps on page 411. The fill trap handler performs a `RESTORED` and a `RETRY` as part of filling the window. When the virtual processor reexecutes the XMONTSQR (due to the handler ending in `RETRY`), another fill trap will result if more than one window needed to be filled.

Implementation Note	<code>MONTSQR</code> and <code>XMONTSQR</code> share a basic opcode (<code>op</code> and <code>opf</code>). They are differentiated by the instruction <code>rd</code> field; <code>rd = 0 0000₂</code> for <code>MONTSQR</code> and <code>rd = 0 0001₂</code> for <code>XMONTSQR</code> .
----------------------------	--

Exceptions

- fp_disabled*
- fill_n_normal* (*n* = 0-7)
- fill_n_other* (*n* = 0-7)

See Also

- `MONTMUL` on page 272
- `MONTSQR` on page 276
- `MPMUL` on page 286
- `XMONTMUL` on page 378
- `XMPMUL` on page 384

XMPMUL

7.150 XMPMUL Crypto

This instruction is new and is not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, it currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	Assembly Language Syntax		Class	Added
XMPMUL ^N	1 0100 1000	Multiple Precision Xor Multiply	xmpmul	imm5	N1	OSA 2015



Description XMPMUL bitwise (XOR) multiplies two values, each of width $(N+1) \times 64$ bits, where N is specified in imm5 field of the instruction. XMPMUL uses an XOR operation instead of the ADD operation when combining partial products. The starting locations of the multiplier and multiplicand are constant relative to CWP in the integer register file (IRF), regardless of size. The starting location of the product is also constant. Below are the locations for $N = 31$. For smaller XMPMULs ($N < 31$), the remaining multiplier and multiplicand locations are unused and the remaining product locations are unchanged.

XMPMUL : product[], multiplier[], and multiplicand[] are arrays of 64-bit doublewords.
 $\text{product}[(2N+1):0] \leftarrow \text{multiplier}[N:0] \times^{\wedge} \text{multiplicand}[N:0]$
 Let $i = \text{CWP}$ when XMPMUL is executed.
 Note: doubleword 0 is the least significant doubleword.

The operands are located in registers as follows:

XMPMUL Operand Locations			
Operand	Window (CWP value)	Register(s)	Notes
<i>Source operands</i>			
multiplier[7:0]	i-6	(F _D [2] :: F _D [0] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	
multiplier[15:8]	i-6	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
multiplier[23:16]	i-6	(F _D [6] :: F _D [4] :: R[29] :: R[28] :: R[27] :: R[26] :: R[25] :: R[24])	f6,f4, i5...i0
multiplier[31:24]	—	(F _D [22] :: F _D [20] :: F _D [18] :: F _D [16] :: F _D [14] :: F _D [12] :: F _D [10] :: F _D [8])	
multiplicand[7:0]	i-5	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
multiplicand[15:8]	i-5	(F _D [26] :: F _D [24] :: R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	f26,f24, o5...o0
multiplicand[23:16]	—	(F _D [42] :: F _D [40] :: F _D [38] :: F _D [36] :: F _D [34] :: F _D [32] :: F _D [30] :: F _D [28])	
multiplicand[31:24]	—	(F _D [58] :: F _D [56] :: F _D [54] :: F _D [52] :: F _D [50] :: F _D [48] :: F _D [46] :: F _D [44])	
<i>Destination operands</i>			
product[7:0]	i-4	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[13:8]	i-4	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...o0
product[21:14]	i-3	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[27:22]	i-3	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...o0

XMPMUL

XMPMUL Operand Locations

Operand	Window (CWP value)	Register(s)	Notes
product[35:28]	i-2	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[41:36]	i-2	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...00
product[49:42]	i-1	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10
product[55:50]	i-1	(R[13] :: R[12] :: R[11] :: R[10] :: R[9] :: R[8])	05...00
product[63:56]	i	(R[23] :: R[22] :: R[21] :: R[20] :: R[19] :: R[18] :: R[17] :: R[16])	17...10

Programming Note The XMPMUL instruction uses seven windows of the integer register file (IRF). Before using the code sequence below, the contents of the IRF must be saved. The code shown below assumes the following window register values prior to execution: CANSAVE=NWINDOWS-2, CANRESTORE=0, OTHERWIN=0.

The following code shows an example usage of XMPMUL with length equal to 31.

```
setx multiplier, %g1, %g4
setx multiplicand, %g1, %g5
```

load_multiplier:

```
ldd [%g4 + 0x000], %f22      !# CWP = i-6
ldd [%g4 + 0x008], %f20
ldd [%g4 + 0x010], %f18
ldd [%g4 + 0x018], %f16
ldd [%g4 + 0x020], %f14
ldd [%g4 + 0x028], %f12
ldd [%g4 + 0x030], %f10
ldd [%g4 + 0x038], %f8

ldd [%g4 + 0x040], %f6
ldd [%g4 + 0x048], %f4
ldx [%g4 + 0x050], %i5
ldx [%g4 + 0x058], %i4
ldx [%g4 + 0x060], %i3
ldx [%g4 + 0x068], %i2
ldx [%g4 + 0x070], %i1
ldx [%g4 + 0x078], %i0

ldx [%g4 + 0x080], %i17
ldx [%g4 + 0x088], %i16
ldx [%g4 + 0x090], %i15
ldx [%g4 + 0x098], %i14
ldx [%g4 + 0x0a0], %i13
ldx [%g4 + 0x0a8], %i12
ldx [%g4 + 0x0b0], %i11
ldx [%g4 + 0x0b8], %i10

ldd [%g4 + 0x0c0], %f2
ldd [%g4 + 0x0c8], %f0
ldx [%g4 + 0x0d0], %o5
ldx [%g4 + 0x0d8], %o4
ldx [%g4 + 0x0e0], %o3
ldx [%g4 + 0x0e8], %o2
ldx [%g4 + 0x0f0], %o1
ldx [%g4 + 0x0f8], %o0
```

XMPMUL

```
save                                     !# CWP = i-5

load_multiplicand:
    ldd    [%g5 + 0x000], %f58
    ldd    [%g5 + 0x008], %f56
    ldd    [%g5 + 0x010], %f54
    ldd    [%g5 + 0x018], %f52
    ldd    [%g5 + 0x020], %f50
    ldd    [%g5 + 0x028], %f48
    ldd    [%g5 + 0x030], %f46
    ldd    [%g5 + 0x038], %f44

    ldd    [%g5 + 0x040], %f42
    ldd    [%g5 + 0x048], %f40
    ldd    [%g5 + 0x050], %f38
    ldd    [%g5 + 0x058], %f36
    ldd    [%g5 + 0x060], %f34
    ldd    [%g5 + 0x068], %f32
    ldd    [%g5 + 0x070], %f30
    ldd    [%g5 + 0x078], %f28

    ldd    [%g5 + 0x080], %f26
    ldd    [%g5 + 0x088], %f24
    ldx    [%g5 + 0x090], %o5
    ldx    [%g5 + 0x098], %o4
    ldx    [%g5 + 0x0a0], %o3
    ldx    [%g5 + 0x0a8], %o2
    ldx    [%g5 + 0x0b0], %o1
    ldx    [%g5 + 0x0b8], %o0

    ldx    [%g5 + 0x0c0], %l17
    ldx    [%g5 + 0x0c8], %l16
    ldx    [%g5 + 0x0d0], %l15
    ldx    [%g5 + 0x0d8], %l14
    ldx    [%g5 + 0x0e0], %l13
    ldx    [%g5 + 0x0e8], %l12
    ldx    [%g5 + 0x0f0], %l11
    ldx    [%g5 + 0x0f8], %l10

    save                                       !# CWP = i-4
    save                                       !# CWP = i-3
    save                                       !# CWP = i-2
    save                                       !# CWP = i-1
    save                                       !# CWP = i

run_xmpmul:
    xmpmul    0x1f                               !# CWP = i

store_result:
    setx     vt_result, %g1, %g4
    stx     %l17, [%g4 + 0x000]                 !# CWP = i
    stx     %l16, [%g4 + 0x008]
    stx     %l15, [%g4 + 0x010]
    stx     %l14, [%g4 + 0x018]
    stx     %l13, [%g4 + 0x020]
    stx     %l12, [%g4 + 0x028]
    stx     %l11, [%g4 + 0x030]
    stx     %l10, [%g4 + 0x038]

    restore                                    !# CWP = i-1
```

XMPMUL

```
stx    %o5, [%g4 + 0x040]
stx    %o4, [%g4 + 0x048]
stx    %o3, [%g4 + 0x050]
stx    %o2, [%g4 + 0x058]
stx    %o1, [%g4 + 0x060]
stx    %o0, [%g4 + 0x068]

stx    %17, [%g4 + 0x070]
stx    %16, [%g4 + 0x078]
stx    %15, [%g4 + 0x080]
stx    %14, [%g4 + 0x088]
stx    %13, [%g4 + 0x090]
stx    %12, [%g4 + 0x098]
stx    %11, [%g4 + 0x0a0]
stx    %10, [%g4 + 0x0a8]

restore                                !# CWP = i-2

stx    %o5, [%g4 + 0x0b0]
stx    %o4, [%g4 + 0x0b8]
stx    %o3, [%g4 + 0x0c0]
stx    %o2, [%g4 + 0x0c8]
stx    %o1, [%g4 + 0x0d0]
stx    %o0, [%g4 + 0x0d8]

stx    %17, [%g4 + 0x0e0]
stx    %16, [%g4 + 0x0e8]
stx    %15, [%g4 + 0x0f0]
stx    %14, [%g4 + 0x0f8]
stx    %13, [%g4 + 0x100]
stx    %12, [%g4 + 0x108]
stx    %11, [%g4 + 0x110]
stx    %10, [%g4 + 0x118]

restore                                !# CWP = i-3

stx    %o5, [%g4 + 0x120]
stx    %o4, [%g4 + 0x128]
stx    %o3, [%g4 + 0x130]
stx    %o2, [%g4 + 0x138]
stx    %o1, [%g4 + 0x140]
stx    %o0, [%g4 + 0x148]

stx    %17, [%g4 + 0x150]
stx    %16, [%g4 + 0x158]
stx    %15, [%g4 + 0x160]
stx    %14, [%g4 + 0x168]
stx    %13, [%g4 + 0x170]
stx    %12, [%g4 + 0x178]
stx    %11, [%g4 + 0x180]
stx    %10, [%g4 + 0x188]

restore                                !# CWP = i-4

stx    %o5, [%g4 + 0x190]
stx    %o4, [%g4 + 0x198]
stx    %o3, [%g4 + 0x1a0]
stx    %o2, [%g4 + 0x1a8]
```

XMPMUL

```
stx    %o1, [%g4 + 0x1b0]
stx    %o0, [%g4 + 0x1b8]

stx    %l7, [%g4 + 0x1c0]
stx    %l6, [%g4 + 0x1c8]
stx    %l5, [%g4 + 0x1d0]
stx    %l4, [%g4 + 0x1d8]
stx    %l3, [%g4 + 0x1e0]
stx    %l2, [%g4 + 0x1e8]
stx    %l1, [%g4 + 0x1f0]
stx    %l0, [%g4 + 0x1f8]

restore                !# CWP = i-5
restore                !# CWP = i-6
```

Exceptions. If $rs1 \neq 0$, an attempt to execute a MONTMUL instruction causes an *illegal_instruction* exception.

If $CFR.xmpmul = 0$, an attempt to execute an XMPMUL instruction causes a *compatibility_feature* exception.

Programming Note Software *must* check that $CFR.xmpmul = 1$ before executing the XMPMUL instruction. If $CFR.xmpmul = 0$, then software should assume that an attempt to execute the XMPMUL instruction either

- (1) will generate an *illegal_instruction* exception because it is not implemented in hardware, or
- (2) will execute, but perform some other operation.

Therefore, if $CFR.xmpmul = 0$, software should perform the XMPMUL operation by other means, such as using a software implementation, a crypto coprocessor, or another set of instructions which implement the desired function.

If the FPU is not enabled ($FPRS.fef = 0$ or $PSTATE.pef = 0$) or if no FPU is present, an attempt to execute an XMPMUL instruction causes an *fp_disabled* exception.

The XMPMUL instruction causes a *fill_n_normal* or *fill_n_other* exception if $CANRESTORE$ is not equal to $NWINDOWS-2$. The fill trap handler is called with CWP set to point to the window to be filled, that is, old $CWP-CANRESTORE-1$. The trap vector for the fill trap is based on the values of $OTHERWIN$ and $WSTATE$, as described in Trap Type for Spill/Fill Traps on page 458. The fill trap handler performs a $RESTORED$ and a $RETRY$ as part of filling the window. When the virtual processor reexecutes XMPMUL (due to the handler ending in $RETRY$), another fill trap results if more than one window needed to be filled.

Implementation Note MPMUL and XMPMUL share a basic opcode (op and opf). They are differentiated by the instruction rd field; $rd = 0\ 0000_2$ for MPMUL and $rd = 0\ 0001_2$ for XMPMUL.

Exceptions *fp_disabled*
fill_n_normal ($n = 0-7$)
fill_n_other ($n = 0-7$)


See Also MONTMUL on page 272
 MONTSQR on page 276
 MPMUL on page 286
 XMONTMUL on page 378
 XMONTSQR on page 381

XMULX[HI]

7.151 XOR Multiply VIS 3

The XMULX instructions are new and are not guaranteed to be implemented on all Oracle SPARC Architecture implementations. Therefore, they currently should only be used in platform-specific software (such as system-supplied runtime libraries or in software created by an implementation-aware runtime code generator).

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class	Added
XMULX	1 0001 0101	Bitwise Multiply (low 64 bits)	i64	i64	i64	xmulx <i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	C2	OSA 2011
XMULXHI	1 0001 0110	Bitwise Multiply (high 64 bits)	i64	i64	i64	xmulxhi <i>reg_{rs1}</i> , <i>reg_{rs2}</i> , <i>reg_{rd}</i>	C2	OSA 2011



Description The XMULX instruction performs a 64-bit by 64-bit bitwise (XOR) multiplication. An XOR multiply uses the XOR operation instead of the ADD operation when combining partial products.

XMULX computes R[rs1] **xor-multiply** R[rs2] and writes the less-significant 64 bits of the 128-bit result into R[rd].

XMULXHI computes R[rs1] **xor-multiply** R[rs2] and writes the more-significant 64 bits of the 128-bit result into R[rd].

Neither instruction modifies any condition codes.


Exceptions None

See Also

XOR / XNOR

7.152 XOR Logical Operation

Instruction	op3	Operation	Assembly Language Syntax	Class
XOR	00 0011	Exclusive or	<code>xor</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
XORcc	01 0011	Exclusive or and modify cc's	<code>xorcc</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
XNOR	00 0111	Exclusive nor	<code>xnor</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1
XNORcc	01 0111	Exclusive nor and modify cc's	<code>xnorcc</code> <i>reg_{rs1}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	A1



Description These instructions implement bitwise logical **xor** operations. They compute “R[rs1] **op** R[rs2]” if $i = 0$, or “R[rs1] **op** **sign_ext**(simm13)” if $i = 1$, and write the result into R[rd].

XORcc and XNORcc modify the integer condition codes (icc and xcc). They set the condition codes as follows:

- `icc.v`, `icc.c`, `xcc.v`, and `xcc.c` are set to 0
- `icc.n` is copied from bit 31 of the result
- `xcc.n` is copied from bit 63 of the result
- `icc.z` is set to 1 if bits 31:0 of the result are zero (otherwise to 0)
- `xcc.z` is set to 1 if all 64 bits of the result are zero (otherwise to 0)

Programming | XNOR (and XNORcc) is identical to the **xor_not** (and set condition **Note** | codes) **xor_not_cc** logical operation, respectively.

An attempt to execute an XOR, XORcc, XNOR, or XNORcc instruction when $i = 0$ and instruction bits 12:5 are nonzero causes an *illegal_instruction* exception.

■ **Exceptions** *illegal_instruction*

IEEE Std 754-1985 Requirements for Oracle SPARC Architecture 2015

The IEEE Std 754-1985 floating-point standard contains a number of implementation dependencies. This chapter specifies choices for these implementation dependencies, to ensure that SPARC V9 implementations are as consistent as possible.

The chapter contains these major sections:

- **Traps Inhibiting Results** on page 391.
- **Underflow Behavior** on page 392.
- **Integer Overflow Definition** on page 393.
- **Floating-Point Nonstandard Mode** on page 393.
- **Arithmetic Result Tables** on page 394.

Exceptions are discussed in this chapter on the assumption that instructions are implemented in hardware. If an instruction is implemented in software, it may not trigger hardware exceptions but its behavior as observed by nonprivileged software (other than timing) must be the same as if it was implemented in hardware.

8.1 Traps Inhibiting Results

As described in *Floating-Point State Register (FSR)* on page 42 and elsewhere, when a floating-point trap occurs, the following conditions are true:

- The destination floating-point register(s) (the F registers) are unchanged.
- The floating-point condition codes (*fcc0*, *fcc1*, *fcc2*, and *fcc3*) are unchanged.
- The *FSR.aexc* (accrued exceptions) field is unchanged.
- The *FSR.cexc* (current exceptions) field is unchanged except for *IEEE_754_exceptions*; in that case, *cexc* contains a bit set to 1, corresponding to the exception that caused the trap. Only one bit shall be set in *cexc*.

Instructions causing an *fp_exception_other* trap because of unfinished FPops execute as if by hardware; that is, such a trap is undetectable by application software, except that timing may be affected.

Programming Note	<p>A user-mode trap handler invoked for an IEEE_754_exception, whether as a direct result of a hardware <i>fp_exception_ieee_754</i> trap or as an indirect result of privileged software handling of an <i>fp_exception_other</i> trap with FSR.ftt = unfinished_FPop, can rely on the following behavior:</p> <ul style="list-style-type: none"> ■ The address of the instruction that caused the exception will be available. ■ The destination floating-point register(s) are unchanged from their state prior to that instruction's execution. ■ The floating-point condition codes (fcc0, fcc1, fcc2, and fcc3) are unchanged. ■ The FSR.aexc field is unchanged. ■ The FSR.cexc field contains exactly one bit set to 1, corresponding to the exception that caused the trap. ■ The FSR.ftt, FSR.qne, and reserved fields of FSR are zero.
-------------------------	---

8.2 Underflow Behavior

An Oracle SPARC Architecture virtual processor detects tininess before rounding occurs. (impl. dep. #55-V8-Cs10)

TABLE 8-1 summarizes what happens when an exact *unrounded* value *u* satisfying

$$0 \leq |u| \leq \text{smallest normalized number}$$

would round, if no trap intervened, to a *rounded* value *r* which might be zero, subnormal, or the smallest normalized value.

TABLE 8-1 Floating-Point Underflow Behavior (Tininess Detected Before Rounding)

		Underflow trap: Inexact trap:	ufm = 1 nxm = x	ufm = 0 nxm = 1	ufm = 0 nxm = 0
<i>u = r</i>	<i>r</i> is minimum normal		None	None	None
	<i>r</i> is subnormal		UF	None	None
	<i>r</i> is zero		None	None	None
<i>u ≠ r</i>	<i>r</i> is minimum normal		UF	NX	uf nx
	<i>r</i> is subnormal		UF	NX	uf nx
	<i>r</i> is zero		UF	NX	uf nx
UF = <i>fp_exception_ieee_754</i> trap with cexc.ufc = 1 NX = <i>fp_exception_ieee_754</i> trap with cexc.nxc = 1 uf = cexc.ufc = 1, aexc.ufa = 1, no <i>fp_exception_ieee_754</i> trap nx = cexc.nxc = 1, aexc.nxa = 1, no <i>fp_exception_ieee_754</i> trap					

8.2.1 Trapped Underflow Definition ($ufm = 1$)

Since tininess is detected before rounding, trapped underflow occurs when the exact unrounded result has magnitude between zero and the smallest normalized number in the destination format.

Note | The wrapped exponent results intended to be delivered on trapped underflows and overflows in IEEE 754 are irrelevant to the Oracle SPARC Architecture at the hardware, and privileged software levels. If they are created at all, it would be by user software in a nonprivileged-mode trap handler.

8.2.2 Untrapped Underflow Definition ($ufm = 0$)

Untrapped underflow occurs when the exact unrounded result has magnitude between zero and the smallest normalized number in the destination format *and* the correctly rounded result in the destination format is inexact.

8.3 Integer Overflow Definition

- **F<sdq>TOi** — When a NaN, infinity, large positive argument $\geq 2^{31}$ or large negative argument $\leq -(2^{31} + 1)$ is converted to an integer, the `invalid_current (nvc)` bit of `FSR.cexc` is set to 1, and if the floating-point invalid trap is enabled (`FSR.tem.nvm = 1`), the *fp_exception_IEEE_754* exception is raised. If the floating-point invalid trap is disabled (`FSR.tem.nvm = 0`), no trap occurs and a numerical result is generated: if the sign bit of the operand is 0, the result is $2^{31} - 1$; if the sign bit of the operand is 1, the result is -2^{31} .
- **F<sdq>TOx** — When a NaN, infinity, large positive argument $\geq 2^{63}$, or large negative argument $\leq -(2^{63} + 1)$ is converted to an extended integer, the `invalid_current (nvc)` bit of `FSR.cexc` is set to 1, and if the floating-point invalid trap is enabled (`FSR.tem.nvm = 1`), the *fp_exception_IEEE_754* exception is raised. If the floating-point invalid trap is disabled (`FSR.tem.nvm = 0`), no trap occurs and a numerical result is generated: if the sign bit of the operand is 0, the result is $2^{63} - 1$; if the sign bit of the operand is 1, the result is -2^{63} .

8.4 Floating-Point Nonstandard Mode

If implemented, floating-point nonstandard mode is enabled by setting `FSR.ns = 1` (see *Nonstandard Floating-Point (ns)* on page 43).

An Oracle SPARC Architecture 2015 processor may choose to implement nonstandard floating-point mode in order to obtain higher performance in certain circumstances. For example, when `FSR.ns = 1` an implementation that processes fully normalized operands more efficiently than subnormal operands may convert a subnormal floating-point operand or result to zero.

The behavior of the following instructions is required to follow IEEE Std. 754 at all times, regardless of the value of `FSR.ns`: `FADD`, `FDIV`, `FLCMP`, `FMAf`, `FHADD`, `FHSUB`, `FNHADD`, `FMUL`, `FNADD`, `FNMUL`, `FSQRT`, and `FSUB`.

Implementation Note	Oracle SPARC Architecture virtual processors are strongly discouraged from implementing a nonstandard floating-point mode. Implementations are encouraged to support standard IEEE 754 floating-point arithmetic with reasonable performance in all cases, even if some cases are slower than others.
----------------------------	--

Assuming that nonstandard floating-point mode is implemented, the effects of $\text{FSR.ns} = 1$ are as follows:

- **IMPL. DEP. #18-V8-Ms10(a):** When $\text{FSR.ns} = 1$ and a floating-point *source operand* is subnormal, an implementation may treat the subnormal operand as if it were a floating-point zero value of the same sign.
The cases in which this replacement is performed are implementation dependent. However, if it occurs,
 - (1) it should *not* apply to FLCMP, FABS, FMOV, or FNEG instructions and
 - (2) FADD, FSUB, FCMPE, and FCMP should give identical treatment to subnormal source operands.
 Treating a subnormal source operand as zero may generate an IEEE 754 floating-point “inexact”, “division by zero”, or “invalid” condition (see *Current Exception (cexc)* on page 46). Whether the generated condition(s) trigger an *fp_exception_ieee_754* exception or not depends on the setting of FSR.tem .
- **IMPL. DEP. #18-V8-Ms10(b):** When a floating-point operation generates a subnormal *result* value, an Oracle SPARC Architecture 2015 implementation may either write the result as a subnormal value or replace the subnormal result by a floating-point zero value of the same sign and generate IEEE 754 floating-point “inexact” and “underflow” conditions. Whether these generated conditions trigger an *fp_exception_ieee_754* exception or not depends on the setting of FSR.tem .
- **IMPL. DEP. #18-V8-Ms10(c):** If an FPop generates an *intermediate* result value, the intermediate value is subnormal, and $\text{FSR.ns} = 1$, it is implementation dependent whether (1) the operation continues, using the subnormal value (possibly with some loss of accuracy), or (2) the virtual processor replaces the subnormal intermediate value with a floating-point zero value of the same sign, generates IEEE 754 floating-point “inexact” and “underflow” conditions, completes the instruction, and writes a final result (possibly with some loss of accuracy). Whether generated IEEE conditions trigger an *fp_exception_ieee_754* exception or not depends on the setting of FSR.tem .

If $\text{GSR.im} = 1$, then the value of FSR.ns is ignored and the processor operates as if $\text{FSR.ns} = 0$ (see page 54).

8.5 Arithmetic Result Tables

This section contains detailed tables, showing the results produced by various floating-point operations, depending on their source operands.

Notes on source types:

- Nn is a number in $F[rsn]$, which may be normal or subnormal.
- $\text{QNaN}n$ and $\text{SNaN}n$ are Quiet and Signaling Not-a-Number values in $F[rsn]$, respectively.

Notes on result types:

- R: (rounded) result of operation, which may be normal, subnormal, zero, or infinity. May also cause OF, UF, NX, unfinished.
- dQNaN is the generated default Quiet NaN (sign = 0, exponent = all 1s, fraction = all 1s). The sign of the default Quiet NaN is zero to distinguish it from storage initialized to all ones.

- QNaNn is the Signalling NaN operand from F[rsn] with the Quiet bit asserted

8.5.1 Floating-Point Add (FADD) and Add and Halve (FHADD)

TABLE 8-2 Floating-Point Add operation (F[rs1] + F[rs2])
and Floating-Point Add and Halve operation ((F[rs1] + F[rs2]) ÷ 2)

		F[rs2]								
		-∞	-N2	-0	+0	+N2	+∞	QNaN2	SNaN2	
F[rs1]	-∞	-∞					dQNaN, NV	QNaN2	QNaN2, NV	
	-N1	-R		-N1		±R*				
	-0	-N2		-0	±0**	+N2				
	+0			±0**	+0					
	+N1	±R*		+N1		+R				
	+∞	dQNaN, NV	+∞							
	QNaN1	QNaN1								
	SNaN1	QNaN1, NV								

* if N1 = -N2, then **

** result is +0 unless rounding mode is round to -∞, in which case the result is -0

For the FADD instructions, R may be any number; its generation may cause OF, UF, and/or NX.

For the FHADD instructions, R may be any number; its generation may cause UF and/or NX.

Floating-point add is not commutative when both operands are NaN.

8.5.2 Floating-Point Negative Add (FNADD) and Negative Add and Halve (FNHADD)

TABLE 8-3 Floating-Point Negative Add operation ($-(F[rs1] + F[rs2])$) and Floating-Point Negative Add and Halve operation ($-(F[rs1] + F[rs2]) \div 2$)

		F[rs2]							QNaN2	SNaN2		
		$-\infty$	-N2	-0	+0	+N2	$+\infty$					
F[rs1]	$-\infty$	+∞					dQNaN, NV	QNaN2	QNaN2, NV			
	-N1		+R	+N1		±R*	QNaN2			QNaN2, NV		
	-0		+N2	+0	±0**	-N2						
	+0			±0**	-0							
	+N1		±R*	-N1		-R						
	+∞	dQNaN, NV								-∞		
	QNaN1	QNaN1										
	SNaN1	QNaN1, NV										

* if N1 = -N2, then **

** result is +0 unless rounding mode is round to $-\infty$, in which case the result is -0

For the FNADD instructions, R may be any number; its generation may cause OF, UF, and/or NX.

For the FNHADD instructions, R may be any number; its generation may cause UF and/or NX.

Floating-point negative add is not commutative when both operands are NaN.

Note that rounding occurs after the negation. Thus, FNADD is not equivalent to FADD followed by FNEG when the rounding mode is towards $\pm\infty$.

8.5.3 Floating-Point Subtract (FSUB) and Subtract and Halve (FHSUB)

TABLE 8-4 Floating-Point Subtract operation ($F[rs1] - F[rs2]$) and Floating-Point Subtract and Halve operation ($(F[rs1] - F[rs2]) \div 2$)

		F[rs2]							QNaN2	SNaN2		
		$-\infty$	-N2	-0	+0	+N2	$+\infty$					
F[rs1]	$-\infty$	dQNaN, NV						-∞	QNaN2	QNaN2, NV		
	-N1		±R*	-N1		-R	QNaN2	QNaN2, NV				
	-0		+N2	±0**	-0	-N2						
	+0			+0	±0**							
	+N1		+R	+N1		±R*						
	+∞	+∞						dQNaN, NV				
	QNaN1	QNaN1										
	SNaN1	QNaN1, NV										

* if N1 = N2, then **

** result is +0 unless rounding mode is round to $-\infty$, in which case the result is -0

For the FSUB instructions, R may be any number; its generation may cause OF, UF, and/or NX.
 For the FHSUB instructions, R may be any number; its generation may cause UF and/or NX.
 Note that $-x \neq 0 - x$ when x is zero or NaN.

8.5.4 Floating-Point Multiply

TABLE 8-5 Floating-Point Multiply operation ($F[rs1] \times F[rs2]$)

		F[rs2]						QNaN2	SNaN2
		$-\infty$	-N2	-0	+0	+N2	$+\infty$		
F[rs1]	$-\infty$	$+\infty$	dQNaN, NV			$-\infty$	QNaN2	QNaN2, NV	
	-N1		+R		-R				
	-0	dQNaN, NV		+0	-0	dQNaN, NV			
	+0			-0	+0				
	+N1		-R			+R			
	$+\infty$	$-\infty$	dQNaN, NV			$+\infty$			
	QNaN1	QNaN1							
	SNaN1	QNaN1, NV							

R may be any number; its generation may cause OF, UF, and/or NX.

Floating-point multiply is not commutative when both operands are NaN.

FsMULd (FdMULq) never causes OF, UF, or NX.

A NaN input operand to FsMULd (FdMULq) must be widened to produce a double-precision (quad-precision) NaN output, by filling the least-significant bits of the NaN result with zeros.

8.5.5 Floating-Point Negative Multiply (FNMUL)

TABLE 8-6 Floating-Point Negative Multiply operation ($-(F[rs1] \times F[rs2])$)

		F[rs2]						QNaN2	SNaN2
		$-\infty$	-N2	-0	+0	+N2	$+\infty$		
F[rs1]	$-\infty$	$-\infty$	dQNaN, NV			$+\infty$	QNaN2	QNaN2, NV	
	-N1		-R		+R				
	-0	dQNaN, NV		-0	+0	dQNaN, NV			
	+0			+0	-0				
	+N1		+R			-R			
	$+\infty$	$+\infty$	dQNaN, NV			$-\infty$			
	QNaN1	QNaN1							
	SNaN1	QNaN1, NV							

R may be any number; its generation may cause OF, UF, and/or NX.

Floating-point negative multiply is not commutative when both operands are NaNs.

Note that rounding occurs after the negation. Thus, FNMUL is not equivalent to FMUL followed by FNEG when the rounding mode is towards $\pm\infty$.

FNSMULD never causes OF, UF, or NX.

A NaN input operand to FNSMULD must be widened to produce a double-precision NaN output, by filling the least-significant bits of the NaN result with zeros.

8.5.6 Floating-Point Multiply-Add (FMAf)

First refer to the Floating-Point Multiply table (TABLE 8-5 on page 397) to select a row in the table below.

TABLE 8-7 Floating-Point Multiply-Add ((F[rs1] × F[rs2]) + F[rs3])

		F[rs3]						QNaN3	SNaN3
		$-\infty$	-N3	-0	+0	+N3	$+\infty$		
F[rs1] × F[rs2]	$-\infty$	$-\infty$				dQNaN, NV		QNaN3	QNaN3, NV
	-N	-R		-N	$\pm R^*$				
	-0	-N3	-0	$\pm 0^{**}$	+N3				
	+0	$\pm 0^{**}$		+0					
	+N	$\pm R^*$		+N	+R				
	$+\infty$	dQNaN, NV					$+\infty$		
	QNaN1	QNaN1							
	QNaN2	QNaN2							
	QNaN ($\pm 0 \times \pm\infty$)	dQNaN, NV***				QNaN3, NV***			
	QNaN1	QNaN1, NV***							
	QNaN2	QNaN2, NV***							

* if N = -N3, then **

** result is +0 unless rounding mode is round to $-\infty$, in which case the result is -0

*** if FSR.nvm = 1, FSR.nvc ← 1, the trap occurs, and FSR.aexc is left unchanged; otherwise, FSR.nvm = 0 so FSR.nva ← 1 and for FMADD FSR.nvc ← 1.

In the above table, R may be any number; its generation may cause OF, UF, and/or NX.

The multiply operation in fused floating-point multiply-add (FMADD) instructions cannot cause inexact, underflow, or overflow exceptions.

See the earlier sections on Nonstandard Mode and unfinished_FPop for additional details.

8.5.7 Floating-Point Negative Multiply-Add (FNMADD)

First refer to the Floating-Point Multiply table (TABLE 8-5 on page 397) to select a row in the table below.

TABLE 8-8 Floating-Point Negative Multiply-Add $-(F[rs1] \times F[rs2]) - F[rs3]$

		F[rs3]						QNaN3	SNaN3
		$-\infty$	-N3	-0	+0	+N3	$+\infty$		
$F[rs1]$ \times $F[rs2]$	$-\infty$	$+\infty$					dQNaN, NV	QNaN3	QNaN3, NV
	-N	+R		+N	$\pm R^*$				
	-0	+N3	+0	$\pm 0^{**}$	-N3				
	+0	$\pm 0^{**}$		-0					
	+N	$\pm R^*$		-N	-R				
	$+\infty$	dQNaN, NV	$-\infty$						
	QNaN1	QNaN1							
	QNaN2	QNaN2							
	QNaN ($\pm 0 \times \pm \infty$)	dQNaN, NV***					QNaN3 NV***		
	QNaN1	QNaN1, NV***							
	QNaN2	QNaN2, NV***							

* if N = -N3, then **

** result is +0 unless rounding mode is round to $-\infty$, in which case the result is -0

*** if FSR.nvm = 1, FSR.nvc \leftarrow -1, the trap occurs, and FSR.aexc is left unchanged; otherwise, FSR.nvm = 0 so FSR.nva \leftarrow -1 and for FMADD FSR.nvc \leftarrow -1.

R may be any number; its generation may cause OF, UF, and/or NX.

The multiply operation in fused floating-point negative multiply-add (FNMADD) instructions cannot cause inexact, underflow, or overflow exceptions.

Note that rounding occurs after the negation. Thus, when the rounding mode is towards $\pm\infty$, FNMADD is not equivalent to FMADD followed by FNEG.

See the earlier sections on Nonstandard Mode and unfinished_FPop for additional details.

8.5.8 Floating-Point Multiply-Subtract (FMSUB)

First refer to the Floating-Point Multiply table (TABLE 8-5 on page 397) to select a row in the table below.

TABLE 8-9 Floating-Point Multiply-Subtract ((F[rs1] × F[rs2])– F[rs3])

		F[rs3]							QNaN3	SNaN3
		–∞	–N3	–0	+0	+N3	+∞			
F[rs1] × F[rs2]	–∞	dQNaN, NV	–∞						QNaN3	QNaN3, NV
	–N		±R*	–N		–R				
	–0		+N3	±0**	–0	–N3				
	+0			+0	±0**					
	+N		+R	+N		±R*				
	+∞	+∞					dQNaN, NV			
	QNaN1	QNaN1								
	QNaN2	QNaN2								
	QNaN (±0 × ±∞)	dQNaN, NV***						QNaN3, NV***		
	QNaN1	QNaN1, NV***								
	QNaN2	QNaN2, NV***								

* if N = N3, then **

** result is +0 unless rounding mode is round to –∞, in which case the result is –0

*** if FSR.nvm = 1, FSR.nvc ← 1, the trap occurs, and FSR.aexc is left unchanged; otherwise, FSR.nvm = 0 so FSR.nva ← 1 and for FMSUB FSR.nvc ← 1.

R may be any number; its generation may cause OF, UF, and/or NX.

The multiply operation in fused floating-point multiply-subtract (FMSUB) instructions cannot cause inexact, underflow, or overflow exceptions.

See the earlier sections on Nonstandard Mode and unfinished_FPop for additional details.

8.5.9 Floating-Point Negative Multiply-Subtract (FNMSUB)

First refer to the Floating-Point Multiply table (TABLE 8-5 on page 397) to select a row in the table below.

TABLE 8-10 Floating-Point Negative Multiply-Subtract ($-(F[rs1] \times F[rs2]) + F[rs3]$)

		F[rs3]							
		$-\infty$	-N3	-0	+0	+N3	$+\infty$	QNaN3	SNaN3
F[rs1] × F[rs2]	$-\infty$	dQNaN, NV					$+\infty$	QNaN3	QNaN3, NV
	-N	$-\infty$	$\pm R^*$	+N		+R			
	-0		-N3	$\pm 0^{**}$	+0	+N3			
	+0			-0	$\pm 0^{**}$				
	+N		-R	-N		$\pm R^*$			
	$+\infty$						dQNaN, NV		
	QNaN1	QNaN1							
	QNaN2	QNaN2							
	QNaN ($\pm 0 \times \pm \infty$)	dQNaN, NV ^{***}						QNaN3, NV ^{***}	
	QNaN1	QNaN1, NV ^{***}							
	QNaN2	QNaN2, NV ^{***}							

* if N = N3, then **

** result is +0 unless rounding mode is round to $-\infty$, in which case the result is -0

*** if FSR.nvm = 1, FSR.nvc ← 1, the trap occurs, and FSR.aexc is left unchanged; otherwise, FSR.nvm = 0 so FSR.nva ← 1 and for FNMSUB FSR.nvc ← 1.

R may be any number; its generation may cause OF, UF, and/or NX.

The multiply operation in fused floating-point negative multiply-subtract (FNMSUB) instructions cannot cause inexact, underflow, or overflow exceptions.

Note that rounding occurs after the negation. Thus, FNMSUB is not equivalent to FMSUB followed by FNEG when the rounding mode is towards $\pm\infty$.

See the earlier sections on Nonstandard Mode and unfinished_FPop for additional details.

8.5.10 Floating-Point Divide (FDIV)

TABLE 8-11 Floating-Point Divide operation ($F[rs1] \div F[rs2]$)

		F[rs2]						QNaN2	SNaN2	
		$-\infty$	-N2	-0	+0	+N2	$+\infty$			
F[rs1]	$-\infty$	dQNaN, NV	$+\infty$		$-\infty$		dQNaN, NV	QNaN2	QNaN2, NV	
	-N1		+R	$+\infty$, DZ	$-\infty$, DZ	-R				
	-0	+0	dQNaN, NV			-0				
	+0	-0				+0				
	+N1		-R	$-\infty$, DZ	$+\infty$, DZ	+R				
	$+\infty$	dQNaN, NV	$-\infty$		$+\infty$		dQNaN, NV			
	QNaN1	QNaN1								
	SNaN1	QNaN1, NV								

R may be any number; its generation may cause OF, UF, and/or NX.

8.5.11 Floating-Point Square Root (FSQRT)

TABLE 8-12 Floating-Point Square Root operation ($\sqrt{F[rs2]}$)

F[rs2]							
$-\infty$	-N2	-0	+0	+N2	$+\infty$	QNaN2	SNaN2
dQNaN, NV		-0	+0	+R	$+\infty$	QNaN2	QNaN2, NV

R may be any number; its generation may cause NX.

Square root cannot cause DZ, OF, or UF.

8.5.12 Floating-Point Compare (FCMP, FCMPE)

TABLE 8-13 Floating-Point Compare (FCMP, FCMPE) operation (F[rs1] ? F[rs2])

		F[rs2]							
		$-\infty$	-N2	-0	+0	+N2	$+\infty$	QNaN2	SNaN2
F[rs1]	$-\infty$	0						3, NV*	3, NV
	-N1	0, 1, 2		1					
	-0	0			1				
	+0	0			1				
	+N1	2		0,1,2		0			
	$+\infty$	0							
	QNaN1	3, NV*							
	SNaN1	3, NV							

* NV for FCMPE, but not for FCMP.

TABLE 8-14 FSR.fcc Encoding for Result of FCMP, FCMPE

fcc result	meaning
0	=
1	<
2	>
3	unordered

NaN is considered to be unequal to anything else, even the identical NaN bit pattern.

FCMP/FCMPE cannot cause DZ, OF, UF, NX.

8.5.13 Floating-Point Lexicographic Compare (FLCMP)

TABLE 8-15 Lexicographic Compare (FLCMP) operation (F[rs1] ?_L F[rs2])

		F[rs2]							
		$-\infty$	-N2	-0	+0	+N2	$+\infty$	QNaN2	SNaN2
F[rs1]	$-\infty$	0						3	
	-N1	0,1		1					
	-0	0			1				
	+0	0			1				
	+N1	0		0,1		0			
	$+\infty$	0							
	QNaN1	2							
	SNaN1	2							

Note | Encoding of the condition code result for FLCMP is different from the encoding for FCMP/FCMPE. Furthermore, for this operation, the sign of zero is significant: $-0 < +0$.

TABLE 8-16 FSR.fcc Encoding for Result of FLCMP

fcc result	Meaning
0	\geq
1	$<$
2	F[rs1] is NaN and F[rs2] is not NaN
3	F[rs2] is NaN

FLCMP does not cause any floating-point exceptions (NV, DZ, OF, UF, NX, or *fp_exception_other* (with FSR.ftt = unfinished_FPop)).

FSR.cexc and FSR.aexc are unchanged by FLCMP.

8.5.14 Floating-Point to Floating-Point Conversions (F<s | d | q>TO<s | d | q>)

TABLE 8-17 Floating-Point to Float-Point Conversions (convert(F[rs2]))

F[rs2]									
-SNaN2	-QNaN2	$-\infty$	-N2	-0	+0	+N2	$+\infty$	+QNaN2	+SNaN2
-QNaN2, NV	-QNaN2	$-\infty$	-R	-0	+0	+R	$+\infty$	+QNaN2	+QNaN2, NV

For FsTOd:

- the least-significant fraction bits of a normal number are filled with zero to fit in double-precision format
- the least-significant bits of a NaN result operand are filled with zero to fit in double-precision format

For FsTOq and FdTOq:

- the least-significant fraction bits of a normal number are filled with zero to fit in quad-precision format
- the least-significant bits of a NaN result operand are filled with zero to fit in quad-precision format

For FqTOs and FdTOs:

- the fraction is rounded according to the current rounding mode
- the lower-order bits of a NaN source are discarded to fit in single-precision format; this discarding is not considered a rounding operation, and will not cause an NX exception

For FqTOd:

- the fraction is rounded according to the current rounding mode
- the least-significant bits of a NaN source are discarded to fit in double-precision format; this discarding is not considered a rounding operation, and will not cause an NX exception

TABLE 8-18 Floating-Point to Float-Point Conversion Exception Conditions

NV	<ul style="list-style-type: none"> • SNaN operand
OF	<ul style="list-style-type: none"> • FdTOs, FqTOs: the input is larger than can be expressed in single precision • FqTOd: the input is larger than can be expressed in double precision • does not occur during other conversion operations
UF	<ul style="list-style-type: none"> • FdTOs, FqTOs: the input is smaller than can be expressed in single precision • FqTOd: the input is smaller than can be expressed in double precision • does not occur during other conversion operations
NX	<ul style="list-style-type: none"> • FdTOs, FqTOs: the input fraction has more significant bits than can be held in a single precision fraction • FqTOd: the input fraction has more significant bits than can be held in a double precision fraction • does not occur during other conversion operations

8.5.15 Floating-Point to Integer Conversions (F<s | d | q>TO<i | x>)

TABLE 8-19 Floating-Point to Integer Conversions (convert(F[rs2]))

	F[rs2]										
	-SNaN2	-QNaN2	-∞	-N2	-0	+0	+N2	+∞	+QNaN2	+SNaN2	
FdTOx FsTOx FqTOx	-2 ⁶³ , NV		-2 ⁶³ , NV	-R	0		+R	2 ⁶³ -1, NV		2 ⁶³ -1, NV	
FdTOi FsTOi FqTOi	-2 ³¹ , NV		-2 ³¹ , NV					2 ³¹ -1, NV		2 ³¹ -1, NV	

R may be any integer, and may cause NV, NX.

Float-to-Integer conversions are always treated as round-toward-zero (truncated).

These operations are invalid (due to integer overflow) under the conditions described in *Integer Overflow Definition* on page 393.

TABLE 8-20 Floating-point to Integer Conversion Exception Conditions

NV	<ul style="list-style-type: none"> • SNaN operand • QNaN operand • ±∞ operand • integer overflow
NX	<ul style="list-style-type: none"> • non-integer source (truncation occurred)

8.5.16 Integer to Floating-Point Conversions (F<i | x>TO<s | d | q>)

TABLE 8-21 Integer to Floating-Point Conversions (convert(F[rs2]))

F[rs2]		
-int	0	+int
-R	+0	+R

R may be any number; its generation may cause NX.

TABLE 8-22 Floating-Point Conversion Exception Conditions

NX	<ul style="list-style-type: none">• FxTOd, FxTOs, FiTOs (possible loss of precision)• not applicable to FiTOd, FxTOq, or FiTOq (FSR.cexc will always be cleared)
----	---

Memory

The Oracle SPARC Architecture *memory models* define the semantics of memory operations. The instruction set semantics require that loads and stores behave *as if* they are performed in the order in which they appear in the dynamic control flow of the program. The *actual* order in which they are processed by the memory may be different. The purpose of the memory models is to specify what constraints, if any, are placed on the order of memory operations.

The memory models apply both to uniprocessor and to shared memory multiprocessors. Formal memory models are necessary for precise definitions of the interactions between multiple virtual processors and input/output devices in a shared memory configuration. Programming shared memory multiprocessors requires a detailed understanding of the operative memory model and the ability to specify memory operations at a low level in order to build programs that can safely and reliably coordinate their activities.

This chapter contains a great deal of theoretical information so that the discussion of the Oracle SPARC Architecture TSO memory model has sufficient background.

This chapter describes memory models in these sections:

- **Memory Location Identification** on page 407.
- **Memory Accesses and Cacheability** on page 407.
- **Memory Addressing and Alternate Address Spaces** on page 409.
- **SPARC V9 Memory Model** on page 412.
- **The Oracle SPARC Architecture Memory Model — TSO** on page 415.
- **Nonfaulting Load** on page 421.
- **Store Coalescing** on page 422.

9.1 Memory Location Identification

A memory location is identified by an 8-bit address space identifier (ASI) and a 64-bit memory address. The 8-bit ASI can be obtained from an ASI register or included in a memory access instruction. The ASI used for an access can distinguish among different 64-bit address spaces, such as Primary memory space, Secondary memory space, and internal control registers. It can also apply attributes to the access, such as whether the access should be performed in big- or little-endian byte order, or whether the address should be taken as a virtual or real.

9.2 Memory Accesses and Cacheability

Memory is logically divided into real memory (cached) and I/O memory (noncached with and without side effects) spaces.

Real memory stores information without side effects. A load operation returns the value most recently stored. Operations are side-effect-free in the sense that a load, store, or atomic load-store to a location in real memory has no program-observable effect, except upon that location (or, in the case of a load or load-store, on the destination register).

I/O locations may not behave like memory and may have side effects. Load, store, and atomic load-store operations performed on I/O locations may have observable side effects, and loads may not return the value most recently stored. The value semantics of operations on I/O locations are *not* defined by the memory models, but the constraints on the order in which operations are performed is the same as it would be if the I/O locations were real memory. The storage properties, contents, semantics, ASI assignments, and addresses of I/O registers are implementation dependent.

9.2.1 Coherence Domains

Two types of memory operations are supported in the Oracle SPARC Architecture: cacheable and noncacheable accesses. The manner in which addresses are differentiated is implementation dependent. In some implementations, it is indicated in the page translation entry (TTE.cp).

Although SPARC V9 does not specify memory ordering between cacheable and noncacheable accesses, the Oracle SPARC Architecture maintains TSO ordering between memory references regardless of their cacheability.

9.2.1.1 Cacheable Accesses

Accesses within the coherence domain are called cacheable accesses. They have these properties:

- Data reside in real memory locations.
- Accesses observe supported cache coherency protocol(s).
- The cache line size is 2^n bytes (where $n \geq 4$), and can be different for each cache.

9.2.1.2 Noncacheable Accesses

Noncacheable accesses are outside of the coherence domain. They have the following properties:

- Data might not reside in real memory locations. Accesses may result in programmer-visible side effects. An example is memory-mapped I/O control registers.
- Accesses do not observe supported cache coherency protocol(s).
- The smallest unit in each transaction is a single byte.

The Oracle SPARC Architecture MMU optionally includes an attribute bit in each page translation, TTE.e, which when set signifies that this page has side effects.

Noncacheable accesses without side effects (TTE.e = 0) are processor-consistent and obey TSO memory ordering. In particular, processor consistency ensures that a noncacheable load that references the same location as a previous noncacheable store will load the data from the previous store.

Noncacheable accesses with side effects (TTE.e = 1) are processor consistent and are strongly ordered. These accesses are described in more detail in the following section.

9.2.1.3 Noncacheable Accesses with Side-Effect

Loads, stores, and load-stores to I/O locations might not behave with memory semantics. Loads and stores could have side effects; for example, a read access could clear a register or pop an entry off a FIFO. A write access could set a register address port so that the next access to that address will read or write a particular internal register. Such devices are considered order sensitive. Also, such devices may only allow accesses of a fixed size, so store merging of adjacent stores or stores within a 16-byte region would cause an error (see *Store Coalescing* on page 422).

Noncacheable accesses (other than block loads and block stores) to pages with side effects (TTE.e = 1) exhibit the following behavior:

- Noncacheable accesses are strongly ordered with respect to each other. Bus protocol should guarantee that IO transactions to the same device are delivered in the order that they are received.
- Noncacheable loads with the TTE.e bit = 1 will not be issued to the system until all previous instructions have completed, and the store queue is empty.
- Noncacheable store coalescing is disabled for accesses with TTE.e = 1.
- A MEMBAR may be needed between side-effect and non-side-effect accesses. See TABLE 9-3 on page 419.

Whether block loads and block stores adhere to the above behavior or ignore TTE.e and always behave as if TTE.e = 0 is implementation-dependent (impl. dep. #410-S10, #411-S10).

On Oracle SPARC Architecture virtual processors, noncacheable and side-effect accesses do not observe supported cache coherency protocols (impl. dep. #120).

Non-faulting loads (using ASI_PRIMARY_NO_FAULT[_LITTLE] or ASI_SECONDARY_NO_FAULT[_LITTLE]) with the TTE.e bit = 1 cause a *DAE_side_effect_page* trap.

Prefetches to noncacheable addresses result in nops.

The processor does speculative instruction memory accesses and follows branches that it predicts are taken. Instruction addresses mapped by the MMU can be accessed even though they are not actually executed by the program. Normally, locations with side effects or that generate timeouts or bus errors are not mapped as instruction addresses by the MMU, so these speculative accesses will not cause problems.

IMPL. DEP. #118-V9: The manner in which I/O locations are identified is implementation dependent.

IMPL. DEP. #120-V9: The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent.

V9 Compatibility Note	Operations to I/O locations are <i>not</i> guaranteed to be sequentially consistent among themselves, as they are in SPARC V8.
------------------------------	--

Systems supporting SPARC V8 applications that use memory-mapped I/O locations must ensure that SPARC V8 sequential consistency of I/O locations can be maintained when those locations are referenced by a SPARC V8 application. The MMU either must enforce such consistency or cooperate with system software or the virtual processor to provide it.

IMPL. DEP. #121-V9: An implementation may choose to identify certain addresses and use an implementation-dependent memory model for references to them.

9.3 Memory Addressing and Alternate Address Spaces

An address in SPARC V9 is a tuple consisting of an 8-bit address space identifier (ASI) and a 64-bit byte-address offset within the specified address space. Memory is byte-addressed, with halfword accesses aligned on 2-byte boundaries, word accesses (which include instruction fetches) aligned on 4-byte boundaries, extended-word and doubleword accesses aligned on 8-byte boundaries, and quadword quantities aligned on 16-byte boundaries. With the possible exception of the cases described in *Memory Alignment Restrictions* on page 79, an improperly aligned address in a load, store, or load-store instruction always causes a trap to occur. The largest datum that is guaranteed to be

atomically read or written is an aligned doubleword¹. Also, memory references to different bytes, halfwords, and words in a given doubleword are treated for ordering purposes as references to the same location. Thus, the unit of ordering for memory is a doubleword.

Notes | The doubleword is the coherency unit for update, but programmers should not assume that doubleword floating-point values are updated as a unit unless they are doubleword-aligned and always updated with double-precision loads and stores. Some programs use pairs of single-precision operations to load and store double-precision floating-point values when the compiler cannot determine that they are doubleword aligned. Also, although quad-precision operations are defined in the SPARC V9 architecture, the granularity of loads and stores for quad-precision floating-point values may be word or doubleword.

9.3.1 Memory Addressing Types

The Oracle SPARC Architecture supports the following types of memory addressing:

Virtual Addresses (VA). Virtual addresses are addresses produced by a virtual processor that maps all systemwide, program-visible memory. Virtual addresses can be presented in nonprivileged mode and privileged mode

Real addresses (RA). A real address is provided to privileged software to describe the underlying physical memory allocated to it. Translation storage buffers (TSBs) maintained by privileged software are used to translate privileged or nonprivileged mode virtual addresses into real addresses. MMU bypass addresses in privileged mode are also real addresses.

Nonprivileged software only uses virtual addresses. Privileged software uses virtual and real addresses.

9.3.2 Memory Address Spaces

The Oracle SPARC Architecture supports accessing memory using virtual or real addresses. Multiple virtual address spaces within the same real address space are distinguished by a *context identifier* (context ID).

Privileged software can create multiple virtual address spaces, using the primary and secondary context registers to associate a context ID with every virtual address. Privileged software manages the allocation of context IDs.

The full representation of a real address is as follows:

$$\text{real_address} = \text{context_ID} :: \text{virtual_address}$$

9.3.3 Address Space Identifiers

The virtual processor provides an address space identifier with every address. This ASI may serve several purposes:

- To identify which of several distinguished address spaces the 64-bit address offset is addressing

¹ Two exceptions to this are the special `ASI_TWIN_DW_NUCLEUS[_L]` and `ASI_TWINX_REAL[_L]` which provide hardware support for an atomic quad load to be used for TTE loads from TSBs.

- To provide additional access control and attribute information, for example, to specify the endianness of the reference
- To specify the address of an internal control register in the virtual processor, cache, or memory management hardware

Memory management hardware can associate an independent 2^{64} -byte memory address space with each ASI. In practice, the three independent memory address spaces (contexts) created by the MMU are Primary, Secondary, and Nucleus.

Programming Note	Independent address spaces, accessible through ASIs, make it possible for system software to easily access the address space of faulting software when processing exceptions or to implement access to a client program's memory space by a server program.
-------------------------	---

Alternate-space load, store, load-store and prefetch instructions specify an *explicit* ASI to use for their data access. The behavior of the access depends on the current privilege mode.

Non-alternate space load, store, load-store, and prefetch instructions use an *implicit* ASI value that is determined by current virtual processor state (the current privilege mode, trap level (TL), and the value of `PSTATE.cle`). Instruction fetches use an implicit ASI that depends only on the current mode and trap level.

The architecturally specified ASIs are listed in Chapter 10, *Address Space Identifiers (ASIs)*. The operation of each ASI in nonprivileged and privileged modes is indicated in TABLE 10-1 on page 425.

Attempts by nonprivileged software (`PSTATE.priv = 0`) to access restricted ASIs (ASI bit 7 = 0) cause a *privileged_action* exception. Attempts by privileged software (`PSTATE.priv = 1`) to access ASIs 30_{16} – $7F_{16}$ cause a *privileged_action* exception.

When `TL = 0`, normal accesses by the virtual processor to memory when fetching instructions and performing loads and stores implicitly specify `ASI_PRIMARY` or `ASI_PRIMARY_LITTLE`, depending on the setting of `PSTATE.cle`.

When `TL = 1` or `2` (> 0 but $\leq \text{MAXPTL}$), the implicit ASI in privileged mode is:

- for instruction fetches, `ASI_NUCLEUS`
- for loads and stores, `ASI_NUCLEUS` if `PSTATE.cle = 0` or `ASI_NUCLEUS_LITTLE` if `PSTATE.cle = 1` (impl. dep. #124-V9)

SPARC V9 supports the `PRIMARY[_LITTLE]`, `SECONDARY[_LITTLE]`, and `NUCLEUS[_LITTLE]` address spaces.

Accesses to other address spaces use the load/store alternate instructions. For these accesses, the ASI is either contained in the instruction (for the register+register addressing mode) or taken from the ASI register (for register+immediate addressing).

ASIs are either nonrestricted or restricted-to-privileged:

- A nonrestricted ASI (ASI range 80_{16} – FF_{16}) is one that may be used independently of the privilege level (`PSTATE.priv`) at which the virtual processor is running.
- A restricted-to-privileged ASI (ASI range 00_{16} – $2F_{16}$) requires that the virtual processor be in privileged mode for a legal access to occur.

The relationship between virtual processor state and ASI restriction is shown in TABLE 9-1.

TABLE 9-1 Allowed Accesses to ASIs

ASI Value	Type	Result of ASI Access in NP Mode	Result of ASI Access in P Mode
00 ₁₆ – 2F ₁₆	Restricted-to-privileged	<i>privileged_action</i> exception	Valid Access
80 ₁₆ – FF ₁₆	Nonrestricted	Valid Access	Valid Access

Some restricted ASIs are provided as mandated by SPARC V9:

ASI_AS_IF_USER_PRIMARY[_LITTLE] and ASI_AS_IF_USER_SECONDARY[_LITTLE]. The intent of these ASIs is to give privileged software efficient, yet secure access to the memory space of nonprivileged software.

The normal address space is *primary address space*, which is accessed by the unrestricted ASI_PRIMARY[_LITTLE] ASIs. The *secondary address space*, which is accessed by the unrestricted ASI_SECONDARY[_LITTLE] ASIs, is provided to allow server software to access client software's address space.

ASI_PRIMARY_NOFAULT[_LITTLE] and ASI_SECONDARY_NOFAULT[_LITTLE] support *nonfaulting loads*. These ASIs may be used to color (that is, distinguish into classes) loads in the instruction stream so that, in combination with a judicious mapping of low memory and a specialized trap handler, an optimizing compiler can move loads outside of conditional control structures.

9.4 SPARC V9 Memory Model

The SPARC V9 processor architecture specified the organization and structure of a central processing unit but did not specify a memory system architecture. This section summarizes the MMU support required by an Oracle SPARC Architecture processor.

The memory models specify the possible order relationships between memory-reference instructions issued by a virtual processor and the order and visibility of those instructions as seen by other virtual processors. The memory model is intimately intertwined with the program execution model for instructions.

9.4.1 SPARC V9 Program Execution Model

The SPARC V9 strand model of a virtual processor consists of three units: an Issue Unit, a Reorder Unit, and an Execute Unit, as shown in FIGURE 9-1.

The Issue Unit reads instructions over the instruction path from memory and issues them in *program order to the Reorder Unit*. Program order is precisely the order determined by the control flow of the program and the instruction semantics, under the assumption that each instruction is performed independently and sequentially.

Issued instructions are collected and potentially reordered in the Reorder Unit, and then dispatched to the Execute Unit. Instruction reordering allows an implementation to perform some operations in parallel and to better allocate resources. The reordering of instructions is constrained to ensure that the results of program execution are the same as they would be if the instructions were performed in program order. This property is called *processor self-consistency*.

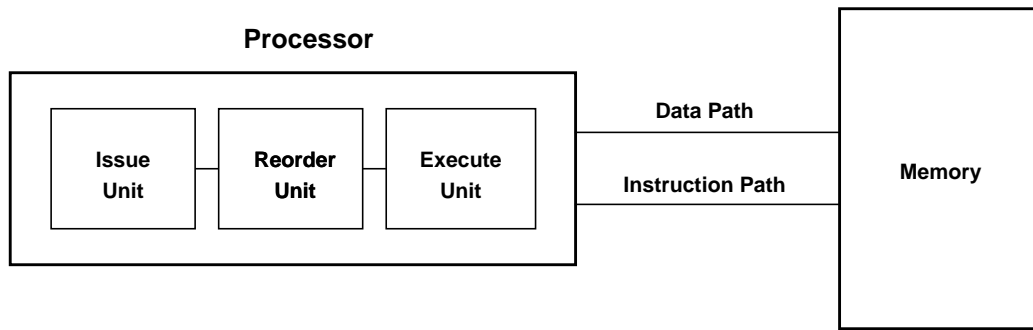


FIGURE 9-1 Processor Model: Uniprocessor System

Processor self-consistency requires that the result of execution, in the absence of any shared memory interaction with another virtual processor, be identical to the result that would be observed if the instructions were performed in program order. In the model in FIGURE 9-1, instructions are issued in program order and placed in the reorder buffer. The virtual processor is allowed to reorder instructions, provided it does not violate any of the data-flow constraints for registers or for memory.

The data-flow order constraints for register reference instructions are these:

1. An instruction that reads from or writes to a register cannot be performed until all earlier instructions that write to that register have been performed (read-after-write hazard; write-after-write hazard).
2. An instruction cannot be performed that writes to a register until all earlier instructions that read that register have been performed (write-after-read hazard).

V9 Compatibility | An implementation can avoid blocking instruction execution in
Note | case 2 and the write-after-write hazard in case 1 by using a
 renaming mechanism that provides the old value of the register
 to earlier instructions and the new value to later uses.

The data-flow order constraints for memory-reference instructions are those for register reference instructions, plus the following additional constraints:

1. A memory-reference instruction that uses (loads or stores) the value at a location cannot be performed until all earlier memory-reference instructions that set (store to) that location have been performed (read-after-write hazard, write-after-write hazard).
2. A memory-reference instruction that writes (stores to) a location cannot be performed until all previous instructions that read (load from) that location have been performed (write-after-read hazard).

Memory-barrier instruction (MEMBAR) and the TSO memory model also constrain the issue of memory-reference instructions. See *Memory Ordering and Synchronization* on page 418 and *The Oracle SPARC Architecture Memory Model — TSO* on page 415 for a detailed description.

The constraints on instruction execution assert a partial ordering on the instructions in the reorder buffer. Every one of the several possible orderings is a legal execution ordering for the program. See Appendix D in the SPARC V9 specification, *Formal Specification of the Memory Models*, for more information.

9.4.2 Virtual Processor/Memory Interface Model

Each Oracle SPARC Architecture virtual processor in a multiprocessor system is modeled as shown in FIGURE 9-2; that is, having two independent paths to memory: one for instructions and one for data.

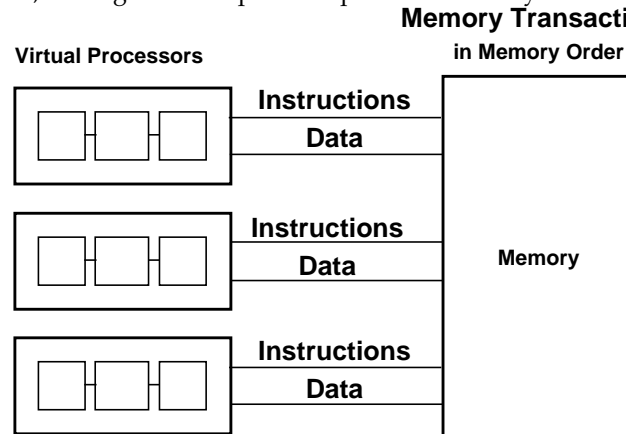


FIGURE 9-2 Data Memory Paths: Multiprocessor System

Data caches are maintained by hardware so their contents always appear to be consistent (coherent). Instruction caches are *not* required to be kept consistent with data caches and therefore require explicit program (software) action to ensure consistency when a program modifies an executing instruction stream. See *Synchronizing Instruction and Data Memory* on page 420 for details. Memory is shared in terms of address space, but it may be nonhomogeneous and distributed in an implementation. Caches are ignored in the model, since their functions are transparent to the memory model¹.

In real systems, addresses may have attributes that the virtual processor must respect. The virtual processor executes loads, stores, and atomic load-stores in whatever order it chooses, as constrained by program order and the memory model.

Instructions are performed in an order constrained by local dependencies. Using this dependency ordering, an execution unit submits one or more pending memory transactions to the memory. The memory performs transactions in *memory order*. The memory unit may perform transactions submitted to it out of order; hence, the execution unit must not concurrently submit two or more transactions that are required to be ordered, unless the memory unit can still guarantee in-order semantics.

The memory accepts transactions, performs them, and then acknowledges their completion. Multiple memory operations may be in progress at any time and may be initiated in a nondeterministic fashion in any order, provided that all transactions to a location preserve the per-virtual processor partial orderings. Memory transactions may complete in any order. Once initiated, all memory operations are performed atomically: loads from one location all see the same value, and the result of stores is visible to all potential requestors at the same instant.

The order of memory operations observed at a single location is a *total order* that preserves the partial orderings of each virtual processor's transactions to this address. There may be many legal total orders for a given program's execution.

¹. The model described here is only a model; implementations of Oracle SPARC Architecture systems are unconstrained as long as their observable behaviors match those of the model.

9.5 The Oracle SPARC Architecture Memory Model — TSO

The Oracle SPARC Architecture is a *model* that specifies the behavior observable by software on Oracle SPARC Architecture systems. Therefore, access to memory can be implemented in any manner, as long as the behavior observed by software conforms to that of the models described here.

The SPARC V9 architecture defines three different memory models: *Total Store Order (TSO)*, *Partial Store Order (PSO)*, and *Relaxed Memory Order (RMO)*.

All SPARC V9 processors must provide Total Store Order (or a more strongly ordered model, for example, Sequential Consistency) to ensure compatibility for SPARC V8 application software.

All Oracle SPARC Architecture virtual processors implement TSO ordering. The PSO and RMO models from SPARC V9 are not described in this Oracle SPARC Architecture specification. Oracle SPARC Architecture 2015 processors do not implement the PSO memory model directly, but all software written to run under PSO will execute correctly on an Oracle SPARC Architecture 2015 processor (using the TSO model).

Whether memory models represented by `PSTATE.mm = 102` or `112` are supported in an Oracle SPARC Architecture processor is implementation dependent (impl. dep. #113-V9-Ms10). If the `102` model is supported, then when `PSTATE.mm = 102` the implementation must correctly execute software that adheres to the RMO model described in *The SPARC Architecture Manual-Version 9*. If the `112` model is supported, its definition is implementation dependent and will be described in implementation-specific documentation.

Programs written for Relaxed Memory Order will work in both Partial Store Order and Total Store Order. Programs written for Partial Store Order will work in Total Store Order. Programs written for a weak model, such as RMO, may execute more quickly when run on hardware directly supporting that model, since the model exposes more scheduling opportunities, but use of that model may also require extra instructions to ensure synchronization. Multiprocessor programs written for a stronger model will behave unpredictably if run in a weaker model.

Machines that implement *sequential consistency* (also called "strong ordering" or "strong consistency") automatically support programs written for TSO. Sequential consistency is not a SPARC V9 memory model. In sequential consistency, the loads, stores, and atomic load-stores of all virtual processors are performed by memory in a serial order that conforms to the order in which these instructions are issued by individual virtual processors. A machine that implements sequential consistency may deliver lower performance than an equivalent machine that implements TSO order. Although particular SPARC V9 implementations may support sequential consistency, portable software must not rely on the sequential consistency memory model.

9.5.1 Memory Model Selection

The active memory model is specified by the 2-bit value in `PSTATE.mm`. The value `002` represents the TSO memory model; increasing values of `PSTATE.mm` indicate increasingly weaker (less strongly ordered) memory models.

Writing a new value into `PSTATE.mm` causes subsequent memory reference instructions to be performed with the order constraints of the specified memory model.

IMPL. DEP. #119-Ms10: The effect of an attempt to write an unsupported memory model designation into `PSTATE.mm` is implementation dependent; however, it should never result in a value of `PSTATE.mm` value greater than the one that was written. In the case of an Oracle SPARC Architecture implementation that only supports the TSO memory model, `PSTATE.mm` always reads as zero and attempts to write to it are ignored.

9.5.2 Programmer-Visible Properties of the Oracle SPARC Architecture TSO Model

Total Store Order must be provided for compatibility with existing SPARC V8 programs. Programs that execute correctly in either RMO or PSO will execute correctly in the TSO model.

The rules for TSO, in addition to those required for self-consistency (see page 413), are:

- Loads are blocking and ordered with respect to earlier loads
- Stores are ordered with respect to stores.
- Atomic load-stores are ordered with respect to loads and stores.
- Stores cannot bypass earlier loads.

Programming Note | Loads *can* bypass earlier stores to other addresses, which maintains processor self-consistency.

Atomic load-stores are treated as both a load and a store and can only be applied to cacheable address spaces.

Thus, TSO ensures the following behavior:

- Each load instruction behaves as if it were followed by a MEMBAR `#LoadLoad` and `#LoadStore`.
- Each store instruction behaves as if it were followed by a MEMBAR `#StoreStore`.
- Each atomic load-store behaves as if it were followed by a MEMBAR `#LoadLoad`, `#LoadStore`, `#StoreLoad`, and `#StoreStore`.

In addition to the above TSO rules, the following rules apply to Oracle SPARC Architecture memory models:

- A MEMBAR `#StoreLoad` must be used to prevent a load from bypassing a prior store, if Strong Sequential Order (as defined in *The Oracle SPARC Architecture Memory Model — TSO* on page 415) is desired.
- Accesses that have side effects are all strongly ordered with respect to each other.
- A MEMBAR `#LookasideD` is not needed between a store and a subsequent load to the same noncacheable address.
- Load (LDXA) and store (STXA) instructions that reference certain internal ASIs perform both an intra-virtual processor synchronization (i.e. an implicit MEMBAR `#Sync` operation before the load or store is executed) and an inter-virtual processor synchronization (that is, all active virtual processors are brought to a point where synchronization is possible, the load or store is executed, and all virtual processors then resume instruction fetch and execution). The model-specific PRM should indicate which ASIs require intra-virtual processor synchronization, inter-virtual processor synchronization, or both.

9.5.3 TSO Ordering Rules

TABLE 9-2 summarizes the cases where a MEMBAR must be inserted between two memory operations on an Oracle SPARC Architecture virtual processor running in TSO mode, to ensure that the operations appear to complete in program order. Memory operation *ordering* is not to be confused with processor consistency or deterministic operation; MEMBARs are required for deterministic operation of certain ASI register updates.

Programming Note To ensure software portability across systems, the MEMBAR rules in this section should be followed (which may be stronger than the rules in SPARC V9).

TABLE 9-2 is to be read as follows: Reading from row to column, the first memory operation in program order in a row is followed by the memory operation found in the column. Symbols used as table entries:

- # — No intervening operation is required.
- M — an intervening MEMBAR #StoreLoad or MEMBAR #Sync or MEMBAR #MemIssue is required
- S — an intervening MEMBAR #Sync or MEMBAR #MemIssue is required
- nc — Noncacheable
- e — Side effect
- ne — No side effect

TABLE 9-2 Summary of Oracle SPARC Architecture Ordering Rules (TSO Memory Model)

From Memory Operation R (row):	To Memory Operation C (column):										
	load	store	atomic	load	bstore	load_nc_e	store_nc_e	load_nc_ne	store_nc_ne	load_nc	bstore_nc
load	#	#	#	S	S	#	#	#	#	S	S
store	M ²	#	#	M	S	M	#	M	#	M	S
atomic	#	#	#	M	S	#	#	#	#	M	S
load	S	S	S	S	S	S	S	S	S	S	S
bstore	M	S	M	M	S	M	S	M	S	M	S
load_nc_e	#	#	#	S	S	# ¹	# ¹	# ¹	# ¹	S	S
store_nc_e	S	#	#	S	S	# ¹	# ¹	M ²	# ¹	M	S
load_nc_ne	#	#	#	S	S	# ¹	# ¹	# ¹	# ¹	S	S
store_nc_ne	S	#	#	S	S	M ²	# ¹	M ²	# ¹	M	S
load_nc	S	S	S	S	S	S	S	S	S	S	S
bstore_nc	S	S	S	S	S	M	S	M	S	M	S

1. This table assumes that both noncacheable operations access the same device.

2. When the store and subsequent load access the *same* location, no intervening MEMBAR is required.

Note that transitivity applies; if operation X is always ordered before operation Y ("#" in TABLE 9-2) and operation Y is always ordered before operation Z (again, "#" in the table), then the sequence of operations X ... Y ... Z may safely be executed with no intervening MEMBAR, even if the table shows that a MEMBAR is normally needed between X and Z. For example, a MEMBAR is normally needed between a store and a load ("M" in TABLE 9-2); however, the sequence "store ... atomic ... load" may be executed safely with no intervening MEMBAR because stores are always ordered before atomics and atomics are always ordered before loads.

9.5.4 Hardware Primitives for Mutual Exclusion

In addition to providing memory-ordering primitives that allow programmers to construct mutual-exclusion mechanisms in software, the Oracle SPARC Architecture provides some hardware primitives for mutual exclusion:

- Compare and Swap (CASA and CASXA)
- Load Store Unsigned Byte (LDSTUB and LDSTUBA)
- Swap (SWAP and SWAPA)

Each of these instructions has the semantics of both a load and a store in all three memory models. They are all *atomic*, in the sense that no other store to the same location can be performed between the load and store elements of the instruction. All of the hardware mutual-exclusion operations conform to the TSO memory model and may require barrier instructions to ensure proper data visibility.

Atomic load-store instructions can be used only in the cacheable domains (not in noncacheable I/O addresses). An attempt to use an atomic load-store instruction to access a noncacheable page results in a *DAE_nc_page* exception.

The atomic load-store alternate instructions can use a limited set of the ASIs. See the specific instruction descriptions for a list of the valid ASIs. An attempt to execute an atomic load-store alternate instruction with an invalid ASI results in a *DAE_invalid_asi* exception.

9.5.4.1 Compare-and-Swap (CASA, CASXA)

Compare-and-swap is an atomic operation that compares a value in a virtual processor register to a value in memory and, if and only if they are equal, swaps the value in memory with the value in a second virtual processor register. Both 32-bit (CASA) and 64-bit (CASXA) operations are provided. The compare-and-swap operation is atomic in the sense that once it begins, no other virtual processor can access the memory location specified until the compare has completed and the swap (if any) has also completed and is potentially visible to all other virtual processors in the system.

Compare-and-swap is substantially more powerful than the other hardware synchronization primitives. It has an infinite consensus number; that is, it can resolve, in a wait-free fashion, an infinite number of contending processes. Because of this property, compare-and-swap can be used to construct wait-free algorithms that do not require the use of locks.

9.5.4.2 Swap (SWAP)

SWAP atomically exchanges the lower 32 bits in a virtual processor register with a word in memory. SWAP has a consensus number of two; that is, it cannot resolve more than two contending processes in a wait-free fashion.

9.5.4.3 Load Store Unsigned Byte (LDSTUB)

LDSTUB loads a byte value from memory to a register and writes the value FF_{16} into the addressed byte atomically. LDSTUB is the classic test-and-set instruction. Like SWAP, it has a consensus number of two and so cannot resolve more than two contending processes in a wait-free fashion.

9.5.5 Memory Ordering and Synchronization

The Oracle SPARC Architecture provides some level of programmer control over memory ordering and synchronization through the MEMBAR and FLUSH instructions.

MEMBAR serves two distinct functions in SPARC V9. One variant of the MEMBAR, the ordering MEMBAR, provides a way for the programmer to control the order of loads and stores issued by a virtual processor. The other variant of MEMBAR, the sequencing MEMBAR, enables the programmer

to explicitly control order and completion for memory operations. Sequencing MEMBARs are needed only when a program requires that the effect of an operation becomes globally visible rather than simply being scheduled.¹ Because both forms are bit-encoded into the instruction, a single MEMBAR can function both as an ordering MEMBAR and as a sequencing MEMBAR.

The SPARC V9 instruction set architecture does not guarantee consistency between instruction and data spaces. A problem arises when instruction space is dynamically modified by a program writing to memory locations containing instructions (Self-Modifying Code). Examples are Lisp, debuggers, and dynamic linking. The FLUSH instruction synchronizes instruction and data memory after instruction space has been modified.

9.5.5.1 Ordering MEMBAR Instructions

Ordering MEMBAR instructions induce an ordering in the instruction stream of a single virtual processor. Sets of loads and stores that appear before the MEMBAR in program order are ordered with respect to sets of loads and stores that follow the MEMBAR in program order. Atomic operations (LDSTUB[A], SWAP[A], CASA, and CASXA) are ordered by MEMBAR as if they were both a load and a store, since they share the semantics of both. An STBAR instruction, with semantics that are a subset of MEMBAR, is provided for SPARC V8 compatibility. MEMBAR and STBAR operate on all pending memory operations in the reorder buffer, independently of their address or ASI, ordering them with respect to all future memory operations. This ordering applies only to memory-reference instructions issued by the virtual processor issuing the MEMBAR. Memory-reference instructions issued by other virtual processors are unaffected.

The ordering relationships are bit-encoded as shown in TABLE 9-3. For example, MEMBAR 01₁₆, written as “membar #LoadLoad” in assembly language, requires that all load operations appearing before the MEMBAR in program order complete before any of the load operations following the MEMBAR in program order complete. Store operations are unconstrained in this case. MEMBAR 08₁₆ (#StoreStore) is equivalent to the STBAR instruction; it requires that the values stored by store instructions appearing in program order prior to the STBAR instruction be visible to other virtual processors before issuing any store operations that appear in program order following the STBAR.

In TABLE 9-3 these ordering relationships are specified by the “<*m*” symbol, which signifies memory order. See Appendix D in the SPARC V9 specification, *Formal Specificatin of the Memory Models*, for a formal description of the <*m* relationship.

TABLE 9-3 Ordering Relationships Selected by Mask

Ordering Relation, Earlier < <i>m</i> Later	Assembly Language Constant Mnemonic	Effective Behavior in TSO model	Mask Value	nmask Bit #
Load < <i>m</i> Load	#LoadLoad	nop	01 ₁₆	0
Store < <i>m</i> Load	#StoreLoad	#StoreLoad	02 ₁₆	1
Load < <i>m</i> Store	#LoadStore	nop	04 ₁₆	2
Store < <i>m</i> Store	#StoreStore	nop	08 ₁₆	3

Implementation Note | An Oracle SPARC Architecture 2015 implementation that only implements the TSO memory model may implement MEMBAR #LoadLoad, MEMBAR #LoadStore, and MEMBAR #StoreStore as nops and MEMBAR #Storeload as a MEMBAR #Sync.

9.5.5.2 Sequencing MEMBAR Instructions

A sequencing MEMBAR exerts explicit control over the completion of operations. The three sequencing MEMBAR options each have a different degree of control and a different application.

¹Sequencing MEMBARs are needed for some input/output operations, forcing stores into specialized stable storage, context switching, and occasional other system functions. Using a sequencing MEMBAR when one is not needed may cause a degradation of performance.

- **Lookaside Barrier (deprecated)** — Ensures that loads following this MEMBAR are from memory and not from a lookaside into a write buffer. Lookaside Barrier requires that pending stores issued prior to the MEMBAR be completed before any load from that address following the MEMBAR may be issued. A Lookaside Barrier MEMBAR may be needed to provide lock fairness and to support some plausible I/O location semantics.
- **Memory Issue Barrier** — Ensures that all memory operations appearing in program order before the sequencing MEMBAR complete before any new memory operation may be initiated.
- **Synchronization Barrier** — Ensures that all instructions (memory reference and others) preceding the MEMBAR complete and that the effects of any fault or error have become visible before any instruction following the MEMBAR in program order is initiated. A Synchronization Barrier MEMBAR fully synchronizes the virtual processor that issues it.

TABLE 9-4 shows the encoding of these functions in the MEMBAR instruction.

TABLE 9-4 Sequencing Barrier Selected by Mask

Sequencing Function	Assembler Tag	Mask Value	cmask Bit #
Lookaside Barrier (deprecated)	#Lookaside ^D	10 ₁₆	0
Memory Issue Barrier	#MemIssue	20 ₁₆	1
Synchronization Barrier	#Sync	40 ₁₆	2

Implementation Note | In Oracle SPARC Architecture 2015 implementations, MEMBAR #Lookaside^D and MEMBAR #MemIssue are typically implemented as a MEMBAR #Sync.

For more details, see the MEMBAR instruction on page 269 of Chapter 7, *Instructions*.

9.5.5.3 Synchronizing Instruction and Data Memory

The SPARC V9 memory models do not require that instruction and data memory images be consistent at all times. The instruction and data memory images may become inconsistent if a program writes into the instruction stream. As a result, whenever instructions are modified by a program in a context where the data (that is, the instructions) in the memory and the data cache hierarchy may be inconsistent with instructions in the instruction cache hierarchy, some special programmatic (software) action must be taken.

The FLUSH instruction will ensure consistency between the in-flight instruction stream and the data references in the virtual processor executing FLUSH. The programmer must ensure that the modification sequence is robust under multiple updates and concurrent execution. Since, in general, loads and stores may be performed out of order, appropriate MEMBAR and FLUSH instructions must be interspersed as needed to control the order in which the instruction data are modified.

The FLUSH instruction ensures that subsequent instruction fetches from the doubleword target of the FLUSH by the virtual processor executing the FLUSH appear to execute after any loads, stores, and atomic load-stores issued by the virtual processor to that address prior to the FLUSH. FLUSH acts as a barrier for instruction fetches in the virtual processor on which it executes and has the properties of a store with respect to MEMBAR operations.

IMPL. DEP. #122-V9: The latency between the execution of FLUSH on one virtual processor and the point at which the modified instructions have replaced outdated instructions in a multiprocessor is implementation dependent.

Programming Note	Because FLUSH is designed to act on a doubleword and because, on some implementations, FLUSH may trap to system software, it is recommended that system software provide a user-callable service routine for flushing arbitrarily sized regions of memory. On some implementations, this routine would issue a series of FLUSH instructions; on others, it might issue a single trap to system software that would then flush the entire region.
-------------------------	--

On an Oracle SPARC Architecture virtual processor:

- A FLUSH instruction causes a synchronization with the virtual processor, which flushes the instruction pipeline in the virtual processor on which the FLUSH instruction is executed.
- Coherency between instruction and data memories may or may not be maintained by hardware. If it is, an Oracle SPARC Architecture implementation may ignore the address in the operands of a FLUSH instruction.

Programming Note	Oracle SPARC Architecture virtual processors are not required to maintain coherency between instruction and data caches in hardware. Therefore, portable software must do the following: (1) must always assume that store instructions (except Block Store with Commit) do not coherently update instruction cache(s); (2) must, in every FLUSH instruction, supply the address of the instruction or instructions that were modified.
-------------------------	---

For more details, see the FLUSH instruction on page 171 of Chapter 7, *Instructions*.

9.6 Nonfaulting Load

A nonfaulting load behaves like a normal load, with the following exceptions:

- A nonfaulting load from a location with side effects (`TTE.e = 1`) causes a *DAE_side_effect_page* exception.
- A nonfaulting load from a page marked for nonfault access only (`TTE.nfo = 1`) is allowed; other types of accesses to such a page cause a *DAE_nfo_page* exception.
- These loads are issued with `ASI_PRIMARY_NO_FAULT[_LITTLE]` or `ASI_SECONDARY_NO_FAULT[_LITTLE]`. A store with a `NO_FAULT` ASI causes a *DAE_invalid_asi* exception.

Typically, optimizers use nonfaulting loads to move loads across conditional control structures that guard their use. This technique potentially increases the distance between a load of data and the first use of that data, in order to hide latency. The technique allows more flexibility in instruction scheduling and improves performance in certain algorithms by removing address checking from the critical code path.

For example, when following a linked list, nonfaulting loads allow the null pointer to be accessed safely in a speculative, read-ahead fashion; the page at virtual address 0_{16} can safely be accessed with no penalty. The `TTE.nfo` bit marks pages that are mapped for safe access by nonfaulting loads but that can still cause a trap by other, normal accesses.

Thus, programmers can trap on “wild” pointer references—many programmers count on an exception being generated when accessing address 0_{16} to debug software—while benefiting from the acceleration of nonfaulting access in debugged library routines.

9.7 Store Coalescing

Cacheable stores may be coalesced with adjacent cacheable stores within an 8 byte boundary offset in the store buffer to improve store bandwidth. Similarly non-side-effect-noncacheable stores may be coalesced with adjacent non-side-effect noncacheable stores within an 8-byte boundary offset in the store buffer.

In order to maintain strong ordering for I/O accesses, stores with side-effect attribute (**e** bit set) will not be combined with any other stores.

Stores that are separated by an intervening MEMBAR #Sync will not be coalesced.

Address Space Identifiers (ASIs)

This appendix describes address space identifiers (ASIs) in the following sections:

- **Address Space Identifiers and Address Spaces** on page 423.
- **ASI Values** on page 423.
- **ASI Assignments** on page 424.
- **Special Memory Access ASIs** on page 432.

10.1 Address Space Identifiers and Address Spaces

An Oracle SPARC Architecture processor provides an address space identifier (ASI) with every address sent to memory. The ASI does the following:

- Distinguishes between different address spaces
- Provides an attribute that is unique to an address space
- Maps internal control and diagnostics registers within a virtual processor

The memory management unit uses a 64-bit virtual address and an 8-bit ASI to generate a memory, I/O, or internal register address.

10.2 ASI Values

The range of address space identifiers (ASIs) is 00_{16}-FF_{16} . That range is divided into restricted and unrestricted portions. ASIs in the range 80_{16}-FF_{16} are unrestricted; they may be accessed by software running in any privilege mode.

ASIs in the range 00_{16}-7F_{16} are restricted; they may only be accessed by software running in a mode with sufficient privilege for the particular ASI. ASIs in the range 00_{16}-2F_{16} may only be accessed by software running in privileged or hyperprivileged mode and ASIs in the range 30_{16}-7F_{16} may only be accessed by software running in hyperprivileged mode.

SPARC V9 Compatibility Note | In SPARC V9, the range of ASIs was evenly divided into restricted (00_{16}-7F_{16}) and unrestricted (80_{16}-FF_{16}) halves.

An attempt by nonprivileged software to access a restricted (privileged or hyperprivileged) ASI (00_{16}-7F_{16}) causes a *privileged_action* trap.

An attempt by privileged software to access a hyperprivileged ASI (30_{16}-7F_{16}) also causes a *privileged_action* trap.

An ASI can be categorized based on how it affects the MMU's treatment of the accompanying address, into one of three categories:

- A *Translating* ASI (the most common type) causes the accompanying address to be treated as a virtual address (which is translated by the MMU).
- A *Non-translating* ASI is not translated by the MMU; instead the address is passed through unchanged. Nontranslating ASIs are typically used for accessing internal registers.
- A *Real* ASI causes the accompanying address to be treated as a real address. An access using a Real ASI can cause exception(s) only visible in hyperprivileged mode. Real ASIs are typically used by privileged software for directly accessing memory using real (as opposed to virtual) addresses.

Implementation-dependent ASIs may or may not be translated by the MMU. See implementation-specific documentation for detailed information about implementation-dependent ASIs.

10.3 ASI Assignments

Every load or store address in an Oracle SPARC Architecture processor has an 8-bit Address Space Identifier (ASI) appended to the virtual address (VA). The VA plus the ASI fully specify the address.

For instruction fetches and for data loads, stores, and load-stores that do not use the load or store alternate instructions, the ASI is an implicit ASI generated by the virtual processor.

If a load alternate, store alternate, or load-store alternate instruction is used, the value of the ASI (an "explicit ASI") can be specified in the ASI register or as an immediate value in the instruction.

In practice, ASIs are not only used to differentiate address spaces but are also used for other functions, like referencing registers in the MMU unit.

10.3.1 Supported ASIs

TABLE 10-1 lists architecturally-defined ASIs; some are in all Oracle SPARC Architecture implementations and some are only present in some implementations.

An ASI marked with a closed bullet (●) is required to be implemented on all Oracle SPARC Architecture 2015 processors.

An ASI marked with an open bullet (○) is defined by the Oracle SPARC Architecture 2015 but is not necessarily implemented in all Oracle SPARC Architecture 2015 processors; its implementation is optional. Across all implementations on which it is implemented, it appears to software to behave identically.

Some ASIs may only be used with certain load or store instructions; see table footnotes for details.

The word "decoded" in the Virtual Address column of TABLE 10-1 indicates that the the supplied virtual address is decoded by the virtual processor.

The "V / non-T / R" column of the table indicates whether each ASI is a Translating ASI(translates from Virtual), non-Translating ASI, or Real ASI, respectively.

ASIs marked "Reserved" are set aside for use in future revisions to the architecture and are not to be used by implemenations. ASIs marked "implementation dependent" may be used for implementation-specific purposes.

Attempting to access an address space described as “Implementation dependent” in TABLE 10-1 produces implementation-dependent results.

TABLE 10-1 Oracle SPARC Architecture ASIs (1 of 7)

ASI Value	req'd(●) opt'l(○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
00 ₁₆ – 01 ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
02 ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
03 ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
04 ₁₆	●	ASI_NUCLEUS (ASI_N)	RW ^{2,4}	(decoded)	V	—	Implicit address space, nucleus context, TL > 0
05 ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
06 ₁₆ – 0B ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
0C ₁₆	●	ASI_NUCLEUS_LITTLE (ASI_NL)	RW ^{2,4}	(decoded)	V	—	Implicit address space, nucleus context, TL > 0, little-endian
0D ₁₆ – 0F ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
10 ₁₆	●	ASI_AS_IF_USER_PRIMARY (ASI_AIUP)	RW ^{2,4,18}	(decoded)	V	—	Primary address space, as if user (nonprivileged)
11 ₁₆	●	ASI_AS_IF_USER_SECONDARY (ASI_AIUS)	RW ^{2,4,18}	(decoded)	V	—	Secondary address space, as if user (nonprivileged)
12 ₁₆	●	ASI_MONITOR_AS_IF_USER_PRIMARY (ASI_MAIUP)	R ^{2,11,18}	(decoded)	V	—	ILoad-Monitor version of ASI_AIUP
13 ₁₆	●	ASI_MONITOR_AS_IF_USER_SECONDARY (ASI_MAIUS)	R ^{2,11,18}	(decoded)	V	—	Load-Monitor version of ASI_AIUS
14 ₁₆	○	ASI_REAL	RW ^{2,4}	(decoded)	R	—	Real address
15 ₁₆	○	ASI_REAL_IO ^D	RW ^{2,5}	(decoded)	R	—	Real address, noncacheable, with side effect (deprecated)
16 ₁₆	○	ASI_BLOCK_AS_IF_USER_PRIMARY (ASI_BLK_AIUP)	RW ^{2,8,14,18}	(decoded)	V	—	Primary address space, block load/store, as if user (nonprivileged)
17 ₁₆	○	ASI_BLOCK_AS_IF_USER_SECONDARY (ASI_BLK_AIUS)	RW ^{2,8,14,18}	(decoded)	V	—	Secondary address space, block load/store, as if user (nonprivileged)
18 ₁₆	●	ASI_AS_IF_USER_PRIMARY_LITTLE (ASI_AIUPL)	RW ^{2,4,18}	(decoded)	V	—	Primary address space, as if user (nonprivileged), little-endian
19 ₁₆	●	ASI_AS_IF_USER_SECONDARY_LITTLE (ASI_AIUSL)	RW ^{2,4,18}	(decoded)	V	—	Secondary address space, as if user (nonprivileged), little-endian
1A ₁₆ – 1B ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
1C ₁₆	○	ASI_REAL_LITTLE (ASI_REAL_L)	RW ^{2,4}	(decoded)	R	—	Real address, little-endian
1D ₁₆	○	ASI_REAL_IO_LITTLE ^D (ASI_REAL_IO_L ^D)	RW ^{2,5}	(decoded)	R	—	Real address, noncacheable, with side effect, little-endian (deprecated)
1E ₁₆	○	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE (ASI_BLK_AIUPL)	RW ^{2,8,14,18}	(decoded)	V	—	Primary address space, block load/store, as if user (nonprivileged), little-endian
1F ₁₆	○	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE (ASI_BLK_AIUSL)	RW ^{2,8,14,18}	(decoded)	V	—	Secondary address space, block load/store, as if user (nonprivileged), little-endian

TABLE 10-1 Oracle SPARC Architecture ASIs (2 of 7)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
20 ₁₆	○	ASI_SCRATCHPAD	RW ^{2,6}	(decoded; see below)	non-T	per strand	Privileged Scratchpad registers; implementation dependent ¹
	○		"	0 ₁₆	"	"	Scratchpad Register 0 ¹
	○		"	8 ₁₆	"	"	Scratchpad Register 1 ¹
	○		"	10 ₁₆	"	"	Scratchpad Register 2 ¹
	○		"	18 ₁₆	"	"	Scratchpad Register 3 ¹
	○		"	20 ₁₆	"	"	Scratchpad Register 4 ¹
	○		"	28 ₁₆	"	"	Scratchpad Register 5 ¹
	○		"	30 ₁₆	"	"	Scratchpad Register 6 ¹
	○		"	38 ₁₆	"	"	Scratchpad Register 7 ¹
21 ₁₆	○	ASI_MMU, ASI_MMU_CONTEXTID	RW ^{2,6}	(decoded; see below)	non-T	per strand	MMU context registers
	○		"	8 ₁₆	"	"	0Read I/D MMU Primary Context ID register 0 Write I/D MMU Primary Context ID registers 0 and 1 (see section 14.6.2)
	○		"	10 ₁₆	"	"	0Read I/D MMU Secondary Context ID register Write I/D MMU Secondary Context ID registers 0 and 1 (see section 14.6.2)
	○		"	28 ₁₆	"	"	Read/write I/D MMU Primary Context ID register 0 (Note: unlike ASI 21 ₁₆ , VA 8 ₁₆ , does not affect I/D MMU Primary Context Register 1; see section 14.6.2)
	○		"	30 ₁₆	"	"	Read/write I/D MMU Secondary Context ID register 0 (Note: unlike ASI 21 ₁₆ , VA 10 ₁₆ , does not affect I/D MMU Secondary Context Register 1; see section 14.6.2)
	○		"	108 ₁₆	"	"	I/D Primary Context ID register 1
	○		"	110 ₁₆	"	"	I/D MMU Secondary Context ID register 1
22 ₁₆	○	ASI_TWIX_AS_IF_USER_PRIMARY (ASI_TWIX_AIUP)	R ^{2,7,11}	(decoded)	V	—	Primary address space, 128-bit atomic load twin extended word, as if user (nonprivileged)
	○	ASI_ST_BLKINIT_AS_IF_USER_PRIMARY (ASI_STBI_AIUP)	W ^{2,20}	(decoded)	V	—	Primary address space, block-initializing store, as if user (nonprivileged); may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
23 ₁₆	○	ASI_TWIX_AS_IF_USER_SECONDARY (ASI_TWIX_AIUS)	R ^{2,7,11}	(decoded)	V	—	Secondary address space, 128-bit atomic load twin extended word, as if user (nonprivileged); may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
	○	ASI_ST_BLKINIT_AS_IF_USER_SECONDARY (ASI_STBI_AIUS)	W ^{2,20}	(decoded)	V	—	Secondary address space, block-initializing store, as if user (nonprivileged)

TABLE 10-1 Oracle SPARC Architecture ASIs (3 of 7)

ASI Value	req'd (●) opt'1 (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
24 ₁₆	○	—	—	—	—	—	Implementation dependent ¹
25 ₁₆	○	ASI_QUEUE	(see below)	(decoded; non-T see below)		per strand	
	○		RW ^{2,6}	3C0 ₁₆	"	"	CPU Mondo Queue Head Pointer
	○		RW ^{2,6,17}	3C8 ₁₆	"	"	CPU Mondo Queue Tail Pointer
	○		RW ^{2,6}	3D0 ₁₆	"	"	Device Mondo Queue Head Pointer
	○		RW ^{2,6,17}	3D8 ₁₆	"	"	Device Mondo Queue Tail Pointer
	○		RW ^{2,6}	3E0 ₁₆	"	"	Resumable Error Queue Head Pointer
	○		RW ^{2,6,17}	3E8 ₁₆	"	"	Resumable Error Queue Tail Pointer
	○		RW ^{2,6}	3F0 ₁₆	"	"	Nonresumable Error Queue Head Pointer
	○		RW ^{2,6,17}	3F8 ₁₆	"	"	Nonresumable Error Queue Tail Pointer
26 ₁₆	○	ASI_TWIX_REAL (ASI_TWIX_R) ASI_QUAD_LDD_REAL ^{pt}	R ^{2,7,11}	(decoded)	R	—	128-bit atomic twin extended-word load from real address
	○	ASI_ST_BLKINIT_REAL (ASI_STBI_R)	W ^{2,20}	(decoded)	V	—	Block-initializing store to real address; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
27 ₁₆	○	ASI_TWIX_NUCLEUS (ASI_TWIX_N)	R ^{2,7,11}	(decoded)	V	—	Nucleus context, 128-bit atomic load twin extended-word
	○	ASI_ST_BLKINIT_NUCLEUS (ASI_STBI_N)	W ^{2,20}	(decoded)	V	—	Nucleus context, block-initializing store; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
28 ₁₆ – 29 ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹
2A ₁₆	○	ASI_TWIX_AS_IF_USER_PRIMARY_LITTLE (ASI_TWIX_AIUP_L)	R ^{2,7,11}	(decoded)	V	—	Primary address space, 128-bit atomic load twin extended-word, as if user (nonprivileged), little-endian
	○	ASI_ST_BLKINIT_AS_IF_USER_PRIMARY_LITTLE (ASI_STBI_AIUPL)	W ^{2,20}	(decoded)	V	—	Primary address space, block-initializing store, as if user (nonprivileged), little-endian; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
2B ₁₆	○	ASI_TWIX_AS_IF_USER_SECONDARY_LITTLE (ASI_TWIX_AIUS_L)	R ^{2,7,11}	(decoded)	V	—	Secondary address space, 128-bit atomic load twin extended-word, as if user (nonprivileged), little-endian
	○	ASI_ST_BLKINIT_AS_IF_USER_SECONDARY_LITTLE (ASI_STBI_AIUS_L)	W ^{2,20}	(decoded)	V	—	Secondary address space, block-initializing store, as if user (nonprivileged), little-endian; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
2C ₁₆	○	—	— ²	—	—	—	Implementation dependent ¹
2D ₁₆	○	—	— ^{2,12}	—	—	—	Implementation dependent ¹

TABLE 10-1 Oracle SPARC Architecture ASIs (4 of 7)

ASI Value	req'd(●) opt'l(○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
2E ₁₆	○	ASI_TWINX_REAL_LITTLE (ASI_TWINX_REAL_L) ASI_QUAD_LDD_REAL_LITTLE ^{D†}	R ^{2,7,11}	(decoded)	R	—	128-bit atomic twin-extended-word load from real address, little-endian
	○	ASI_ST_BLKINIT_REAL_LITTLE (ASI_STBI_RL)	W ^{2,20}	(decoded)	V	—	Block-initializing store to real address, little-endian; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
2F ₁₆	○	ASI_TWINX_NUCLEUS_LITTLE (ASI_TWINX_NL)	R ^{2,7,11}	(decoded)	V	—	Nucleus context, 128-bit atomic load twin extended-word, little-endian
	○	ASI_ST_BLKINIT_NUCLEUS_LITTLE (ASI_STBI_NL)	W ^{2,20}	(decoded)	V	—	Nucleus context, block-initializing store, little-endian; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
30 ₁₆ – 7F ₁₆	●	—	— ³	—	—	—	Reserved for use in hyperprivilege mode
34 ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
42 ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
43 ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
45 ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
46 ₁₆ – 48 ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
49 ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
4A ₁₆ – 4B ₁₆	○	—	— ^{3,13}	—	—	—	Implementation dependent ¹
4C ₁₆	○	Error Status and Enable Registers	—	—	—	—	Implementation dependent ¹
75 ₁₆ – 7F ₁₆	○	—	— ^{3,13}	—	—	—	Reserved
80 ₁₆	●	ASI_PRIMARY (ASI_P)	RW ⁴	(decoded)	V	—	Implicit primary address space
81 ₁₆	●	ASI_SECONDARY (ASI_S)	RW ⁴	(decoded)	V	—	Secondary address space
82 ₁₆	●	ASI_PRIMARY_NO_FAULT (ASI_PNF)	R ^{9,11}	(decoded)	V	—	Primary address space, no fault
83 ₁₆	●	ASI_SECONDARY_NO_FAULT (ASI_SNF)	R ^{9,11}	(decoded)	V	—	Secondary address space, no fault
84 ₁₆	●	ASI_MONITOR_PRIMARY (ASI_MP)	R ¹¹	(decoded)	V	—	Load-Monitor version of ASI_P
85 ₁₆	●	ASI_MONITOR_SECONDARY (ASI_MS)	R ¹¹	(decoded)	V	—	Load-Monitor version of ASI_S
86 ₁₆ – 87 ₁₆	○	—	— ¹⁶	—	—	—	Reserved
88 ₁₆	●	ASI_PRIMARY_LITTLE (ASI_PL)	RW ⁴	(decoded)	V	—	Implicit primary address space, little-endian
89 ₁₆	●	ASI_SECONDARY_LITTLE (ASI_SL)	RW ⁴	(decoded)	V	—	Secondary address space, little-endian
8A ₁₆	●	ASI_PRIMARY_NO_FAULT_LITTLE (ASI_PNFL)	R ^{9,11}	(decoded)	V	—	Primary address space, no fault, little-endian
8B ₁₆	●	ASI_SECONDARY_NO_FAULT_LITTLE (ASI_SNFL)	R ^{9,11}	(decoded)	V	—	Secondary address space, no fault, little-endian
8C ₁₆ – 8F ₁₆	○	—	— ¹⁶	—	—	—	Reserved
90 ₁₆ – 9C ₁₆	○	—	— ¹⁶	—	—	—	Reserved

TABLE 10-1 Oracle SPARC Architecture ASIs (5 of 7)

ASI Value	req'd (●) opt'1 (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
93 ₁₆ – 9C ₁₆	○	—	— ₁₆	—	—	—	Reserved
9D ₁₆ – 9F ₁₆	○	—	— ₁₆	—	—	—	Reserved
A0 ₁₆ – AF ₁₆	○	—	— ₁₆	—	—	—	Reserved
B0 ₁₆	○	ASI_PIC	RW	0 ₁₆	non-T	per strand	Performance Instrumentation Counter 0
	○		RW	8 ₁₆	non-T	per strand	Performance Instrumentation Counter 1
	○		RW	10 ₁₆	non-T	per strand	Performance Instrumentation Counter 2
	○		RW	18 ₁₆	non-T	per strand	Performance Instrumentation Counter 3
B1	○	—	— ₁₆	—	—	—	Reserved
B2 ₁₆ – BF ₁₆	○	—	— ₁₆	—	—	—	Reserved
C0 ₁₆	○	ASI_PST8_PRIMARY (ASI_PST8_P)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 8x8-bit partial store
C1 ₁₆	○	ASI_PST8_SECONDARY (ASI_PST8_S)	W ^{8,10,14}	(decoded)	v	—	Secondary address space, 8x8-bit partial store
C2 ₁₆	○	ASI_PST16_PRIMARY (ASI_PST16_P)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 4x16-bit partial store
C3 ₁₆	○	ASI_PST16_SECONDARY (ASI_PST16_S)	W ^{8,10,14}	(decoded)	v	—	Secondary address space, 4x16-bit partial store
C4 ₁₆	○	ASI_PST32_PRIMARY (ASI_PST32_P)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 2x32-bit partial store
C5 ₁₆	○	ASI_PST32_SECONDARY (ASI_PST32_S)	W ^{8,10,14}	(decoded)	v	—	Secondary address space, 2x32-bit partial store
C6 ₁₆ – C7 ₁₆	○	—	— ₁₅	—	—	—	Implementation dependent ¹
C8 ₁₆	○	ASI_PST8_PRIMARY_LITTLE (ASI_PST8_PL)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 8x8-bit partial store, little-endian
C9 ₁₆	○	ASI_PST8_SECONDARY_LITTLE (ASI_PST8_SL)	W ^{8,10,14}	(decoded)	v	—	Secondary address space, 8x8-bit partial store, little-endian
CA ₁₆	○	ASI_PST16_PRIMARY_LITTLE (ASI_PST16_PL)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 4x16-bit partial store, little-endian
CB ₁₆	○	ASI_PST16_SECONDARY_LITTLE (ASI_PST16_SL)	W ^{8,10,14}	(decoded)	v	—	Secondary address space, 4x16-bit partial store, little-endian
CC ₁₆	○	ASI_PST32_PRIMARY_LITTLE (ASI_PST32_PL)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 2x32-bit partial store, little-endian
CD ₁₆	○	ASI_PST32_SECONDARY_LITTLE (ASI_PST32_SL)	W ^{8,10,14}	(decoded)	v	—	Second address space, 2x32-bit partial store, little-endian
CE ₁₆ – CF ₁₆	○	—	— ₁₅	—	—	—	Implementation dependent ¹
D0 ₁₆	○	ASI_FL8_PRIMARY (ASI_FL8_P)	RW ^{8,14}	(decoded)	v	—	Primary address space, one 8-bit floating-point load/store
D1 ₁₆	○	ASI_FL8_SECONDARY (ASI_FL8_S)	RW ^{8,14}	(decoded)	v	—	Second address space, one 8-bit floating-point load/store
D2 ₁₆	○	ASI_FL16_PRIMARY (ASI_FL16_P)	RW ^{8,14}	(decoded)	v	—	Primary address space, one 16-bit floating-point load/store

TABLE 10-1 Oracle SPARC Architecture ASIs (6 of 7)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
D3 ₁₆	○	ASI_FL16_SECONDARY (ASI_FL16_S)	RW ^{8,14}	(decoded)	v	—	Second address space, one 16-bit floating-point load/store
D4 ₁₆ – D7 ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
D8 ₁₆	○	ASI_FL8_PRIMARY_LITTLE (ASI_FL8_PL)	RW ^{8,14}	(decoded)	v	—	Primary address space, one 8-bit floating point load/store, little-endian
D9 ₁₆	○	ASI_FL8_SECONDARY_LITTLE (ASI_FL8_SL)	RW ^{8,14}	(decoded)	v	—	Second address space, one 8-bit floating point load/store, little-endian
DA ₁₆	○	ASI_FL16_PRIMARY_LITTLE (ASI_FL16_PL)	RW ^{8,14}	(decoded)	v	—	Primary address space, one 16-bit floating-point load/store, little-endian
DB ₁₆	○	ASI_FL16_SECONDARY_LITTLE (ASI_FL16_SL)	RW ^{8,14}	(decoded)	v	—	Second address space, one 16-bit floating point load/store, little-endian
DC ₁₆ – DF ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
E0 ₁₆	○	ASI_BLOCK_COMMIT_PRIMARY (ASI_BLK_COMMIT_P)	W ^{8,10,14}	(decoded)	v	—	Primary address space, 8x8-byte block store commit operation
E1 ₁₆	○	ASI_BLOCK_COMMIT_SECONDARY (ASI_BLK_COMMIT_S)	W ^{8,10,14}	(decoded)	v	—	Secondary address space, 8x8-byte block store commit operation
E2 ₁₆	○	ASI_TWIXX_PRIMARY (ASI_TWIXX_P)	R ¹⁹	(decoded)	v	—	Primary address space, 128-bit atomic load twin extended word
	○	ASI_ST_BLKINIT_PRIMARY (ASI_STBI_P)	W ²⁰	(decoded)	v	—	Primary address space, block-initializing store; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
E3 ₁₆	○	ASI_TWIXX_SECONDARY (ASI_TWIXX_S)	R ¹⁹	(decoded)	v	—	Secondary address space, 128-bit atomic load twin extended-word
	○	ASI_ST_BLKINIT_SECONDARY (ASI_STBI_S)	W ²⁰	(decoded)	v	—	Secondary address space, block-initializing store; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
E4 ₁₆ – E6 ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
E7 ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
E8 ₁₆ – E9 ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
EA ₁₆	○	ASI_TWIXX_PRIMARY_LITTLE (ASI_TWIXX_PL)	R ¹⁹	(decoded)	v	—	Primary address space, 128-bit atomic load twin extended word, little endian
	○	ASI_ST_BLKINIT_PRIMARY_LITTLE (ASI_STBI_PL)	W ²⁰	(decoded)	v	—	Primary address space, block-initializing store, little endian; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
EB ₁₆	○	ASI_TWIXX_SECONDARY_LITTLE (ASI_TWIXX_SL)	R ¹⁹	(decoded)	v	—	Secondary address space, 128-bit atomic load twin extended word, little endian
	○	ASI_ST_BLKINIT_SECONDARY_LITTLE (ASI_STBI_SL)	W ²⁰	(decoded)	v	—	Secondary address space, block-initializing store, little endian; may be evicted from cache quickly (installed in cache with LRU [least-recently-used] or equivalent attribute)
EC ₁₆ – EF ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
F0 ₁₆	○	ASI_BLOCK_PRIMARY (ASI_BLK_P)	RW ^{8,14}	(decoded)	v	—	Primary address space, 8x8-byte block load/store

TABLE 10-1 Oracle SPARC Architecture ASIs (7 of 7)

ASI Value	req'd (●) opt'1 (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	V/ non-T/ R	Shared /per strand	Description
F1 ₁₆	○	ASI_BLOCK_SECONDARY (ASI_BLK_S)	RW ^{8,14}	(decoded)	V	—	Secondary address space, 8x8- byte block load/store
F2 ₁₆	○	ASI_ST_BLKINIT_PRIMARY_MRU (ASI_STBI_PM)	W ²⁰	(decoded)	V	—	Primary address space, block-initializing store (installed in cache with MRU [most-recently-used] or equivalent attribute)
F3 ₁₆	○	ASI_ST_BLKINIT_SECONDARY_MRU (ASI_STBI_SM)	W ²⁰	(decoded)	V	—	Secondary address space, block-initializing store (installed in cache with MRU [most-recently-used] or equivalent attribute)
F4 ₁₆ – F5 ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
F6 ₁₆ – F7 ₁₆	○	—	—	—	—	—	Implementation dependent ¹
F8 ₁₆	○	ASI_BLOCK_PRIMARY_LITTLE (ASI_BLK_PL)	RW ^{8,14}	(decoded)	V	—	Primary address space, 8x8- byte block load/store, little endian
F9 ₁₆	○	ASI_BLOCK_SECONDARY_LITTLE (ASI_BLK_SL)	RW ^{8,14}	(decoded)	V	—	Secondary address space, 8x8- byte block load/store, little endian
FA ₁₆	○	ASI_ST_BLKINIT_PRIMARY_LITTLE_MR U (ASI_STBI_PLM)	W ²⁰	(decoded)	V	—	Primary address space, block-initializing store, little endian (installed in cache with MRU [most-recently-used] or equivalent attribute)
FB ₁₆	○	ASI_ST_BLKINIT_SECONDARY_LITTLE_MRU (ASI_STBI_SLM)	W ²⁰	(decoded)	V	—	Secondary address space, block-initializing store, little endian (installed in cache with MRU [most-recently-used] or equivalent attribute)
FC ₁₆ – FD ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹
FE ₁₆ – FF ₁₆	○	—	— ¹⁵	—	—	—	Implementation dependent ¹

† This ASI name has been changed, for consistency; although use of this name is deprecated and software should use the new name, the old name is listed here for compatibility.

- 1 Implementation dependent ASI (impl. dep. #29); available for use by implementors. Software that references this ASI may not be portable.
- 2 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in nonprivileged mode causes a *privileged_action* exception.
- 3 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in nonprivileged mode or privileged mode causes a *privileged_action* exception.
- 4 May be used with all load alternate, store alternate, atomic alternate and prefetch alternate instructions (CASA, CASXA, LDSTUBA, LDTWA, LDDFA, LDFA, LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, LDUWA, LDXA, PREFETCHA, STBA, STTWA, STDFA, STFA, STHA, STWA, STXA, SWAPA).
- 5 May be used with all of the following load alternate and store alternate instructions: LDTWA, LDDFA, LDFA, LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, LDUWA, LDXA, STBA, STTWA, STDFA, STFA, STHA, STWA, STXA.
Use with an atomic alternate or prefetch alternate instruction (CASA, CASXA, LDSTUBA, SWAPA or PREFETCHA) causes a *DAE_invalid_asi* exception.
- 6 May only be used in a LDXA or STXA instruction for RW ASIs, LDXA for read-only ASIs and STXA for write-only ASIs. Use of LDXA for write-only ASIs, STXA for read-only ASIs, or any other load alternate, store alternate, atomic alternate or prefetch alternate instruction causes a *DAE_invalid_asi* exception.

- 7 May only be used in an LDTXA instruction. Use of this ASI in any other load alternate, store alternate, atomic alternate or prefetch alternate instruction causes a *DAE_invalid_asi* exception.
- 8 May only be used in a LDDFA or STDFA instruction for RW ASIs, LDDFA for read-only ASIs and STDFA for write-only ASIs. Use of LDDFA for write-only ASIs, STDFA for read-only ASIs, or any other load alternate, store alternate, atomic alternate or prefetch alternate instruction causes a *DAE_invalid_asi* exception.
- 9 May be used with all of the following load and prefetch alternate instructions: LDTWA, LDDFA, LDFA, LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, LDUWA, LDXA, PREFETCHA. Use with an atomic alternate or store alternate instruction causes a *DAE_invalid_asi* exception.
- 10 Write(store)-only ASI; an attempted load alternate, atomic alternate, or prefetch alternate instruction to this ASI causes a *DAE_invalid_asi* exception.
- 11 Read(load)-only ASI; an attempted store alternate or atomic alternate instruction to this ASI causes a *DAE_invalid_asi* exception.
- 12 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in privileged mode causes a *DAE_invalid_asi* exception.
- 14 An attempted access to this ASI may cause an exception (see *Special Memory Access ASIs* on page 432 for details).
- 15 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in any mode causes a *DAE_invalid_asi* exception if this ASI is not implemented by the model dependent implementation.
- 16 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to a reserved ASI in any mode causes a *DAE_invalid_asi* exception.
- 17 The Queue Tail Registers (ASI 25₁₆) are read-only. An attempted write to the Queue Tail Registers causes a *DAE_invalid_asi* exception
- 19 May only be used in an LDTXA (load twin-extended-word) instruction (which shares an opcode with LDTWA). Use of this ASI in any other load instruction causes a *DAE_invalid_asi* exception.

10.4 Special Memory Access ASIs

This section describes special memory access ASIs that are not described in other sections.

10.4.1 ASIs 10₁₆, 11₁₆, 16₁₆, and 17₁₆ (ASI_*AS_IF_USER_*)

These ASI are intended to be used in accesses from privileged mode, but are processed as if they were issued from nonprivileged mode. Therefore, they are subject to privilege-related exceptions. They are distinguished from each other by the context from which the access is made, as described in TABLE 10-2.

When one of these ASIs is specified in a load alternate or store alternate instruction, the virtual processor behaves as follows:

- In nonprivileged mode, a *privileged_action* exception occurs
- In any other privilege mode:
 - If U/DMMU TTE.p = 1, a *DAE_privilege_violation* exception occurs
 - Otherwise, the access occurs and its endianness is determined by the U/DMMU TTE.ie bit. If U/DMMU TTE.ie = 0, the access is big-endian; otherwise, it is little-endian.

TABLE 10-2 Privileged ASI_*AS_IF_USER_* ASIs

ASI	Names	Addressing (Context)	Endianness of Access
10 ₁₆	ASI_AS_IF_USER_PRIMARY (ASI_AIUP)	Virtual (Primary)	Big-endian when U/DMMU TTE.ie = 0; little-endian when U/DMMU TTE.ie = 1
11 ₁₆	ASI_AS_IF_USER_SECONDARY (ASI_AIUS)	Virtual (Secondary)	
16 ₁₆	ASI_BLOCK_AS_IF_USER_PRIMARY (ASI_BLK_AIUP)	Virtual (Primary)	
17 ₁₆	ASI_BLOCK_AS_IF_USER_SECONDARY (ASI_BLK_AIUS)	Virtual (Secondary)	

10.4.2 ASIs 18₁₆, 19₁₆, 1E₁₆, and 1F₁₆ (ASI_*AS_IF_USER_*_LITTLE)

These ASIs are little-endian versions of ASIs 10₁₆, 11₁₆, 16₁₆, and 17₁₆ (ASI_AS_IF_USER_*), described in section 10.4.1. Each operates identically to the corresponding non-little-endian ASI, except that if an access occurs its endianness is the opposite of that for the corresponding non-little-endian ASI.

These ASI are intended to be used in accesses from privileged mode, but are processed as if they were issued from nonprivileged mode. Therefore, they are subject to privilege-related exceptions. They are distinguished from each other by the context from which the access is made, as described in TABLE 10-3.

When one of these ASIs is specified in a load alternate or store alternate instruction, the virtual processor behaves as follows:

- In nonprivileged mode, a *privileged_action* exception occurs
- In any other privilege mode:
 - If U/DMMU TTE.p = 1, a *DAE_privilege_violation* exception occurs
 - Otherwise, the access occurs and its endianness is determined by the U/DMMU TTE.ie bit. If U/DMMU TTE.ie = 0, the access is little-endian; otherwise, it is big-endian.

TABLE 10-3 Privileged ASI_*AS_IF_USER_*_LITTLE ASIs

ASI	Names	Addressing (Context)	Endianness of Access
18 ₁₆	ASI_AS_IF_USER_PRIMARY_LITTLE (ASI_AIUPL)	Virtual (Primary)	Little-endian when U/DMMU TTE.ie = 0; big-endian when U/DMMU TTE.ie = 1
19 ₁₆	ASI_AS_IF_USER_SECONDARY_LITTLE (ASI_AIUSL)	Virtual (Secondary)	
1E ₁₆	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE (ASI_BLK_AIUP)	Virtual (Primary)	
1F ₁₆	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE (ASI_BLK_AIUSL)	Virtual (Secondary)	

10.4.3 ASI 14₁₆ (ASI_REAL)

When ASI_REAL is specified in any load alternate, store alternate or prefetch alternate instruction, the virtual processor behaves as follows:

- In nonprivileged mode, a *privileged_action* exception occurs
- In any other privilege mode:

- VA is passed through to RA
- During the address translation, context values are disregarded.
- The endianness of the access is determined by the U/DMMU TTE.ie bit; if U/DMMU TTE.ie = 0, the access is big-endian, otherwise it is little-endian.

Even if data address translation is disabled, an access with this ASI is still a cacheable access.

10.4.4 ASI 15₁₆ (ASI_REAL_IO)

Accesses with ASI_REAL_IO bypass the external cache and behave as if the side effect bit (TTE.e bit) is set. When this ASI is specified in any load alternate or store alternate instruction, the virtual processor behaves as follows:

- In nonprivileged mode, a *privileged_action* exception occurs
- Used with any other load alternate or store alternate instruction, in privileged mode:
 - VA is passed through to RA
 - During the address translation, context values are disregarded.
 - The endianness of the access is determined by the U/DMMU TTE.ie bit; if U/DMMU TTE.ie = 0, the access is big-endian, otherwise it is little-endian.

10.4.5 ASI 1C₁₆ (ASI_REAL_LITTLE)

ASI_REAL_LITTLE is a little-endian version of ASI 14₁₆ (ASI_REAL). It operates identically to ASI_REAL, except if an access occurs, its endianness the opposite of that for ASI_REAL.

10.4.6 ASI 1D₁₆ (ASI_REAL_IO_LITTLE)

ASI_REAL_IO_LITTLE is a little-endian version of ASI 15₁₆ (ASI_REAL_IO). It operates identically to ASI_REAL_IO, except if an access occurs, its endianness the opposite of that for ASI_REAL_IO.

10.4.7 ASIs 22₁₆, 23₁₆, 26₁₆, 27₁₆, 2A₁₆, 2B₁₆, 2E₁₆, 2F₁₆ (Privileged Load Integer Twin Extended Word)

ASIs 22₁₆, 23₁₆, 27₁₆, 2A₁₆, 2B₁₆ and 2F₁₆ exist for use with the (nonportable) LDTXA instruction as atomic Load Integer Twin Extended Word operations (see *Load Integer Twin Extended Word from Alternate Space* on page 262). These ASIs are distinguished by the context from which the access is made and the endianness of the access, as described in TABLE 10-4.

TABLE 10-4 Privileged Load Integer Twin Extended WordASIs

ASI	Names	Addressing (Context)	Endianness of Access
22 ₁₆	ASI_TWIXN_AS_IF_USER_PRIMARY (ASI_TWIXN_AIUP)	Virtual (Primary)	Big-endian when U/ DMMU
23 ₁₆	ASI_TWIXN_AS_IF_USER_SECONDARY (ASI_TWIXN_AIUS)	Virtual (Secondary)	TTE.ie = 0;
26 ₆	ASI_TWIXN_REAL (ASI_TWIXN_R)	Virtual (Primary)	little-endian when U/ DMMU
27 ₁₆	ASI_TWIXN_NUCLEUS (ASI_TWIXN_N)	Virtual (Nucleus)	TTE.ie = 1
2A ₁₆	ASI_TWIXN_AS_IF_USER_PRIMARY_LITTLE (ASI_TWIXN_AIUP_L)	Virtual (Primary)	Little-endian when U/ DMMU
2B ₁₆	ASI_TWIXN_AS_IF_USER_SECONDARY_LITTLE (ASI_TWIXN_AIUS_L)	Virtual (Secondary)	TTE.ie = 0;
2E ₁₆	ASI_TWIXN_REAL_LITTLE (ASI_TWIXN_RL)	Virtual (Primary)	big-endian when U/ DMMU
2F ₁₆	ASI_TWIXN_NUCLEUS_LITTLE (ASI_TWIXN_NL)	Virtual (Nucleus)	TTE.ie = 1

When these ASIs are used with LDTXA, a *mem_address_not_aligned* exception is generated if the operand address is not 16-byte aligned.

If these ASIs are used with any other Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *DAE_invalid_asi* exception is always generated and *mem_address_not_aligned* is not generated.

Compatibility Note | These ASIs replaced ASIs 24₁₆ and 2C₁₆ used in earlier UltraSPARC implementations; see the detailed Compatibility Note on page 553 for details.

10.4.8 ASIs 26₁₆ and 2E₁₆ (Privileged Load Integer Twin Extended Word, Real Addressing)

ASIs 26₁₆ and 2E₁₆ exist for use with the LDTXA instruction as atomic Load Integer Twin Extended Word operations using Real addressing (see *Load Integer Twin Extended Word from Alternate Space* on page 262). These two ASIs are distinguished by the endianness of the access, as described in TABLE 10-5.

TABLE 10-5 Load Integer Twin Extended Word (Real) ASIs

ASI	Name	Addressing (Context)	Endianness of Access
26 ₁₆	ASI_TWIXN_REAL (ASI_TWIXN_R)	Real (—)	Big-endian when U/DMMU TTE.ie = 0; little-endian when U/ DMMU TTE.ie = 1
2E ₁₆	ASI_TWIXN_REAL_LITTLE (ASI_TWIXN_REAL_L)	Real (—)	Little-endian when U/DMMU TTE.ie = 0; big-endian when U/ DMMU TTE.ie = 1

When these ASIs are used with LDTXA, a *mem_address_not_aligned* exception is generated if the operand address is not 16-byte aligned.

If these ASIs are used with any other Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *DAE_invalid_asi* exception is always generated and *mem_address_not_aligned* is not generated.

Compatibility Note | These ASIs replaced ASIs 34₁₆ and 3C₁₆ used in earlier UltraSPARC implementations; see the Compatibility Note on page 553 for details.

10.4.9 ASIs E2₁₆, E3₁₆, EA₁₆, EB₁₆ (Nonprivileged Load Integer Twin Extended Word)

ASIs E2₁₆, E3₁₆, EA₁₆, and EB₁₆ exist for use with the (nonportable) LDTXA instruction as atomic Load Integer Twin Extended Word operations (see *Load Integer Twin Extended Word from Alternate Space* on page 262). These ASIs are distinguished by the address space accessed (Primary or Secondary) and the endianness of the access, as described in TABLE 10-6.

TABLE 10-6 Load Integer Twin Extended Word ASIs

ASI	Names	Addressing (Context)	Endianness of Access
E2 ₁₆	ASI_TWIXX_PRIMARY (ASI_TWIXX_P)	Virtual (Primary)	Big-endian when U/DMMU TTE.ie = 0, little-endian when U/DMMU TTE.ie = 1
E3 ₁₆	ASI_TWIXX_SECONDARY (ASI_TWIXX_S)	Virtual (Secondary)	
EA ₁₆	ASI_TWIXX_PRIMARY_LITTLE (ASI_TWIXX_PL)	Virtual (Primary)	Little-endian when U/DMMU TTE.ie = 0, big-endian when U/DMMU TTE.ie = 1
EB ₁₆	ASI_TWIXX_SECONDARY_LITTLE (ASI_TWIXX_SL)	Virtual (Secondary)	

When these ASIs are used with LDTXA, a *mem_address_not_aligned* exception is generated if the operand address is not 16-byte aligned.

If these ASIs are used with any other Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *DAE_invalid_asi* exception is always generated and *mem_address_not_aligned* is not generated.

10.4.10 Block Load and Store ASIs

ASIs 16₁₆, 17₁₆, 1E₁₆, 1F₁₆, E0₁₆, E1₁₆, F0₁₆, F1₁₆, F8₁₆, and F9₁₆ exist for use with LDDFA and STDFA instructions as Block Load (LDBLOCKF^D) and Block Store (STBLOCKF^D) operations (see *Block Load* on page 243 and *Block Store* on page 337).

When these ASIs are used with the LDDFA (STDFA) opcode for Block Load (Store), a *mem_address_not_aligned* exception is generated if the operand address is not 64-byte aligned.

ASIs E0₁₆ and E1₁₆ are only defined for use in Block Store with Commit operations (see page 337). Neither ASI E0₁₆ nor E1₁₆ should be used with the LDDFA opcode; however, if either *is* used, the resulting behavior is specified in the LDDFA instruction description on page 249.

If a Block Load or Block Store ASI is used with any other Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *DAE_invalid_asi* exception is always generated and *mem_address_not_aligned* is not generated.

10.4.11 Partial Store ASIs

ASIs C0₁₆–C5₁₆ and C8₁₆–CD₁₆ exist for use with the STDFFA instruction as Partial Store (STPARTIALF) operations (see *DAE_invalid_asi*, *eDAE_invalid_asi*, *eStore Partial Floating-Point* on page 347).

When these ASIs are used with STDFFA for Partial Store, a *mem_address_not_aligned* exception is generated if the operand address is not 8-byte aligned and an *illegal_instruction* exception is generated if *i* = 1 in the instruction and the ASI register contains one of the Partial Store ASIs.

If one of these ASIs is used with a Store Alternate instruction other than STDFFA, a Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *DAE_invalid_asi* exception is generated and *mem_address_not_aligned*, *LDDF_mem_address_not_aligned*, and *illegal_instruction* (for *i* = 1) are not generated.

ASIs C0₁₆–C5₁₆ and C8₁₆–CD₁₆ are only defined for use in Partial Store operations (see page 347). None of them should be used with LDDFA; however, if any of those ASIs is used with LDDFA, the resulting behavior is specified in the LDDFA instruction description on page 250.

10.4.12 Short Floating-Point Load and Store ASIs

ASIs D0₁₆–D3₁₆ and D8₁₆–DB₁₆ exist for use with the LDDFA and STDFFA instructions as Short Floating-point Load and Store operations (see *Load Floating-Point Register* on page 246 and *Store Floating-Point* on page 340).

When ASI D2₁₆, D3₁₆, DA₁₆, or DB₁₆ is used with LDDFA (STDFFA) for a 16-bit Short Floating-point Load (Store), a *mem_address_not_aligned* exception is generated if the operand address is not halfword-aligned.

If any of these ASIs are used with any other Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *DAE_invalid_asi* exception is always generated and *mem_address_not_aligned* is not generated.

10.4.13 Load-Monitor ASIs

ASIs 12₁₆ (ASI_MONITOR_AS_IF_USER_PRIMARY), 13₁₆ (ASI_MONITOR_AS_IF_USER_SECONDARY), 84₁₆ (ASI_MONITOR_PRIMARY), and 85₁₆ (ASI_MONITOR_SECONDARY) exist for use with the LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, LDUWA and LDXA instructions.

When ASI 12₁₆, 13₁₆, 84₁₆, or 85₁₆ is used with one of the above instructions, the instruction functions as for ASI_AS_IF_USER_PRIMARY, ASI_AS_IF_USER_SECONDARY, ASI_PRIMARY, or ASI_SECONDARY, respectively, except that the load also behaves as a Load-Monitor for a subsequent MWait instruction (see *MWait* on page 292).

10.5 ASI-Accessible Registers

In this section the scratchpad registers are described.

A list of Oracle SPARC Architecture 2015 ASIs is shown in TABLE 10-1 on page 425.

10.5.1 Privileged Scratchpad Registers (ASI_SCRATCHPAD) (D1)

An Oracle SPARC Architecture virtual processor includes eight Scratchpad registers (64 bits each, read/write accessible) (impl.dep. #302-U4-Cs10). The use of the Scratchpad registers is completely defined by software.

The Scratchpad registers are intended to be used by performance-critical trap handler code.

The addresses of the privileged scratchpad registers are defined in TABLE 10-7.

TABLE 10-7 Scratchpad Registers

Assembly Language ASI Name	ASI #	Virtual Address	Privileged Scratchpad Register #
ASI_SCRATCHPAD	20 ₁₆	00 ₁₆	0
		08 ₁₆	1
		10 ₁₆	2
		18 ₁₆	3
		20 ₁₆	4
		28 ₁₆	5
		30 ₁₆	6
		38 ₁₆	7

IMPL. DEP. #404-S10: The degree to which Scratchpad registers 4–7 are accessible to privileged software is implementation dependent. Each may be

- (1) fully accessible,
- (2) accessible, with access much slower than to scratchpad registers 0–3, or
- (3) inaccessible (cause a *DAE_invalid_asi* exception).

V9 Compatibility Note | Privileged scratchpad registers are an Oracle SPARC Architecture extension to SPARC V9.

10.6 ASI Changes in the Oracle SPARC Architecture

The following Compatibility Note summarize ASI changes in the Oracle SPARC Architecture.

Compatibility Note	ASI#	Previous UltraSPARC	Oracle SPARC Architecture
The names of several ASIs used in earlier UltraSPARC implementations have changed in Oracle SPARC Architecture. Their functions have not changed; just their names have changed.	14 ₁₆	ASI_PHYS_USE_EC	ASI_REAL
	15 ₁₆	ASI_PHYS_BYPASS_EC_WITH_EBIT	ASI_REAL_IO
	1C ₁₆	ASI_PHYS_USE_EC_LITTLE (ASI_PHYS_USE_EC_L)	ASI_REAL_LITTLE
	1D ₁₆	ASI_PHYS_BYPASS_EC_WITH_ EBIT_LITTLE (ASI_PHY_BYPASS_EC_WITH_EBIT_L)	ASI_REAL_IO_LITTLE

Performance Instrumentation

This chapter describes the architecture for performance monitoring hardware on Oracle SPARC Architecture processors. The architecture is based on the design of performance instrumentation counters in previous Oracle SPARC Architecture processors, with an extension for the selective sampling of instructions.

11.1 High-Level Requirements

11.1.1 Usage Scenarios

The performance monitoring hardware on Oracle SPARC Architecture processors addresses the needs of various kinds of users. There are four scenarios envisioned:

- *System-wide performance monitoring.* In this scenario, someone skilled in system performance analysis (e.g., a Systems Engineer) is using analysis tools to evaluate the performance of the entire system. An example of such a tool is `cpustat`. The objective is to obtain performance data relating to the configuration and behavior of the system, e.g., the utilization of the memory system.
- *Self-monitoring of performance by the operating system.* In this scenario the OS is gathering performance data in order to tune the operation of the system. Some examples might be:
 - (a) determining whether the processors in the system should be running in single- or multi-stranded mode.
 - (b) determining the affinity of a process to a processor by examining that process's memory behavior.
- *Performance analysis of an application by a developer.* In this scenario a developer is trying to optimize the performance of a specific application, by altering the source code of the application or the compilation options. The developer needs to know the performance characteristics of the components of the application at a coarse grain, and where these are problematic, to be able to determine fine-grained performance information. Using this information, the developer will alter the source or compilation parameters, re-run the application, and observe the new performance characteristics. This process is repeated until performance is acceptable, or no further improvements can be found.

An example might be that a loop nest is measured to be not performing well. Upon closer inspection, the developer determines that the loop has poor cache behavior, and upon more detailed inspection finds a specific operation which repeatedly misses the cache. Reorganizing the code and/or data may improve the cache behavior.

- *Monitoring of an application's performance, e.g., by a Java Virtual Machine.* In this scenario the application is not executing directly on the hardware, but its execution is being mediated by a piece of system software, which for the purposes of this document is called a Virtual Machine. This may

be a Java VM, or a binary translation system running software compiled for another architecture, or for an earlier version of the Oracle SPARC Architecture. One goal of the VM is to optimize the behavior of the application by monitoring its performance and dynamically reorganizing the execution of the application (e.g., by selective recompilation of the application).

This scenario differs from the previous one principally in the time allowed to gather performance data. Because the data are being gathered during the execution of the program, the measurements must not adversely affect the performance of the application by more than, say, a few percent, and must yield insight into the performance of the application in a relatively short time (otherwise, optimization opportunities are deferred for too long). This implies an observation mechanism which is of very low overhead, so that many observations can be made in a short time.

In contrast, a developer optimizing an application has the luxury of running or re-running the application for a considerable period of time (minutes or even hours) to gather data. However, the developer will also expect a level of precision and detail in the data which would overwhelm a virtual machine, so the accuracy of the data required by a virtual machine need not be as high as that supplied to the developer.

Scenarios 1 and 2 are adequately dealt with by a suitable set of performance counters capable of counting a variety of performance-related events. Counters are ideal for these situations because they provide low-overhead statistics without any intrusion into the behavior of the system or disruption to the code being monitored. However, counters may not adequately address the latter two scenarios, in which detailed and timely information is required at the level of individual instructions. Therefore, Oracle SPARC Architecture processors may also implement an instruction sampling mechanism.

11.1.2 Metrics

There are two classes of data reported by a performance instrumentation mechanism:

- *Architectural performance metrics.* These are metrics related to the observable execution of code at the architectural level (Oracle SPARC Architecture). Examples include:
 - The number of instructions executed
 - The number of floating point instructions executed
 - The number of conditional branch instructions executed
- *Implementation performance metrics.* These describe the behavior of the microprocessor in terms of its implementation, and would not necessarily apply to another implementation of the architecture.

In optimizing the performance of an application or system, attention will first be paid to the first class of metrics, and so these are more important. Only in performance-critical cases would the second class receive attention, since using these metrics requires a fairly extensive understanding of the specific implementation of the Oracle SPARC Architecture.

11.1.3 Accuracy Requirements

Accuracy requirements for performance instrumentation vary depending on the scenario. The requirements are complicated by the possibly speculative nature of Oracle SPARC Architecture processor implementations. For example, an implementation may include in its cache miss statistics the misses induced by speculative executions which were subsequently flushed, or provide two separate statistics, one for the misses induced by flushed instructions and one for misses induced by retired instructions. Although the latter would be desirable, the additional implementation complexity of associating events with specific instructions is significant, and so all events may be counted without distinction. The instruction sampling mechanism may distinguish between instructions that retired and those that were flushed, in which case sampling can be used to obtain statistical estimates of the frequencies of operations induced by mis-speculation.

For critical performance measurements, architectural event counts must be accurate to a high degree (1 part in 10^5). Which counters are considered performance-critical (and therefore accurate to 1 part in 10^5) are specified in implementation-specific documentation.

Implementation event counts must be accurate to 1 part in 10^3 , not including the speculative effects mentioned above. An upper bound on counter skew must be stated in implementation-specific documentation.

Programming Note	Increasing the time between counter reads will mitigate the inaccuracies that could be introduced by counter skew (due to speculative effects).
-------------------------	---

11.2 Performance Counters and Controls

The performance instrumentation hardware provides performance instrumentation counters (PICs). The number and size of performance counters is implementation dependent, but each performance counter register contains at least one 32-bit counter. It is implementation dependent whether the performance counter registers are accessed as ASRs or are accessed through ASIs.

There are one or more performance counter control registers (PCRs) associated with the counter registers. It is implementation dependent whether the PCRs are accessed as ASRs or are accessed through ASIs.

Each counter in a counter register can count one kind of event at a time. The number of the kinds of events that can be counted is implementation dependent. For each performance counter register, the corresponding control register is used to select the event type being counted. A counter is incremented whenever an event of the matching type occurs. A counter may be incremented by an event caused by an instruction which is subsequently flushed (for example, due to mis-speculation). Counting of events may be controlled based on privilege mode or on the strand in which they occur. Masking may be provided to allow counting of subgroups of events (for example, various occurrences of different opcode groups).

A field that indicates when a counter has overflowed must be present in either each performance instrumentation counter or in a separate performance counter control register.

Performance counters are usually provided on a per-strand basis.

11.2.1 Counter Overflow

Overflow of a counter is recorded in the overflow-indication field of either a performance instrumentation counter or a separate performance counter control register.

Programming Note	Counter overflow traps can also be used for sampling, by setting the initial counter value so that an interrupt occurs n counts later.
-------------------------	--

Counter overflow traps are provided so that large counts can be maintained in software, beyond the range directly supported in hardware. The counters continue to count after an overflow, and software can utilize the overflow traps to maintain additional high-order bits.

Traps

A *trap* is a vectored transfer of control to software running in a privilege mode (see page 444) with (typically) greater privileges. A trap in nonprivileged mode can be delivered to privileged mode or hyperprivileged mode. A trap that occurs while executing in privileged mode can be delivered to privileged mode or hyperprivileged mode.

The actual transfer of control occurs through a trap table that contains the first eight instructions (32 instructions for *clean_window*, window spill, and window fill, traps) of each trap handler. The virtual base address of the trap table for traps to be delivered in privileged mode is specified in the Trap Base Address (TBA) register. The displacement within the table is determined by the trap type and the current trap level (TL). One-half of each table is reserved for hardware traps; the other half is reserved for software traps generated by Tcc instructions.

A trap behaves like an unexpected procedure call. It causes the hardware to do the following:

1. Save certain virtual processor state (such as program counters, CWP, ASI, CCR, PSTATE, and the trap type) on a hardware register stack.
2. Enter privileged execution mode with a predefined PSTATE.
3. Begin executing trap handler code in the trap vector.

When the trap handler has finished, it uses either a DONE or RETRY instruction to return.

A trap may be caused by a Tcc instruction, an instruction-induced exception, a reset, an asynchronous error, or an interrupt request not directly related to a particular instruction. The virtual processor must appear to behave as though, before executing each instruction, it determines if there are any pending exceptions or interrupt requests. If there are pending exceptions or interrupt requests, the virtual processor selects the highest-priority exception or interrupt request and causes a trap.

Thus, an *exception* is a condition that makes it impossible for the virtual processor to continue executing the current instruction stream without software intervention. A *trap* is the action taken by the virtual processor when it changes the instruction flow in response to the presence of an exception, interrupt, reset, or Tcc instruction.

V9 Compatibility	Exceptions referred to as “catastrophic error exceptions” in the
Note	SPARC V9 specification do not exist in the Oracle SPARC
	Architecture; they are handled using normal error-reporting
	exceptions. (impl. dep. #31-V8-Cs10)

An *interrupt* is a request for service presented to a virtual processor by an external device.

Traps are described in these sections:

- **Virtual Processor Privilege Modes** on page 444.
- **Virtual Processor States and Traps** on page 445.
- **Trap Categories** on page 445.
- **Trap Control** on page 449.
- **Trap-Table Entry Addresses** on page 450.
- **Trap Processing** on page 458.
- **Exception and Interrupt Descriptions** on page 460.

- Register Window Traps on page 465.

12.1 Virtual Processor Privilege Modes

An Oracle SPARC Architecture virtual processor is always operating in a discrete privilege mode. The privilege modes are listed below in order of increasing privilege:

- Nonprivileged mode (also known as “user mode”)
- Privileged mode, in which supervisor (operating system) software primarily operates
- Hyperprivileged mode (not described in this document)

The virtual processor’s operating mode is determined by the state of two mode bits, as shown in TABLE 12-1.

TABLE 12-1 Virtual Processor Privilege Modes

PSTATE.priv	Virtual Processor Privilege Mode
0	Nonprivileged
1	Privileged

A trap is delivered to the virtual processor in either privileged mode or hyperprivileged mode; in which mode the trap is delivered depends on:

- Its trap type
- The trap level (TL) at the time the trap is taken
- The privilege mode at the time the trap is taken

Traps detected in nonprivileged and privileged mode can be delivered to the virtual processor in privileged mode or hyperprivileged mode.

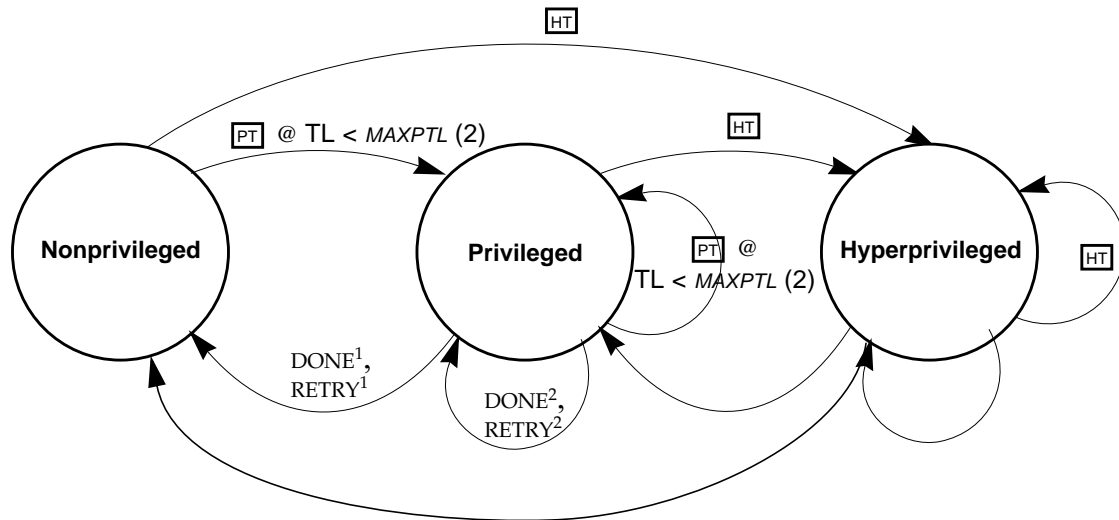
TABLE 12-4 on page 453 indicates in which mode each trap is processed, based on the privilege mode at which it was detected.

A trap delivered to privileged mode uses the privileged-mode trap vector, based upon the TBA register. See *Trap-Table Entry Address to Privileged Mode* on page 450 for details.

The maximum trap level at which privileged software may execute is *MAXPTL* (which, on an Oracle SPARC Architecture 2015 virtual processor, is 2)..

Notes | Execution in nonprivileged mode with TL > 0 is an invalid condition that privileged software should never allow to occur.

FIGURE 12-1 shows how a virtual processor transitions between privilege modes, excluding transitions that can occur due to direct software writes to `PSTATE.priv`. In this figure, `PT` indicates a “trap destined for privileged mode” and `HT` indicates a “trap destined for hyperprivileged mode”.



¹ if (TSTATE[TL].PSTATE.priv = 0)

² if (TSTATE[TL].PSTATE.priv = 1)

FIGURE 12-1 Virtual Processor Privilege Mode Transition Diagram

12.2 Virtual Processor States and Traps

The value of `TL` affects the generated trap vector address. `TL` also determines where (that is, into which element of the `TSTATE` array) the states are saved.

12.2.0.1 Usage of Trap Levels

If `MAXPTL = 2` in an Oracle SPARC Architecture implementation, the trap levels might be used as shown in TABLE 12-2.

TABLE 12-2 Typical Usage for Trap Levels

TL	Corresponding Execution Mode	Usage
0	Nonprivileged	Normal execution
1	Privileged	System calls; interrupt handlers; instruction emulation
2	Privileged	Window spill/fill handler

12.3 Trap Categories

An exception, error, or interrupt request can cause any of the following trap types:

- Precise trap
- Deferred trap
- Disrupting trap

- Reset trap

12.3.1 Precise Traps

A *precise trap* is induced by a particular instruction and occurs before any program-visible state has been changed by the trap-inducing instructions. When a precise trap occurs, several conditions must be true:

- The PC saved in TPC[TL] points to the instruction that induced the trap and the NPC saved in TNPC[TL] points to the instruction that was to be executed next.
- All instructions issued before the one that induced the trap have completed execution.
- Any instructions issued after the one that induced the trap remain unexecuted.

Among the actions that trap handler software might take when processing a precise trap are:

- Return to the instruction that caused the trap and reexecute it by executing a RETRY instruction ($PC \leftarrow \text{old PC}$, $NPC \leftarrow \text{old NPC}$).
- Emulate the instruction that caused the trap and return to the succeeding instruction by executing a DONE instruction ($PC \leftarrow \text{old NPC}$, $NPC \leftarrow \text{old NPC} + 4$).
- Terminate the program or process associated with the trap.

12.3.2 Deferred Traps

A *deferred trap* is also induced by a particular instruction, but unlike a precise trap, a deferred trap may occur after program-visible state has been changed. Such state may have been changed by the execution of either the trap-inducing instruction itself or by one or more other instructions.

There are two classes of deferred traps:

- *Termination deferred traps* — The instruction (usually a store) that caused the trap has passed the retirement point of execution (the TPC has been updated to point to an instruction beyond the one that caused the trap). The trap condition is an error that prevents the instruction from completing and its results becoming globally visible. A termination deferred trap has high trap priority, second only to the priority of resets.

Programming Note | Not enough state is saved for execution of the instruction stream to resume with the instruction that caused the trap. Therefore, the trap handler must terminate the process containing the instruction that caused the trap.

- *Restartable deferred traps* — The program-visible state has been changed by the trap-inducing instruction or by one or more other instructions after the trap-inducing instruction.

SPARC V9 Compatibility Note | A *restartable* deferred trap is the “deferred trap” defined in the SPARC V9 specification.

The fundamental characteristic of a *restartable* deferred trap is that the state of the virtual processor on which the trap occurred may not be consistent with any precise point in the instruction sequence being executed on that virtual processor. When a restartable deferred trap occurs, TPC[TL] and TNPC[TL] contain a PC value and an NPC value, respectively, corresponding to a point in the instruction sequence being executed on the virtual processor. This PC may correspond to the trap-inducing instruction or it may correspond to an instruction following the trap-inducing instruction. With a restartable deferred trap, program-visible updates may be missing from instructions prior to the instruction to which TPC[TL] refers. The missing updates are limited to instructions in the range from (and including) the actual trap-inducing instruction up to (but not including) the instruction to which TPC[TL] refers. By definition, the instruction to which TPC[TL] refers has not yet executed, therefore it cannot have any updates, missing or otherwise.

With a restartable deferred trap there must exist sufficient information to report the error that caused the deferred trap. If system software can recover from the error that caused the deferred trap, then there must be sufficient information to generate a consistent state within the processor so that execution can resume. Included in that information must be an indication of the mode (nonprivileged, privileged, or hyperprivileged) in which the trap-inducing instruction was issued.

How the information necessary for repairing the state to make it consistent state is maintained and how the state is repaired to a consistent state are implementation dependent. It is also implementation dependent whether execution resumes at the point of the trap-inducing instruction or at an arbitrary point between the trap-inducing instruction and the instruction pointed to by the TPC[TL], inclusively.

Associated with a particular restartable deferred trap implementation, the following must exist:

- An instruction that causes a potentially outstanding restartable deferred trap exception to be taken as a trap
- Instructions with sufficient privilege to access the state information needed by software to emulate the restartable deferred trap-inducing instruction and to resume execution of the trapped instruction stream.

Programming Note	Resuming execution may require the emulation of instructions that had not completed execution at the time of the restartable deferred trap, that is, those instructions in the deferred-trap queue.
-------------------------	---

Software should resume execution with the instruction starting at the instruction to which TPC[TL] refers. Hardware should provide enough information for software to recreate virtual processor state and update it to the point just before execution of the instruction to which TPC[TL] refers. After software has updated virtual processor state up to that point, it can then resume execution by issuing a RETRY instruction.

IMPL. DEP. #32-V8-Ms10: Whether any restartable deferred traps (and, possibly, associated deferred-trap queues) are present is implementation dependent.

Among the actions software can take after a restartable deferred trap are these:

- Emulate the instruction that caused the exception, emulate or cause to execute any other execution-deferred instructions that were in an associated restartable deferred trap state queue, and use RETRY to return control to the instruction at which the deferred trap was invoked.
- Terminate the program or process associated with the restartable deferred trap.

A deferred trap (of either of the two classes) is always delivered to the virtual processor in hyperprivileged mode.

12.3.3 Disrupting Traps

12.3.3.1 Disrupting versus Precise and Deferred Traps

A *disrupting trap* is caused by a condition (for example, an interrupt) rather than directly by a particular instruction. This distinguishes it from *precise* and *deferred* traps.

When a disrupting trap has been serviced, trap handler software normally arranges for program execution to resume where it left off. This distinguishes disrupting traps from *reset* traps, since a reset trap vectors to a unique reset address and execution of the program that was running when the reset occurred is generally not expected to resume.

When a disrupting trap occurs, the following conditions are true:

1. The PC saved in TPC[TL] points to an instruction in the disrupted program stream and the NPC value saved in TNPC[TL] points to the instruction that was to be executed after that one.
2. All instructions issued before the instruction indicated by TPC[TL] have retired.
3. The instruction to which TPC[TL] refers and any instruction(s) that were issued after it remain unexecuted.

A disrupting trap may be due to an interrupt request directly related to a previously-executed instruction; for example, when a previous instruction sets a bit in the SOFTINT register.

12.3.3.2 Causes of Disrupting Traps

A disrupting trap may occur due to either an interrupt request or an error not directly related to instruction processing. The source of an interrupt request may be either internal or external. An interrupt request can be induced by the assertion of a signal not directly related to any particular virtual processor or memory state, for example, the assertion of an “I/O done” signal.

A condition that causes a disrupting trap persists until the condition is cleared.

12.3.3.3 Conditioning of Disrupting Traps

How disrupting traps are conditioned is affected by:

- The privilege mode in effect when the trap is outstanding, just before the trap is actually taken (regardless of the privilege mode that was in effect when the exception was detected).
- The privilege mode for which delivery of the trap is destined

Outstanding in Nonprivileged or Privileged mode, destined for delivery in Privileged mode. An outstanding disrupting trap condition in either nonprivileged mode or privileged mode and destined for delivery to privileged mode is held pending while the Interrupt Enable (*ie*) field of PSTATE is zero ($PSTATE.ie = 0$). *interrupt_level_n* interrupts are further conditioned by the Processor Interrupt Level (PIL) register. An interrupt is held pending while either $PSTATE.ie = 0$ or the condition’s interrupt level is less than or equal to the level specified in PIL. When delivery of this disrupting trap is enabled by $PSTATE.ie = 1$, it is delivered to the virtual processor in privileged mode if $TL < MAXPTL$ (2, in Oracle SPARC Architecture 2015 implementations).

Outstanding in Nonprivileged or Privileged mode, destined for delivery in Hyperprivileged mode. An outstanding disrupting trap condition detected while in either nonprivileged mode or privileged mode and destined for delivery in hyperprivileged mode is never masked; it is delivered immediately.

The above is summarized in TABLE 12-3.

TABLE 12-3 Conditioning of Disrupting Traps

Type of Disrupting Trap Condition	Current Virtual Processor Privilege Mode	Disposition of Disrupting Traps, based on privilege mode in which the trap is destined to be delivered	
		Privileged	Hyperprivileged
<i>Interrupt_level_n</i>	Nonprivileged or Privileged	Held pending while $PSTATE.ie = 0$ or interrupt level \leq PIL	—
All other disrupting traps	Nonprivileged or Privileged	Held pending while $PSTATE.ie = 0$	Delivered immediately

12.3.3.4 Trap Handler Actions for Disrupting Traps

Among the actions that trap-handler software might take to process a disrupting trap are:

- Use RETRY to return to the instruction at which the trap was invoked (PC ← old PC, NPC ← old NPC).
- Terminate the program or process associated with the trap.

12.3.4 Uses of the Trap Categories

The SPARC V9 *trap model* stipulates the following:

1. Reset traps occur asynchronously to program execution.
2. When recovery from an exception can affect the interpretation of subsequent instructions, such exceptions shall be precise. See TABLE 12-4, TABLE 12-5, and *Exception and Interrupt Descriptions* on page 460 for identification of which traps are precise.
3. In an Oracle SPARC Architecture implementation, all exceptions that occur as the result of program execution are precise (impl. dep. #33-V8-Cs10).
4. An error detected after the initial access of a multiple-access load instruction (for example, LDTX or LDBLOCKF^D) should be precise. Thus, a trap due to the second memory access can occur. However, the processor state should not have been modified by the first access.
5. Exceptions caused by external events unrelated to the instruction stream, such as interrupts, are disrupting.

A deferred trap may occur one or more instructions after the trap-inducing instruction is dispatched.

12.4 Trap Control

Several registers control how any given exception is processed, for example:

- The interrupt enable (ie) field in PSTATE and the Processor Interrupt Level (PIL) register control interrupt processing. See *Disrupting Traps* on page 447 for details.
- The enable floating-point unit (fef) field in FPRS, the floating-point unit enable (pef) field in PSTATE, and the trap enable mask (tem) in the FSR control floating-point traps.
- The TL register, which contains the current level of trap nesting, affects whether the trap is processed in privileged mode or hyperprivileged mode.
- PSTATE.tle determines whether implicit data accesses in the trap handler routine will be performed using big-endian or little-endian byte order.

Between the execution of instructions, the virtual processor prioritizes the outstanding exceptions, errors, and interrupt requests. At any given time, only the highest-priority exception, error, or interrupt request is taken as a trap. When there are multiple interrupts outstanding, the interrupt with the highest interrupt level is selected. When there are multiple outstanding exceptions, errors, and/or interrupt requests, a trap occurs based on the exception, error, or interrupt with the highest priority (numerically lowest priority number in TABLE 12-5). See *Trap Priorities* on page 458.

12.4.1 PIL Control

When an interrupt request occurs, the virtual processor compares its interrupt request level against the value in the Processor Interrupt Level (PIL) register. If the interrupt request level is greater than PIL and no higher-priority exception is outstanding, then the virtual processor takes a trap using the appropriate *interrupt_level_n* trap vector.

12.4.2 FSR.tem Control

The occurrence of floating-point traps of type IEEE_754_exception can be controlled with the user-accessible trap enable mask (tem) field of the FSR. If a particular bit of FSR.tem is 1, the associated IEEE_754_exception can cause an *fp_exception_ieee_754* trap.

If a particular bit of FSR.tem is 0, the associated IEEE_754_exception does not cause an *fp_exception_ieee_754* trap. Instead, the occurrence of the exception is recorded in the FSR's accrued exception field (aexc).

If an IEEE_754_exception results in an *fp_exception_ieee_754* trap, then the destination F register, FSR.fccn, and FSR.aexc fields remain unchanged. However, if an IEEE_754_exception does not result in a trap, then the F register, FSR.fccn, and FSR.aexc fields are updated to their new values.

12.5 Trap-Table Entry Addresses

Traps are delivered to the virtual processor in either privileged mode or hyperprivileged mode, depending on the trap type, the value of TL at the time the trap is taken, and the privilege mode at the time the exception was detected. See TABLE 12-4 on page 453 and TABLE 12-5 on page 456 for details.

Unique trap table base addresses are provided for traps being delivered in privileged mode and in hyperprivileged mode.

12.5.1 Trap-Table Entry Address to Privileged Mode

Privileged software initializes bits 63:15 of the Trap Base Address (TBA) register (its most significant 49 bits) with bits 63:15 of the desired 64-bit privileged trap-table base address.

At the time a trap to privileged mode is taken:

- Bits 63:15 of the trap vector address are taken from TBA{63:15}.
- Bit 14 of the trap vector address (the "TL>0" field) is set based on the value of TL just before the trap is taken; that is, if TL = 0 then bit 14 is set to 0 and if TL > 0 then bit 14 is set to 1.
- Bits 13:5 of the trap vector address contain a copy of the contents of the TT register (TT[TL]).
- Bits 4:0 of the trap vector address are always 0; hence, each trap table entry is at least 2⁵ or 32 bytes long. Each entry in the trap table may contain the first eight instructions of the corresponding trap handler.

FIGURE 12-2 illustrates the trap vector address for a trap delivered to privileged mode. In FIGURE 12-2, the "TL>0" bit is 0 if TL = 0 when the trap was taken, and 1 if TL > 0 when the trap was taken. This implies, as detailed in the following section, that there are two trap tables for traps to privileged mode: one for traps from TL = 0 and one for traps from TL > 0.




FIGURE 12-2 Privileged Mode Trap Vector Address

12.5.2 Privileged Trap Table Organization

The layout of the privileged-mode trap table (which is accessed using virtual addresses) is illustrated in FIGURE 12-3.

Value of TL (before trap)	Software Trap Type	Hardware Trap Type (TT[TL])	Trap Table Offset (from TBA)	Contents of Trap Table
TL = 0	—	000 ₁₆ –07F ₁₆	0 ₁₆ – FE0 ₁₆	Hardware traps
	—	080 ₁₆ –0FF ₁₆	1000 ₁₆ –1FE0 ₁₆	Spill / fill traps
	0 ₁₆ – 7F ₁₆	100 ₁₆ –17F ₁₆	2000 ₁₆ –2FE0 ₁₆	Software traps to Privileged level
	—	180 ₁₆ –1FF ₁₆	3000 ₁₆ –3FE0 ₁₆	<i>unassigned</i>
TL = 1 (TL = MAXPTL–1)	—	000 ₁₆ –07F ₁₆	4000 ₁₆ –4FE0 ₁₆	Hardware traps
	—	080 ₁₆ –0FF ₁₆	5000 ₁₆ –5FE0 ₁₆	Spill / fill traps
	0 ₁₆ – 7F ₁₆	100 ₁₆ –17F ₁₆	6000 ₁₆ –6FE0 ₁₆	Software traps to Privileged level
	—	180 ₁₆ –1FF ₁₆	7000 ₁₆ –7FE0 ₁₆	<i>unassigned</i>

FIGURE 12-3 Privileged-mode Trap Table Layout

The trap table for TL = 0 comprises 512 thirty-two-byte entries; the trap table for TL > 0 comprises 512 more thirty-two-byte entries. Therefore, the total size of a full privileged trap table is $2 \times 512 \times 32$ bytes (32 Kbytes). However, if privileged software does not use software traps (Tcc instructions) at TL > 0, the table can be made 24 Kbytes long.

12.5.3 Trap Type (TT)

When a normal trap occurs, a value that uniquely identifies the type of the trap is written into the current 9-bit TT register (TT[TL]) by hardware. Control is then transferred into the trap table to an address formed by the trap's destination privilege mode:

- The TBA register, (TL > 0), and TT[TL] (see *Trap-Table Entry Address to Privileged Mode* on page 450)

TT values 000₁₆–0FF₁₆ are reserved for hardware traps. TT values 100₁₆–17F₁₆ are reserved for software traps (caused by execution of a Tcc instruction) to privileged-mode trap handlers.

IMPL. DEP. #35-V8-Cs20: TT values 060₁₆ to 07F₁₆ were reserved for *implementation_dependent_exception_n* exceptions in the SPARC V9 specification, but are now all defined as standard Oracle SPARC Architecture exceptions. See TABLE 12-4 for details.

The assignment of TT values to traps is shown in TABLE 12-4; TABLE 12-5 provides the same list, but sorted in order of trap priority. The key to both tables follows:

Symbol	Meaning
●	This trap type is associated with a feature that is architecturally required in an implementation of Oracle SPARC Architecture 2015. Hardware must detect this exception or interrupt, trap on it (if not masked), and set the specified trap type value in the TT register.
○	This trap type is associated with a feature that is architecturally defined in Oracle SPARC Architecture 2015, but its implementation is optional.
P	Trap is taken via the Privileged trap table, in Privileged mode (PSTATE.priv = 1)
H	Trap is taken in Hyperprivileged mode
-x-	Not possible. Hardware cannot generate this trap in the indicated running mode. For example, all privileged instructions can be executed in privileged mode, therefore a <i>privileged_opcode</i> trap cannot occur in privileged mode.
—	This trap is reserved for future use.
(ie)	When the outstanding disrupting trap condition occurs in this privilege mode, it may be conditioned (masked out) by PSTATE.ie = 0 (but remains pending).
(nm)	Never Masked — when the condition occurs in this running mode, it is never masked out and the trap is always taken.
(pend)	Held Pending — the condition can <i>occur</i> in this running mode, but can't be <i>serviced</i> in this mode. Therefore, it is held pending until the mode changes to one in which the exception <i>can</i> be serviced.

TABLE 12-4 Exception and Interrupt Requests, by TT Value (1 of 3)

UA-2015 ●=Req'd. ○=Opt'l	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = High- est)	Mode in which Trap is Delivered (and Conditioning Applied), based on Current Privilege Mode	
					NP	Priv
—	<i>Reserved</i>	000 ₁₆	—	—	—	—
●	<i>(used at higher privilege levels)</i>	001 ₁₆ – 005 ₁₆	—	—	—	—
—	<i>Reserved</i>	005 ₁₆	—	—	—	—
—	<i>implementation-dependent</i>	006 ₁₆	—	—	—	—
●	<i>IAE_privilege_violation</i>	008 ₁₆	precise	3.1	H	-x-
●	<i>(used at higher privilege levels)</i>	009 ₁₆	—	—	—	—
●	<i>(used at higher privilege levels)</i>	00A ₁₆	—	—	—	—
●	<i>IAE_unauth_access</i>	00B ₁₆	precise	3.2	H	H
●	<i>IAE_nfo_page</i>	00C ₁₆	precise	3.3	H	H
—	<i>Reserved</i>	00F ₁₆	—	—	—	—
●	<i>illegal_instruction</i>	010 ₁₆	precise	6.2	H	H
●	<i>privileged_opcode</i>	011 ₁₆	precise	7	P (nm)	-x-
○	<i>unimplemented_LDTW</i>	012 ₁₆	precise	6.3	H	H
○	<i>unimplemented_STTW</i>	013 ₁₆	precise	6.3	H	H
●	<i>DAE_invalid_asi</i>	014 ₁₆	precise	11.3	H	H
	<i>DAE_privilege_violation</i>	015 ₁₆	<i>precise</i>	12.06	<i>H</i>	<i>H</i>
●	<i>DAE_nc_page</i>	016 ₁₆	precise	12.07	H	H
●	<i>DAE_nfo_page</i>	017 ₁₆	precise	12.08	H	H
—	<i>Reserved</i>	018 ₁₆	—	—	—	—
—	<i>Reserved</i>	019 ₁₆ – 01B ₁₆	—	—	—	—
—	<i>Reserved</i>	01C ₁₆	—	—	—	—
—	<i>Reserved</i>	01D ₁₆ – 01F ₁₆	—	—	—	—
●	<i>fp_disabled</i>	020 ₁₆	precise	8	P (nm)	P (nm)
○	<i>fp_exception_ieee_754</i>	021 ₁₆	precise	11.1	P (nm)	P (nm)
○	<i>fp_exception_other</i>	022 ₁₆	precise	11.1	P (nm)	P (nm)
●	<i>tag_overflow^D</i>	023 ₁₆	precise	14	P (nm)	P (nm)
●	<i>clean_window</i>	024 ₁₆ [‡] – 027 ₁₆	precise	10.1	P (nm)	P (nm)

TABLE 12-4 Exception and Interrupt Requests, by TT Value (2 of 3)

UA-2015 ●=Req'd. ○=Opt'l	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = High- est)	Mode in which Trap is Delivered (and Conditioning Applied), based on Current Privilege Mode	
					NP	Priv
●	<i>division_by_zero</i>	028 ₁₆	precise	15	P (nm)	P (nm)
□		02C ₁₆	disrupting	16.03	H	H
—	<i>compatibility_feature</i>	02F ₁₆	precise	6.3	H	H
●	<i>DAE_side_effect_page</i>	030 ₁₆	precise	12.08	H	H
—	<i>Reserved</i>	032 ₁₆	—	—	—	—
●	<i>mem_address_not_aligned</i>	034 ₁₆	precise	10.2	H	H
●	<i>LDDF_mem_address_not_aligned</i>	035 ₁₆	precise	10.1	H	H
●	<i>STDF_mem_address_not_aligned</i>	036 ₁₆	precise	10.1	H	H
●	<i>privileged_action</i>	037 ₁₆	precise	11.1	H	H
○	<i>LDQF_mem_address_not_aligned</i>	038 ₁₆	precise	10.1	H	H
○	<i>STQF_mem_address_not_aligned</i>	039 ₁₆	precise	10.1	H	H
○	<i>STTW_exception</i>	03A ₁₆	precise	12.13	H	H
—	<i>Reserved</i>	03B ₁₆	—	—	—	—
○	<i>BLD_exception</i>	03C ₁₆	precise	12.13	H	H
○	<i>BST_exception</i>	03D ₁₆	precise	12.13	H	H
○	<i>(used at higher privilege levels)</i>	040 ₁₆	—	—	—	—
●	<i>interrupt_level_n (n = 1–15)</i>	041 ₁₆ – 04F ₁₆	disrupting	32- <i>n</i> (31 to 17)	P (ie)	P (ie)
—	<i>Reserved</i>	050 ₁₆	—	—	—	—
—	<i>Reserved</i>	051 ₁₆ – 05D ₁₆	—	—	—	—
●	<i>(used at higher privilege levels)</i>	05F ₁₆ – 061 ₁₆	—	—	—	—
○	<i>(used at higher privilege levels)</i>	060 ₁₆	—	—	—	—
○	<i>VA_watchpoint</i>	062 ₁₆	precise	11.2	P (nm)	P (nm)
●	<i>(used at higher privilege levels)</i>	063 ₁₆ – 06C ₁₆	—	—	—	—
○	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	070 ₁₆	—	∇	—	—
●	<i>(used at higher privilege levels)</i>	071 ₁₆ – 072 ₁₆	—	—	—	—
○	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	073 ₁₆	—	∇	—	—
●	<i>control_transfer_instruction</i>	074 ₁₆	precise	11.1	P	P
○	<i>instruction_VA_watchpoint</i>	075 ₁₆	precise	2.05	P (nm)	P (nm)

TABLE 12-4 Exception and Interrupt Requests, by TT Value (3 of 3)

UA-2015 ●=Req'd. ○=Opt'l	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = High- est)	Mode in which Trap is Delivered (and Conditioning Applied), based on Current Privilege Mode	
					NP	Priv
□	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	077 ₁₆ – 078 ₁₆	—	∇	—	—
□	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	079 ₁₆ – 07B ₁₆	—	∇	—	—
●	<i>cpu_mondo</i>	07C ₁₆	disrupting	16.08	P (ie)	P (ie)
●	<i>dev_mondo</i>	07D ₁₆	disrupting	16.11	P (ie)	P (ie)
●	<i>resumable_error</i>	07E ₁₆	disrupting	33.3	P (ie)	P (ie)
—	<i>nonresumable_error</i>	07F ₁₆	—	—	—	—
●	<i>spill_n_normal</i> (<i>n</i> = 0–7)	080 ₁₆ [‡] – 09F ₁₆	precise	9	P (nm)	P (nm)
●	<i>spill_n_other</i> (<i>n</i> = 0–7)	0A0 ₁₆ [‡] – 0BF ₁₆	precise	9	P (nm)	P (nm)
●	<i>fill_n_normal</i> (<i>n</i> = 0–7)	0C0 ₁₆ [‡] – 0DF ₁₆	precise	9	P (nm)	P (nm)
●	<i>fill_n_other</i> (<i>n</i> = 0–7)	0E0 ₁₆ [‡] – 0FF ₁₆	precise	9	P (nm)	P (nm)
●	<i>trap_instruction</i>	100 ₁₆ – 17F ₁₆	precise	16.05	P (nm)	P (nm)
●	<i>htrap_instruction</i>	180 ₁₆ – 1FF ₁₆	precise	16.05	-x-	

* Although these trap priorities are recommended, all trap priorities are implementation dependent (impl. dep. #36-V8 on page 458), including relative priorities within a given priority level.

‡ The trap vector entry (32 bytes) for this trap type plus the next three trap types (total of 128 bytes) are permanently reserved for this exception.

^D This exception is deprecated, because the only instructions that can generate it have been deprecated.

TABLE 12-5 Exception and Interrupt Requests, by Priority (1 of 2)

UA-2015 ●=Req'd. ○=Opt'l □.=Impl-Specific	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = Highest)	Mode in which Trap is Delivered and (and Conditioning Applied), based on Current Privilege Mode	
					NP	Priv
○	<i>instruction_VA_watchpoint</i>	075 ₁₆	precise	2.05	P (nm)	P (nm)
●	<i>IAE_privilege_violation</i>	008 ₁₆	precise	3.1	H	-x-
●	<i>IAE_unauth_access</i>	00B ₁₆	precise	3.2	H	H
●	<i>IAE_nfo_page</i>	00C ₁₆	precise	3.3	H	H
●	<i>illegal_instruction</i>	010 ₁₆	precise	6.2	H	H
○	<i>unimplemented_LDTW</i>	012 ₁₆	precise	6.3	H	H
○	<i>unimplemented_STTW</i>	013 ₁₆	precise		H	H
—	<i>compatibility_feature</i>	02F ₁₆	precise		H	H
●	<i>privileged_opcode</i>	011 ₁₆	precise	7	P (nm)	-x-
●	<i>fp_disabled</i>	020 ₁₆	precise	8	P (nm)	P (nm)
●	<i>spill_n_normal</i> (n = 0–7)	080 ₁₆ [±] – 09F ₁₆	precise	9	P (nm)	P (nm)
●	<i>spill_n_other</i> (n = 0–7)	0A0 ₁₆ [±] – 0BF ₁₆	precise		P (nm)	P (nm)
●	<i>fill_n_normal</i> (n = 0–7)	0C0 ₁₆ [±] – 0DF ₁₆	precise		P (nm)	P (nm)
●	<i>fill_n_other</i> (n = 0–7)	0E0 ₁₆ [±] – 0FF ₁₆	precise		P (nm)	P (nm)
●	<i>clean_window</i>	024 ₁₆ [±] – 027 ₁₆	precise	10.1	P (nm)	P (nm)
●	<i>LDDF_mem_address_not_aligned</i>	035 ₁₆	precise		H	H
●	<i>STDF_mem_address_not_aligned</i>	036 ₁₆	precise		H	H
○	<i>LDQF_mem_address_not_aligned</i>	038 ₁₆	precise		H	H
○	<i>STQF_mem_address_not_aligned</i>	039 ₁₆	precise		H	H
●	<i>mem_address_not_aligned</i>	034 ₁₆	precise	10.2	H	H
○	<i>fp_exception_other</i>	022 ₁₆	precise	11.1	P (nm)	P (nm)
○	<i>fp_exception_ieee_754</i>	021 ₁₆	precise		P (nm)	P (nm)
●	<i>privileged_action</i>	037 ₁₆	precise		H	H
●	<i>control_transfer_instruction</i>	074 ₁₆	precise		P	P

TABLE 12-5 Exception and Interrupt Requests, by Priority (2 of 2)

UA-2015 ●=Req'd. ○=Opt'l □=Impl-Specific	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = Highest)	Mode in which Trap is Delivered and (and Conditioning Applied), based on Current Privilege Mode	
					NP	Priv
○	<i>VA_watchpoint</i>	062 ₁₆	precise	11.2	P (nm)	P (nm)
●	<i>DAE_invalid_asl</i>	014 ₁₆	precise	11.3	H	H
●	<i>DAE_privilege_violation</i>	015 ₁₆	precise	12.06	H	H
●	<i>DAE_nc_page</i>	016 ₁₆	precise	12.07	H	H
●	<i>DAE_nfo_page</i>	017 ₁₆	precise	12.08	H	H
●	<i>DAE_side_effect_page</i>	030 ₁₆	precise		H	H
○	<i>STTW_exception</i>	03A ₁₆	precise	12.13	H	H
○	<i>BLD_exception</i>	03C ₁₆	precise		H	H
○	<i>BST_exception</i>	03D ₁₆	precise		H	H
●	<i>tag_overflow^D</i>	023 ₁₆	precise	14	P (nm)	P (nm)
●	<i>division_by_zero</i>	028 ₁₆	precise	15	P (nm)	P (nm)
□		02C ₁₆	disrupting	16.03	H	H
●	<i>trap_instruction</i>	100 ₁₆ – 17F ₁₆	precise	16.05	P (nm)	P (nm)
●	<i>htrap_instruction</i>	180 ₁₆ – 1FF ₁₆	precise		-x-	
●	<i>cpu_mondo</i>	07C ₁₆	disrupting	16.08	P (ie)	P (ie)
●	<i>dev_mondo</i>	07D ₁₆	disrupting	16.11	P (ie)	P (ie)
●	<i>interrupt_level_n</i> (<i>n</i> = 1–15)	041 ₁₆ – 04F ₁₆	disrupting	32- <i>n</i> (31 to 17)	P (ie)	P (ie)
●	<i>resumable_error</i>	07E ₁₆	disrupting	33.3	P (ie)	P (ie)
—	<i>nonresumable_error</i>	07F ₁₆	—	—	—	—

* Although these trap priorities are recommended, all trap priorities are implementation dependent (impl. dep. #36-V8 on page 458), including relative priorities within a given priority level.

‡ The trap vector entry (32 bytes) for this trap type plus the next three trap types (total of 128 bytes) are permanently reserved for this exception.

^D This exception is deprecated, because the only instructions that can generate it have been deprecated.

12.5.3.1 Trap Type for Spill/Fill Traps

The trap type for window *spill/fill* traps is determined on the basis of the contents of the OTHERWIN and WSTATE registers as described below and shown in FIGURE 12-4.

Bit	Field	Description
8:6	spill_or_fill	010 ₂ for spill traps; 011 ₂ for fill traps
5	other	(OTHERWIN ≠ 0)
4:2	wtype	If (other) then WSTATE.other; else WSTATE.normal




FIGURE 12-4 Trap Type Encoding for Spill/Fill Traps

12.5.4 Trap Priorities

TABLE 12-4 on page 453 and TABLE 12-5 on page 456 show the assignment of traps to TT values and the relative priority of traps and interrupt requests. A trap priority is an ordinal number, with 0 indicating the highest priority and greater priority numbers indicating decreasing priority; that is, if $x < y$, a pending exception or interrupt request with priority x is taken instead of a pending exception or interrupt request with priority y . Traps within the same priority class (0 to 33) are listed in priority order in TABLE 12-5 (impl. dep. #36-V8).

IMPL. DEP. #36-V8: The relative priorities of traps defined in the Oracle SPARC Architecture are fixed. However, the absolute priorities of those traps are implementation dependent (because a future version of the architecture may define new traps). The priorities (both absolute and relative) of any new traps are implementation dependent.

However, the TT values for the exceptions and interrupt requests shown in TABLE 12-4 and TABLE 12-5 must remain the same for every implementation.

The trap priorities given above always need to be considered within the context of how the virtual processor actually issues and executes instructions.

12.6 Trap Processing

The virtual processor's action during trap processing depends on various virtual processor states, including the trap type, the current level of trap nesting (given in the TL register), and PSTATE. When a trap occurs, the GL register is normally incremented by one (described later in this section), which replaces the set of eight global registers with the next consecutive set.

During normal operation, the virtual processor is in `execute_state`. It processes traps in `execute_state` and continues.

TABLE 12-6 describes the virtual processor mode and trap-level transitions involved in handling traps.

TABLE 12-6 Trap Received While in `execute_state`

Original State	New State, After Receiving Trap or Interrupt
<code>execute_state</code> $TL < MAXPTL - 1$	<code>execute_state</code> $TL \leftarrow TL + 1$

12.6.1 Normal Trap Processing

A trap is delivered in either privileged mode or hyperprivileged mode, depending on the type of trap, the trap level (TL), and the privilege mode in effect when the exception was detected.

During normal trap processing, the following state changes occur (conceptually, in this order):

- The trap level is updated. This provides access to a fresh set of privileged trap-state registers used to save the current state, in effect, pushing a frame on the trap stack.

$$TL \leftarrow TL + 1$$

- Existing state is preserved.

$$TSTATE[TL].gl \leftarrow GL$$

$$TSTATE[TL].ccr \leftarrow CCR$$

$$TSTATE[TL].asi \leftarrow ASI$$

$$TSTATE[TL].pstate \leftarrow PSTATE$$

$$TSTATE[TL].cwp \leftarrow CWP$$

$$TPC[TL] \leftarrow PC \text{ // (upper 32 bits zeroed if } PSTATE.am = 1)$$

$$TNPC[TL] \leftarrow NPC \text{ // (upper 32 bits zeroed if } PSTATE.am = 1)$$

- The trap type is preserved.

$$TT[TL] \leftarrow \text{the trap type}$$

- The Global Level register (GL) is updated. This normally provides access to a fresh set of global registers:

$$GL \leftarrow \min(GL + 1, MAXPGL)$$

- The PSTATE register is updated to a predefined state:

`PSTATE.mm` is unchanged

`PSTATE.pef` $\leftarrow 1$ // if an FPU is present, it is enabled

`PSTATE.am` $\leftarrow 0$ // address masking is turned off

`PSTATE.priv` $\leftarrow 1$ // the virtual processor enters privileged mode

`PSTATE.cle` $\leftarrow PSTATE.tle$ //set endian mode for traps

endif

`PSTATE.ie` $\leftarrow 0$ // interrupts are disabled

`PSTATE.tle` is unchanged

`PSTATE.tct` $\leftarrow 0$ // trap on CTI disabled

- For a register-window trap (*clean_window*, window spill, or window fill) only, CWP is set to point to the register window that must be accessed by the trap-handler software, that is:

if $TT[TL] = 024_{16}$ // a *clean_window* trap

then $CWP \leftarrow (CWP + 1) \bmod N_REG_WINDOWS;$

endif

if $(080_{16} \leq TT[TL] \leq 0BF_{16})$ // window spill trap

then $CWP \leftarrow (CWP + CANSERVE + 2) \bmod N_REG_WINDOWS;$

endif

if $(0C0_{16} \leq TT[TL] \leq 0FF_{16})$ // window fill trap

then $CWP \leftarrow (CWP - 1) \bmod N_REG_WINDOWS;$

endif

For non-register-window traps, CWP is not changed.

- Control is transferred into the trap table:


```
// Note that at this point, TL has already been incremented (above)
if ( (trap is to privileged mode) and (TL ≤ MAXPTL) )
then
    //the trap is handled in privileged mode
    //Note: The expression "(TL > 1)" below evaluates to the
    //value 02 if TL was 0 just before the trap (in which
    //case, TL = 1 now, since it was incremented above,
    //during trap entry). "(TL > 1)" evaluates to 12 if
    //TL was > 0 before the trap.
    PC ← TBA{63:15} :: (TL > 1) :: TT[TL] :: 0 00002
    NPC ← TBA{63:15} :: (TL > 1) :: TT[TL] :: 0 01002
else { trap is handled in hyperprivileged mode }
endif
```

Interrupts are ignored as long as PSTATE.ie = 0.

Programming Note State in TPC[*n*], TNPC[*n*], TSTATE[*n*], and TT[*n*] is only changed autonomously by the processor when a trap is taken while TL = *n* - 1; however, software can change any of these values with a WRPR instruction when TL = *n*.

12.7 Exception and Interrupt Descriptions

The following sections describe the various exceptions and interrupt requests and the conditions that cause them. Each exception and interrupt request describes the corresponding trap type as defined by the trap model.

All other trap types are reserved.

Note The encoding of trap types in the Oracle SPARC Architecture differs from that shown in *The SPARC Architecture Manual-Version 9*. Each trap is marked as precise, deferred, disrupting, or reset. Example exception conditions are included for each exception type. Chapter 7, *Instructions*, enumerates which traps can be generated by each instruction.

The following traps are generally expected to be supported in all Oracle SPARC Architecture 2015 implementations. A given trap is not required to be supported in an implementation in which the conditions that cause the trap can never occur.

- **clean_window** [TT = 024₁₆-027₁₆] (Precise) — A SAVE instruction discovered that the window about to be used contains data from another address space; the window must be cleaned before it can be used.

IMPL. DEP. #102-V9: An implementation may choose either to implement automatic cleaning of register windows in hardware or to generate a *clean_window* trap, when needed, so that window(s) can be cleaned by software. If an implementation chooses the latter option, then support for this trap type is mandatory.
- **control_transfer_instruction** [TT = 074₁₆] (Precise) — This exception is generated if PSTATE.tct = 1 and the processor determines that a successful control transfer will occur as a result of execution of that instruction. If such a transfer will occur, the processor generates a *control_transfer_instruction* precise trap (trap type = 74₁₆) instead of completing the control transfer. The pc stored in TPC[TL] is the address of the CTI, and the TNPC[TL] is set to the value of NPC before the CTI is executed. (impl. dep. #450-S20). PSTATE.tct is always set to 0 as part of

normal entry into a trap handler. When this exception occurs in nonprivileged or privileged mode, the trap is delivered in privileged mode. If it occurs in hyperprivileged mode, the trap is delivered in hyperprivileged mode.

- **cpu_mondo** [TT = 07C₁₆] (Disrupting) — This interrupt is generated when another virtual processor has enqueued a message for this virtual processor. It is used to deliver a trap in privileged mode, to inform privileged software that an interrupt report has been appended to the virtual processor's CPU mondo queue. A direct message between virtual processors is sent via a CPU mondo interrupt. When the CPU mondo queue contains a valid entry, a *cpu_mondo* exception is sent to the target virtual processor.

Programming Note	It is possible that an implementation may occasionally cause a <i>cpu_mondo</i> interrupt when the CPU Mondo queue is empty (CPU Mondo Queue Head pointer = CPU Mondo Queue Tail pointer). A guest operating system running in privileged mode should handle this by ignoring any CPU Mondo interrupt with an empty queue.
-------------------------	--

SPARC V9 Compatibility Note	The <i>data_access_exception</i> exception from SPARC V9 and UltraSPARC Architecture 2005 has been replaced by more specific exceptions, such as <i>DAE_invalid_asl</i> , <i>DAE_nc_page</i> , <i>DAE_nfo_page</i> , <i>DAE_privilege_violation</i> , and <i>DAE_side_effect_page</i> .
------------------------------------	---

- **compatibility_feature** [TT = 02F₁₆] (Precise) — This exception is generated if an attempt is made to execute an extended-capability instruction that is currently disabled by the Compatibility Feature Register (CFR)
- **DAE_invalid_asl** [TT = 014₁₆] (Precise) — An attempt was made to execute an invalid combination of instruction and ASI. See the instruction descriptions in Chapter 7 for a detailed list of valid ASIs for each instruction that can access alternate address spaces. The following invalid combinations of instruction, ASI, and virtual address cause a *DAE_invalid_asl* exception:
 - A load, store, load-store, or PREFETCHA instruction with either an invalid ASI or an invalid virtual address for a valid ASI.
 - A disallowed combination of instruction and ASI (see *Block Load and Store ASIs* on page 436 and *Partial Store ASIs* on page 437). This includes the following:
 - an attempt to use a (deprecated) atomic quad load ASI (24₁₆, 2C₁₆, 34₁₆, or 3C₁₆) with any load alternate opcode other than LDTXA's (which is shared by LDDA)
 - an attempt to use a nontranslating ASI value with any load or store alternate instruction other than LDXA, LDDFA, STXA, or STDFA
 - an attempt to read from a write-only ASI-accessible register, or load from a store-only ASI (for example, a block commit store ASI, E0₁₆ or E1₁₆)
 - an attempt to write to a read-only ASI-accessible register
- **DAE_nc_page** [TT = 016₁₆] (Precise) — An access to a noncacheable page (TTE.cp = 0) was attempted by an atomic load-store instruction (CASA, CASXA, SWAP, SWAPA, LDSTUB, or LDSTUBA), an LDTXA instruction, a LDBLOCKF^D instruction, a STPARTIALF instruction.
- **DAE_nfo_page** [TT = 017₁₆] (Precise) — An attempt was made to access a non-faulting-only page (TTE.nfo = 1) by any type of load, store, load-store, or FLUSH instruction with an ASI other than a nonfaulting ASI (PRIMARY_NO_FAULT[_LITTLE] or SECONDARY_NO_FAULT[_LITTLE]).
- **DAE_privilege_violation** [TT = 015₁₆] (Precise) — A privilege violation occurred, due to an attempt to access a privileged page (TTE.p = 1) by any type of load, store, or load-store instruction when executing in nonprivileged mode (PSTATE.priv = 0). This includes the special case of an access by privileged software using one of the ASI_AS_IF_USER_PRIMARY[_LITTLE] or ASI_AS_IF_USER_SECONDARY[_LITTLE] ASIs.

- **DAE_side_effect_page** [TT = 030₁₆] (Precise) — An attempt was made to access a page which may cause side effects (TTE.e = 1) by any type of load instruction with nonfaulting ASI.
- **dev_mondo** [TT = 07D₁₆] (Disrupting) — This interrupt causes a trap to be delivered in privileged mode, to inform privileged software that an interrupt report has been appended to its device mondo queue. When a virtual processor has appended a valid entry to a target virtual processor's device mondo queue, it sends a *dev_mondo* exception to the target virtual processor. The interrupt report contents are device specific.

Programming Note	It is possible that an implementation may occasionally cause a <i>dev_mondo</i> interrupt when the Device Mondo queue is empty (Device Mondo Queue Head pointer = Device Mondo Queue Tail pointer). A guest operating system running in privileged mode should handle this by ignoring any Device Mondo interrupt with an empty queue.
-------------------------	--

- **disrupting_performance_event** [TT = 02C₁₆] (Disrupting) — One of the virtual processor's performance control registers is configured for a disrupting performance event and the associated counter has overflowed.
 - **division_by_zero** [TT = 028₁₆] (Precise) — An integer divide instruction attempted to divide by zero.
 - **fill_n_normal** [TT = 0C0₁₆–0DF₁₆] (Precise)
 - **fill_n_other** [TT = 0E0₁₆–0FF₁₆] (Precise)
- A RESTORE or RETURN instruction has determined that the contents of a register window must be restored from memory.
- **fp_disabled** [TT = 020₁₆] (Precise) — An attempt was made to execute an FPop, a floating-point branch, or a floating-point load/store instruction while an FPU was disabled (PSTATE.pef = 0 or FPRS.fef = 0).
 - **fp_exception_ieee_754** [TT = 021₁₆] (Precise) — An FPop instruction generated an IEEE_754_exception and its corresponding trap enable mask (FSR.tem) bit was 1. The floating-point exception type, IEEE_754_exception, is encoded in the FSR.ftt, and specific IEEE_754_exception information is encoded in FSR.cexc.
 - **fp_exception_other** [TT = 022₁₆] (Precise) — An FPop instruction generated an exception other than an IEEE_754_exception. Example: execution of an FPop requires software assistance to complete. The floating-point exception type is encoded in FSR.ftt.
 - **htrap_instruction** [TT = 180₁₆–1FF₁₆] (Precise) — A Tcc instruction was executed in privileged mode, the trap condition evaluated to TRUE, and the software trap number was greater than 127. The trap is delivered in hyperprivileged mode. See also *trap_instruction* on page 464.
 - **IAE_nfo_page** [TT = 00C₁₆] (Precise) — An instruction-access exception occurred as a result of an attempt to fetch an instruction from a memory page which was marked for access only by nonfaulting loads (TTE.nfo = 1).
 - **IAE_privilege_violation** [TT = 008₁₆] (Precise) — An instruction-access exception occurred as a result of an attempt to fetch an instruction from a privileged memory page (TTE.p = 1) while the virtual processor was executing in nonprivileged mode.
 - **IAE_unauth_access** [TT = 00B₁₆] (Precise) — An instruction-access exception occurred as a result of an attempt to fetch an instruction from a memory page which was missing "execute" permission (TTE.ep = 0).
 - **illegal_instruction** [TT = 010₁₆] (Precise) — An attempt was made to execute an ILLTRAP instruction, an instruction with an unimplemented opcode, an instruction with invalid field usage, or an instruction that would result in illegal processor state.

Examples of cases in which *illegal_instruction* is generated include the following:

- An instruction encoding does not match any of the opcode map definitions (see Appendix A, *Opcode Maps*).

- An instruction is not implemented in hardware.
- A reserved instruction field in Tcc instruction is nonzero.
If a reserved instruction field in an instruction other than Tcc is nonzero, an *illegal_instruction* exception should be, but is not required to be, generated. (See *Reserved Opcodes and Instruction Fields* on page 93.)
- An illegal value is present in an instruction i field.
- An illegal value is present in a field that is explicitly defined for an instruction, such as cc2, cc1, cc0, fcn, impl, rcond, or opf_cc.
- Illegal register alignment (such as odd rd value in a doubleword load instruction).
- Illegal rd value for LDXFSR, STXFSR, or the deprecated instructions LDFSR or STFSR.
- ILLTRAP instruction.
- DONE or RETRY when TL = 0.

All causes of an *illegal_instruction* exception are described in individual instruction descriptions in Chapter 7, *Instructions*.

SPARC V9 Compatibility Note	The <i>instruction_access_exception</i> exception from SPARC V9 has been replaced by more specific exceptions, such as <i>IAE_privilege_violation</i> and <i>IAE_unauth_access</i> .
------------------------------------	--

- ***instruction_VA_watchpoint*** [TT = 075₁₆] (Precise) — The virtual processor has detected that the Program Counter (PC) matches the VA Watchpoint register, when instruction VA watchpoints are enabled and the PC is being translated from a virtual address to a hardware address. If the PC is not being translated from a virtual address (for example, the PC is being treated as a hardware address), then an *instruction_VA_watchpoint* exception will not be generated, even if a match is detected between the VA Watchpoint register and the PC.
- ***interrupt_level_n*** [TT = 041₁₆–04F₁₆] (Disrupting) — SOFTINT{n} was set to 1 or an external interrupt request of level n was presented to the virtual processor and n > PIL.

Implementation Note	interrupt_level_14 can be caused by (1) setting SOFTINT{14} to 1, (2) occurrence of a "TICK match", or (3) occurrence of a "STICK match" (see <i>softintP Register (ASRs 20, 21, 22)</i> on page 55).
----------------------------	---
- ***LDDF_mem_address_not_aligned*** [TT = 035₁₆] (Precise) — An attempt was made to execute an LDDF or LDDFA instruction and the effective address was not doubleword aligned. (impl. dep. #109)
- ***mem_address_not_aligned*** [TT = 034₁₆] (Precise) — A load/store instruction generated a memory address that was not properly aligned according to the instruction, or a JMWL or RETURN instruction generated a non-word-aligned address. (See also *Special Memory Access ASIs* on page 432.)
- ***nonresumable_error*** [TT = 07F₁₆] (Disrupting) — This interrupt indicates that there is a valid entry in the nonresumable error queue. This interrupt is not generated by hardware, but is used by hyperprivileged software to inform privileged software that an error report has been appended to the nonresumable error queue.
- ***privileged_action*** [TT = 037₁₆] (Precise) — An action defined to be privileged has been attempted while in nonprivileged mode (PSTATE.priv = 0), or an action defined to be hyperprivileged has been attempted while in nonprivileged or privileged mode. Examples:
 - A data access by nonprivileged software using a restricted (privileged or hyperprivileged) ASI, that is, an ASI in the range 00₁₆ to 7F₁₆ (inclusively)
 - A data access by nonprivileged or privileged software using a hyperprivileged ASI, that is, an ASI in the range 30₁₆ to 7F₁₆ (inclusively)
 - Execution by nonprivileged software of an instruction with a privileged operand value
 - An attempt to read the TICK register by nonprivileged software when nonprivileged access to TICK is disabled .

- An attempt to execute a nonprivileged instruction with an operand value requiring more privilege than available in the current privilege mode.
- **privileged_opcode** [TT = 011₁₆] (Precise) — An attempt was made to execute a privileged instruction while in nonprivileged mode (PSTATE.priv = 0).
- **resumable_error** [TT = 07E₁₆] (Disrupting) — There is a valid entry in the resumable error queue. This interrupt is used to inform privileged software that an error report has been appended to the resumable error queue, and the current instruction stream is in a consistent state so that execution can be resumed after the error is handled.
- **spill_n_normal** [TT = 080₁₆–09F₁₆] (Precise)
- **spill_n_other** [TT = 0A0₁₆–0BF₁₆] (Precise)
A SAVE or FLUSHW instruction has determined that the contents of a register window must be saved to memory.
- **STDF_mem_address_not_aligned** [TT = 036₁₆] (Precise) — An attempt was made to execute an STDF or STDFA instruction and the effective address was not doubleword aligned. (impl. dep. #110)
- **tag_overflow** [TT = 023₁₆] (Precise) (deprecated (C2)) — A TADDccTV or TSUBccTV instruction was executed, and either 32-bit arithmetic overflow occurred or at least one of the tag bits of the operands was nonzero.
- **trap_instruction** [TT = 100₁₆–17F₁₆] (Precise) — A Tcc instruction was executed and the trap condition evaluated to TRUE, and the software trap number operand of the instruction is 127 or less.
- **unimplemented_LDTW** [TT = 012₁₆] (Precise) — An attempt was made to execute an LDTW instruction that is not implemented in hardware on this implementation (impl. dep. #107-V9).
- **unimplemented_STTW** [TT = 013₁₆] (Precise) — An attempt was made to execute an STTW instruction that is not implemented in hardware on this implementation (impl. dep. #108-V9).
- **VA_watchpoint** [TT = 062₁₆] (Precise) — The virtual processor has detected an attempt to access (load from or store to) a virtual address specified by the VA Watchpoint register, while VA watchpoints are enabled and the address is being translated from a virtual address to a hardware address. If the load or store address is not being translated from a virtual address (for example, the address is being treated as a real address), then a *VA_watchpoint* exception will not be generated even if a match is detected between the VA Watchpoint register and a load or store address.

12.7.1 SPARC V9 Traps Not Used in Oracle SPARC Architecture 2015

The following traps were optional in the SPARC V9 specification and are not used in Oracle SPARC Architecture 2015:

- **implementation_dependent_exception_n** [TT = 077₁₆ - 07B₁₆] This range of implementation-dependent exceptions has been replaced by a set of architecturally-defined exceptions. (impl. dep. #35-V8-Cs20)
- **LDQF_mem_address_not_aligned** [TT = 038₁₆] (Precise) — An attempt was made to execute an LDQF instruction and the effective address was word aligned but not quadword aligned. Use of this exception is implementation dependent (impl. dep. #111-V9-Cs10). A separate trap entry for this exception supports fast software emulation of the LDQF instruction when the effective address is word aligned but not quadword aligned. See *Load Floating-Point Register on page 246*. (impl. dep. #111)
- **STQF_mem_address_not_aligned** [TT = 039₁₆] (Precise) — An attempt was made to execute an STQF instruction and the effective address was word aligned but not quadword aligned. Use of this exception is implementation dependent (impl. dep. #112-V9-Cs10). A separate trap entry for the exception supports fast software emulation of the STQF instruction when the effective address is word aligned but not quadword aligned. See *Store Floating-Point on page 340*. (impl. dep. #112)

12.8 Register Window Traps

Window traps are used to manage overflow and underflow conditions in the register windows, support clean windows, and implement the FLUSHW instruction.

12.8.1 Window Spill and Fill Traps

A window overflow occurs when a SAVE instruction is executed and the next register window is occupied ($CANSAVE = 0$). An overflow causes a spill trap that allows privileged software to save the occupied register window in memory, thereby making it available for use.

A window underflow occurs when a RESTORE instruction is executed and the previous register window is not valid ($CANRESTORE = 0$). An underflow causes a fill trap that allows privileged software to load the registers from memory.

12.8.2 *clean_window* Trap

The virtual processor provides the *clean_window* trap so that system software can create a secure environment in which it is guaranteed that data cannot inadvertently leak through register windows from one software program to another.

A clean register window is one in which all of the registers, including uninitialized registers, contain either 0 or data assigned by software executing in the address space to which the window belongs. A clean window cannot contain register values from another process, that is, from software operating in a different address space.

Supervisor software specifies the number of windows that are clean with respect to the current address space in the CLEANWIN register. This number includes register windows that can be restored (the value in the CANRESTORE register) and the register windows following CWP that can be used without cleaning. Therefore, the number of clean windows available to be used by the SAVE instruction is

$$\text{CLEANWIN} - \text{CANRESTORE}$$

The SAVE instruction causes a *clean_window* exception if this value is 0. This behavior allows supervisor software to clean a register window before it is accessed by a user.

12.8.3 Vectoring of Fill/Spill Traps

To make handling of fill and spill traps efficient, the SPARC V9 architecture provides multiple trap vectors for the fill and spill traps. These trap vectors are determined as follows:

- Supervisor software can mark a set of contiguous register windows as belonging to an address space different from the current one. The count of these register windows is kept in the OTHERWIN register. A separate set of trap vectors (*fill_n_other* and *spill_n_other*) is provided for spill and fill traps for these register windows (as opposed to register windows that belong to the current address space).
- Supervisor software can specify the trap vectors for fill and spill traps by presetting the fields in the WSTATE register. This register contains two subfields, each three bits wide. The WSTATE.normal field determines one of eight spill (fill) vectors to be used when the register window to be spilled (filled) belongs to the current address space ($OTHERWIN = 0$). If the OTHERWIN register is nonzero, the WSTATE.other field selects one of eight *fill_n_other* (*spill_n_other*) trap vectors.

See *Trap-Table Entry Addresses* on page 450, for more details on how the trap address is determined.

12.8.4 CWP on Window Traps

On a window trap, the CWP is set to point to the window that must be accessed by the trap handler, as follows.

Note | All arithmetic on CWP is done **modulo** $N_REG_WINDOWS$.

- If the spill trap occurs because of a SAVE instruction (when $CANSAVE = 0$), there is an overlap window between the CWP and the next register window to be spilled:

$$CWP \leftarrow (CWP + 2) \bmod N_REG_WINDOWS$$

If the spill trap occurs because of a FLUSHW instruction, there can be unused windows ($CANSAVE$) in addition to the overlap window between the CWP and the window to be spilled:

$$CWP \leftarrow (CWP + CANSAVE + 2) \bmod N_REG_WINDOWS$$

Implementation | All spill traps can set CWP by using the calculation:

Note | $CWP \leftarrow (CWP + CANSAVE + 2) \bmod N_REG_WINDOWS$
since $CANSAVE$ is 0 whenever a trap occurs because of a SAVE instruction.

- On a fill trap, the window preceding CWP must be filled:

$$CWP \leftarrow (CWP - 1) \bmod N_REG_WINDOWS$$

- On a *clean_window* trap, the window following CWP must be cleaned. Then

$$CWP \leftarrow (CWP + 1) \bmod N_REG_WINDOWS$$

12.8.5 Window Trap Handlers

The trap handlers for fill, spill, and *clean_window* traps must handle the trap appropriately and return, by using the REPLY instruction, to reexecute the trapped instruction. The state of the register windows must be updated by the trap handler, and the relationships among CLEANWIN, CANSAVE, CANRESTORE, and OTHERWIN must remain consistent. Follow these recommendations:

- A spill trap handler should execute the SAVED instruction for each window that it spills.
- A fill trap handler should execute the RESTORED instruction for each window that it fills.
- A *clean_window* trap handler should increment CLEANWIN for each window that it cleans:

$$CLEANWIN \leftarrow (CLEANWIN + 1)$$

Interrupt Handling

Virtual processors and I/O devices can interrupt a selected virtual processor by assembling and sending an interrupt packet. The contents of the interrupt packet are defined by software convention. Thus, hardware interrupts and cross-calls can have the same hardware mechanism for interrupt delivery and share a common software interface for processing.

The interrupt mechanism is a two-step process:

- sending of an interrupt request (through an implementation-specific hardware mechanism) to an interrupt queue of the target virtual processor
- receipt of the interrupt request on the target virtual processor and scheduling software handling of the interrupt request

Privileged software running on a virtual processor can schedule interrupts to *itself* (typically, to process queued interrupts at a later time) by setting bits in the privileged SOFTINT register (see *Software Interrupt Register (softint)* on page 468).

Programming Note	An interrupt request packet is sent by an interrupt source and is received by the specified target in an interrupt queue. Upon receipt of an interrupt request packet, a special trap is invoked on the target virtual processor. The trap handler software invoked in the target virtual processor then schedules itself to later handle the interrupt request by posting an interrupt in the SOFTINT register at the desired interrupt level.
-------------------------	---

In the following sections, the following aspects of interrupt handling are described:

- **Interrupt Packets** on page 467.
- **Software Interrupt Register (softint)** on page 468.
- **Interrupt Queues** on page 468.
- **Interrupt Traps** on page 470.

13.1 Interrupt Packets

Each interrupt is accompanied by data, referred to as an “interrupt packet”. An interrupt packet is 64 bytes long, consisting of eight 64-bit doublewords. The contents of these data are defined by software convention.

13.2 Software Interrupt Register (SOFTINT)

To schedule interrupt vectors for processing at a later time, privileged software running on a virtual processor can send itself signals (interrupts) by setting bits in the privileged SOFTINT register.

See *softintP Register (ASRs 20, 21, 22)* on page 55 for a detailed description of the SOFTINT register.

Programming Note | The SOFTINT register (ASR 16₁₆) is used for communication from nucleus (privileged, TL > 0) software to privileged software running with TL = 0. Interrupt packets and other service requests can be scheduled in queues or mailboxes in memory by the nucleus, which then sets SOFTINT{n} to cause an interrupt at level *n*.

Programming Note | The SOFTINT mechanism is independent of the “mondo” interrupt mechanism mentioned in *Interrupt Queues* on page 468. The two mechanisms do not interact.

13.2.1 Setting the Software Interrupt Register

SOFTINT{n} is set to 1 by executing a WRSOFTINT_SET^P instruction (WRAsr using ASR 20) with a ‘1’ in bit *n* of the value written (bit *n* corresponds to interrupt level *n*). The value written to the SOFTINT_SET register is effectively **ored** into the SOFTINT register. This approach allows the interrupt handler to set one or more bits in the SOFTINT register with a single instruction.

See *softint_setP Pseudo-Register (ASR 20)* on page 55 for a detailed description of the SOFTINT_SET pseudo-register.

13.2.2 Clearing the Software Interrupt Register

When all interrupts scheduled for service at level *n* have been serviced, kernel software executes a WRSOFTINT_CLR^P instruction (WRAsr using ASR 21) with a ‘1’ in bit *n* of the value written, to clear interrupt level *n* (impl. dep. 34-V8a). The complement of the value written to the SOFTINT_CLR register is effectively **anded** with the SOFTINT register. This approach allows the interrupt handler to clear one or more bits in the SOFTINT register with a single instruction.

Programming Note | To avoid a race condition between operating system kernel software clearing an interrupt bit and nucleus software setting it, software should (again) examine the queue for any valid entries after clearing the interrupt bit.

See *softint_clrP Pseudo-Register (ASR 21)* on page 56 for a detailed description of the SOFTINT_CLR pseudo-register.

13.3 Interrupt Queues

Interrupts are indicated to privileged mode via circular interrupt queues, each with an associated trap vector. There are 4 interrupt queues, one for each of the following types of interrupts:

- Device mondos¹

- CPU mondos
- Resumable errors
- Nonresumable errors

New interrupt entries are appended to the tail of a queue and privileged software reads them from the head of the queue.

Programming Note | Software conventions for cooperative management of interrupt queues and the format of queue entries are specified in the separate *Hypervisor API Specification* document.

13.3.1 Interrupt Queue Registers

The active contents of each queue are delineated by a 64-bit head register and a 64-bit tail register.

The interrupt queue registers are accessed through ASI `ASI_QUEUE` (25_{16}). The ASI and address assignments for the interrupt queue registers are provided in TABLE 13-1.

TABLE 13-1 Interrupt Queue Register ASI Assignments

Register	ASI	Virtual Address	Privileged mode Access
CPU Mondo Queue Head	25_{16} (<code>ASI_QUEUE</code>)	$3C0_{16}$	RW
CPU Mondo Queue Tail	25_{16} (<code>ASI_QUEUE</code>)	$3C8_{16}$	R or RW†
Device Mondo Queue Head	25_{16} (<code>ASI_QUEUE</code>)	$3D0_{16}$	RW
Device Mondo Queue Tail	25_{16} (<code>ASI_QUEUE</code>)	$3D8_{16}$	R or RW†
Resumable Error Queue Head	25_{16} (<code>ASI_QUEUE</code>)	$3E0_{16}$	RW
Resumable Error Queue Tail	25_{16} (<code>ASI_QUEUE</code>)	$3E8_{16}$	R or RW†
Nonresumable Error Queue Head	25_{16} (<code>ASI_QUEUE</code>)	$3F0_{16}$	RW
Nonresumable Error Queue Tail	25_{16} (<code>ASI_QUEUE</code>)	$3F8_{16}$	R or RW†

† see **IMPL. DEP.#422-S10**

The status of each queue is reflected by its head and tail registers:

- A Queue Head Register indicates the location of the oldest interrupt packet in the queue
- A Queue Tail Register indicates the location where the next interrupt packet will be stored


An event that results in the insertion of a queue entry causes the tail register for that queue to refer to the following entry in the circular queue. Privileged code is responsible for updating the head register appropriately when it removes an entry from the queue.

A queue is *empty* when the contents of its head and tail registers are equal. A queue is *full* when the insertion of one more entry would cause the contents of its head and tail registers to become equal.

¹ “mundo” is a historical term, referring to the name of the original UltraSPARC 1 bus transaction in which these interrupts were introduced

Programming Note

By current convention, the format of a Queue Head or Tail register is as follows:



Under this convention:

- updating a Queue Head register involves incrementing it by 64 (size of a queue entry, in bytes)
- Queue Head and Tail registers are updated using modular arithmetic (modulo the size of the circular queue, in bytes)
- bits 5:0 always read as zeros, and attempts to write to them are ignored
- the maximum queue offset for an interrupt queue is implementation dependent
- behavior when a queue register is written with a value larger than the maximum queue offset (queue length minus the length of the last entry) is undefined

This is merely a convention and is subject to change.

13.4 Interrupt Traps

The following interrupt traps are defined in the Oracle SPARC Architecture 2015: *cpu_mondo*, *dev_mondo*, *resumable_error*, and *nonresumable_error*. See Chapter 12, *Traps*, for details.

Oracle SPARC Architecture 2015 also supports the *interrupt_level_n* traps defined in the SPARC V9 specification.pt trans

How interrupts are delivered is implementation-specific; see the relevant implementation-specific Supplement to this specification for details.

Memory Management

An Oracle SPARC Architecture Memory Management Unit (MMU) conforms to the requirements set forth in the *SPARC V9 Architecture Manual*. In particular, it supports a 64-bit virtual address space, simplified protection encoding, and multiple page sizes.

IMPL. DEP. # 451-S20: The width of the virtual address supported is implementation dependent. If fewer than 64 bits are supported, then unsupported bits must have the same value as the most significant supported bit. For example, if the model supports 48 virtual address bits, then bits 63:48 must have the same value as bit 47.

This appendix describes the Memory Management Unit, as observed by privileged software, in these sections:

- **Virtual Address Translation** on page 471.
- **Context ID** on page 474.
- **TSB Translation Table Entry (TTE)** on page 475.
- **Translation Storage Buffer (TSB)** on page 478.

14.1 Virtual Address Translation

The MMUs may support up to eight page sizes: 8 KBytes, 64 KBytes, 512 KBytes, 4 MBytes, 32 MBytes, 256 MBytes, 2 GBytes, and 16 GBytes and 1 TByte. 8 KByte, 64 KByte and 4 MByte page sizes must be supported; the other page sizes are optional.

IMPL. DEP. #310-U4: Which, if any, of the following optional page sizes are supported by the MMU in an Oracle SPARC Architecture 2015 implementation is implementation dependent: 512 KBytes, 32 MBytes, 256 MBytes, 2 GBytes, and 16 GBytes.

An Oracle SPARC Architecture MMU supports a 64-bit virtual address (VA) space.

IMPL. DEP. #452-S20: The number of real address (RA) bits supported is implementation dependent. A minimum of 40 bits and maximum of 56 bits can be provided for real addresses (RA). See implementation-specific documentation for details.

In each translation, the virtual page number is replaced by a physical page number, which is concatenated with the page offset to form the full hardware address, as illustrated in FIGURE 14-1 and FIGURE 14-2.

IMPL. DEP. #453-S20: It is implementation dependent whether there is a unified MMU (UMMU) or a separate IMMU (for instruction accesses) and DMMU (for data accesses). The Oracle SPARC Architecture supports both configurations.

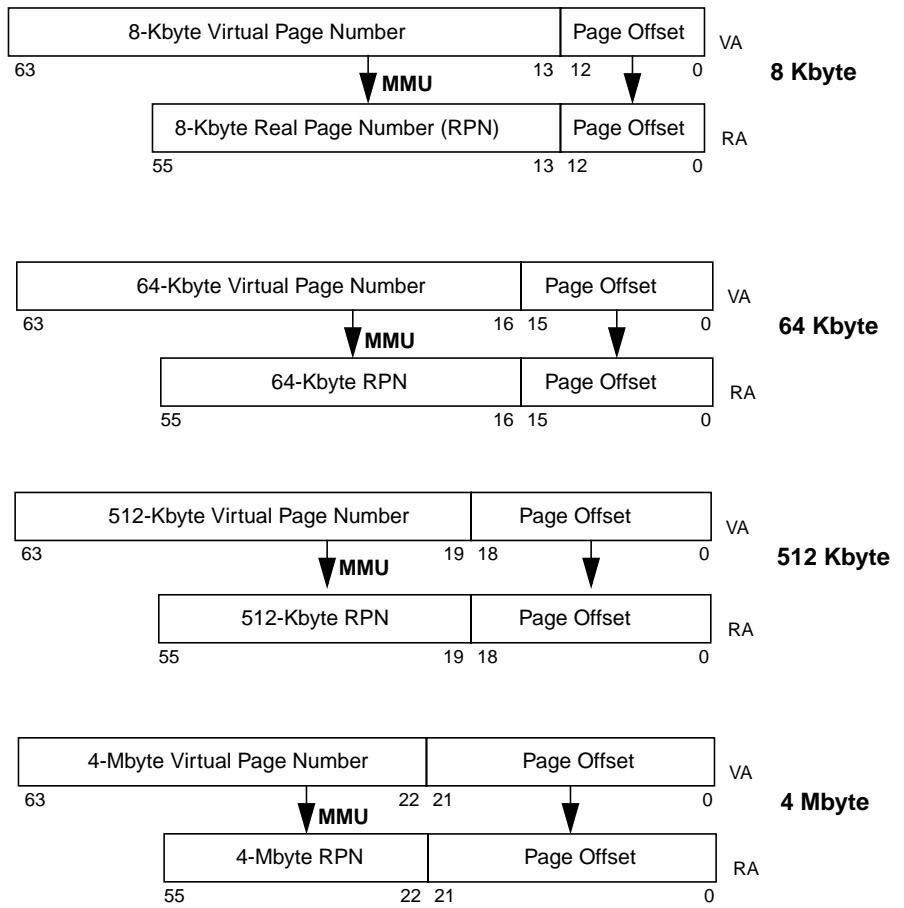


FIGURE 14-1 Virtual-to-Real Address Translation for 8-Kbyte, 64-Kbyte, 512-Kbyte, and 4-Mbyte Page Sizes




FIGURE 14-2 Virtual-to-Real Address Translation for 32-Mbyte, 256-Mbyte, 2-Gbyte, and 16-Gbyte Page Sizes

Privileged software manages virtual-to-real address translations.

Privileged software maintains translation information in an arbitrary data structure, called the *software translation table*.

The Translation Storage Buffer (TSB) is an array of Translation Table Entries which serves as a cache of the software translation table, used to quickly reload the TLB in the event of a TLB miss.

A conceptual view of privileged-mode memory management the MMU is shown in FIGURE 14-3. The software translation table is likely to be large and complex. The translation storage buffer (TSB), which acts like a direct-mapped cache, is the interface between the software translation table and the underlying memory management hardware. The TSB can be shared by all processes running on a virtual processor or can be process specific; the hardware does not require any particular scheme. There can be several TSBs.




FIGURE 14-3 Conceptual View of the MMU

14.2 Context ID

The MMU supports three contexts:

- Primary Context
- Secondary Context
- Nucleus Context (which has a fixed Context ID value of zero)

The context used for each access depends on the type of access, the ASI used, the current privilege mode, and the current trap level (TL). Details are provided in the following paragraphs and in TABLE 14-1.

For instruction fetch accesses, in nonprivileged and privileged mode when TL = 0 the Primary Context is used; when TL > 0, the Nucleus Context is used.

For data accesses using *implicit* ASIs, in nonprivileged and privileged mode when TL = 0 the Primary Context is used; when TL > 0, the Nucleus Context is used.

For data accesses using *explicit* ASIs:

- In nonprivileged mode the Primary Context is used for the ASI_PRIMARY* ASIs, and the Secondary Context is used for the ASI_SECONDARY* ASIs.
- In privileged mode, the Primary Context is used for the ASI_PRIMARY* and the ASI_AS_IF_USER_PRIMARY* ASIs, the Secondary Context is used for the ASI_SECONDARY* and the ASI_AS_IF_USER_SECONDARY* ASIs, and the Nucleus Context is used for ASI_NUCLEUS* ASIs.

The above paragraphs are summarized in TABLE 14-1.

TABLE 14-1 Context Usage

Access Type	Privilege Mode	Under What Conditions each Context is Used		
		Primary Context	Secondary Context	Nucleus Context
Instruction Access	Nonprivileged or Privileged	(when TL = 0)	— †	(when TL > 0)
Data access using <i>implicit</i>	Nonprivileged or Privileged	(when TL = 0)	— †	(when TL > 0)
ASI	Nonprivileged	ASI_PRIMARY*	ASI_SECONDARY*	†
Data access using <i>explicit</i> ASI	Privileged	ASI_PRIMARY* ASI_AS_IF_USER_PRIMARY*	ASI_SECONDARY* ASI_AS_IF_USER_SECONDARY*	ASI_NUCLEUS*

†, no context is listed because this case cannot occur

Note | The Oracle SPARC Architecture provides the capability of private and shared contexts. Multiple primary and secondary context IDs, which allow different processes to share TTEs, are defined. See *Context ID Registers* on page 482 for details.

Programming Note | Privileged software (operating systems) intended to be portable across all Oracle SPARC Architecture implementations should always ensure that, for memory accesses made in privileged mode, private and shared context IDs are set to the same value. The exception to this is privileged-mode accesses using the ASI_AS_IF_USER* ASIs, which remain portable even if the private and shared context IDs differ.

IMPL. DEP. #___: The Oracle SPARC Architecture defines a 16-bit context ID. The size of the context ID field is implementation dependent. At least 13 bits must be implemented. If fewer than 16 bits are supported, the unused high order bits are ignored on writes to the context ID, and read as zeros.

14.3 TSB Translation Table Entry (TTE)

Each Translation Table Entry (TTE) in a Translation Storage Buffer (TSB) is the equivalent of a page table entry as defined in the *Sun4v Architecture Specification*; it holds information for a single page mapping. The TTE is divided into two 64-bit words representing the *tag* and *data* of the translation. Just as in a hardware cache, the tag is used to determine whether there is a hit in the TSB; if there is a hit, the data are used by either the hardware tablewalker or privileged software.

The TTE configuration is illustrated in FIGURE 14-4 and described in TABLE 14-2.




FIGURE 14-4 Translation Storage Buffer (TSB) Translation Table Entry (TTE)

TABLE 14-2 TSB TTE Bit Description (1 of 3)

Bit	Field	Description
Tag– 63:48	context_id	The 16-bit context ID associated with the TTE.
Tag– 47:42	—	These bits must be zero for a tag match.
Tag– 41:0	va	Bits 63:22 of the Virtual Address (the virtual page number). Bits 21:13 of the VA are not maintained because these bits index the minimally sized, direct-mapped TSBs.
Data – 63	v	Valid. If $v = 1$, then the remaining fields of the TTE are meaningful, and the TTE can be used; otherwise, the TTE cannot be used to translate a virtual address. <div style="border-left: 1px solid black; padding-left: 10px;"> Programming Note The explicit Valid bit is (intentionally) redundant with the software convention of encoding an invalid TTE with an unused context ID. The encoding of the <code>context_id</code> field is necessary to cause a failure in the TTE tag comparison, while the explicit Valid bit in the TTE data simplifies the TTE miss handler. </div>
Data – 62	nfo	No Fault Only. If $nfo = 1$, loads with <code>ASI_PRIMARY_NO_FAULT{ _LITTLE }</code> or <code>ASI_SECONDARY_NO_FAULT{ _LITTLE }</code> are translated. Any other data access with the D/UMMU TTE. $nfo = 1$ will trap with a <code>DAE_nfo_page</code> exception. An instruction fetch access to a page with the IMMU TTE. $nfo = 1$ results in an <code>IAE_nfo_page</code> exception.
Data –61:56	soft2	Software-defined field, provided for use by the operating system. The <code>soft2</code> field can be written with any value in the TSB. Hardware is not required to maintain this field in any TLB (or uTLB), so when it is read from the TLB (uTLB), it may read as zero.
Data – 55:13	taddr	Target address; the underlying address (Real Address {55:13}) to which the MMU will map the page. IMPL. DEP. # 238-U3: When page offset bits for larger page sizes are stored in the TLB, it is implementation dependent whether the data returned from those fields by a Data Access read is zero or the data previously written to them.

TABLE 14-2 TSB TTE Bit Description (2 of 3)

Bit	Field	Description											
Data – 12	ie	<p>Invert Endianness. If <code>ie = 1</code> for a page, accesses to the page are processed with inverse endianness from that specified by the instruction (big for little, little for big). See page 479 for details.</p> <p>Programming Notes</p> <p>(1) The primary purpose of this bit is to aid in the mapping of I/O devices (through <i>noncacheable</i> memory addresses) whose registers contain and expect data in little-endian format. Setting <code>TTE.ie = 1</code> allows those registers to be accessed correctly by big-endian programs using ordinary loads and stores, such as those typically issued by compilers; otherwise little-endian loads and stores would have to be issued by hand-written assembler code.</p> <p>(2) This bit can also be used when mapping <i>cacheable</i> memory. However, cacheable accesses to pages marked with <code>TTE.ie = 1</code> may be slower than accesses to the page with <code>TTE.ie = 0</code>. For example, an access to a cacheable page with <code>TTE.ie = 1</code> may perform as if there was a miss in the first-level data cache.</p> <p>Implementation Note</p> <p>Some implementations may require cacheable accesses to pages tagged with <code>TTE.ie = 1</code> to bypass the data cache, adding latency to those accesses.</p> <p>IMPL. DEP. #_: The <code>ie</code> bit in the IMMU is ignored during ITLB operation. It is implementation dependent if it is implemented and how it is read and written.</p>											
Data – 11	e	<p>Side effect. If the side-effect bit is set to 1, loads with <code>ASI_PRIMARY_NO_FAULT</code>, <code>ASI_SECONDARY_NO_FAULT</code>, and their <code>*_LITTLE</code> variations will trap for addresses within the page, noncacheable memory accesses other than block loads and stores are strongly ordered against other <code>e</code>-bit accesses, and noncacheable stores are not merged. This bit should be set to 1 for pages that map I/O devices having side effects. Note, also, that the <code>e</code> bit causes the prefetch instruction to be treated as a nop, but does not prevent normal (hardware) instruction prefetching.</p> <p>Note 1: The <code>e</code> bit does not force a noncacheable access. It is expected, but not required, that the <code>cp</code> bit will be set to 0 when the <code>e</code> bit is set to 1. If the <code>cp</code> bit is set to 1 along with the <code>e</code> bit, the result is undefined.</p> <p>Note 2: The <code>e</code> bit and the <code>nfo</code> bit are mutually exclusive; both bits should never be set to 1 in any TTE.</p>											
Data – 10	cp	<p>The cacheable-in-physically-indexed-cache bit determines the placement of data in caches, according to the following table. The MMU's operation is not affected by the <code>cp</code> bit; the MMU merely passes it through to the cache subsystem:</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th rowspan="2">Cacheable (cp)</th> <th colspan="2">Meaning of TTE when placed in:</th> </tr> <tr> <th>I-TLB (Instruction Cache PA-indexed)</th> <th>D-TLB (Data Cache PA-indexed)</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Cacheable in L2 and L3 caches only</td> <td>Cacheable in L2 and L3 caches only</td> </tr> <tr> <td>1</td> <td>Cacheable in L3 cache, L2 cache, and L1 I-cache</td> <td>Cacheable in L3 cache, L2 cache, and L1 D-cache</td> </tr> </tbody> </table>	Cacheable (cp)	Meaning of TTE when placed in:		I-TLB (Instruction Cache PA-indexed)	D-TLB (Data Cache PA-indexed)	0	Cacheable in L2 and L3 caches only	Cacheable in L2 and L3 caches only	1	Cacheable in L3 cache, L2 cache, and L1 I-cache	Cacheable in L3 cache, L2 cache, and L1 D-cache
Cacheable (cp)	Meaning of TTE when placed in:												
	I-TLB (Instruction Cache PA-indexed)	D-TLB (Data Cache PA-indexed)											
0	Cacheable in L2 and L3 caches only	Cacheable in L2 and L3 caches only											
1	Cacheable in L3 cache, L2 cache, and L1 I-cache	Cacheable in L3 cache, L2 cache, and L1 D-cache											
Data – 9	—												

TABLE 14-2 TSB TTE Bit Description (3 of 3)

Bit	Field	Description																								
Data – 8	p	Privileged. If p = 1, only privileged software can access the page mapped by the TTE. If p = 1 and an access to the page is attempted by nonprivileged mode (PSTATE.priv = 0), then the MMU signals an <i>IAE_privilege_violation</i> exception or <i>DAE_privilege_violation</i> exception.																								
Data – 7	ep	Executable. If ep = 1, the page mapped by this TTE has execute permission granted. Instructions may be fetched and executed from this page. If ep = 0, an attempt to execute an instruction from this page results in an <i>IAE_unauth_access</i> exception. IMPL. DEP. #___: An Oracle SPARC Architecture ITLB implementation may elect to not implement the ep bit, and instead present the <i>IAE_unauth_access</i> exception if there is an attempt to load an ITLB entry with ep = 0 during a hardware tablewalk. In this case, the MMU miss trap handler software must also detect the ep = 0 case when the IMMU miss is handled by software.																								
Data – 6	w	IMPL. DEP. # Writable. If w = 1, the page mapped by this TTE has write permission granted. Otherwise, write permission is not granted																								
Data – 5:4	soft	Software-defined field, provided for use by the operating system. The soft field can be written with any value in the TSB. Hardware is not required to maintain this field in any TLB (or uTLB), so when it is read from the TLB (or uTLB), it may read as zero.																								
Data – 3:0	sz	The page size of this entry, encoded as shown below. <table border="1"> <thead> <tr> <th>sz</th> <th>Page Size</th> </tr> </thead> <tbody> <tr> <td>0000</td> <td>8 Kbytes</td> </tr> <tr> <td>0001</td> <td>64 Kbytes</td> </tr> <tr> <td>0010</td> <td>512 Kbytes</td> </tr> <tr> <td>0011</td> <td>4 Mbytes</td> </tr> <tr> <td>0100</td> <td>32 Mbytes</td> </tr> <tr> <td>0101</td> <td>256 Mbytes</td> </tr> <tr> <td>0110</td> <td>2 Gbytes</td> </tr> <tr> <td>0111</td> <td>16 Gbytes</td> </tr> <tr> <td>1000</td> <td>128 Gbytes</td> </tr> <tr> <td>1001</td> <td>Reserved for 1 Tbyte</td> </tr> <tr> <td>1010-1111</td> <td>Reserved</td> </tr> </tbody> </table>	sz	Page Size	0000	8 Kbytes	0001	64 Kbytes	0010	512 Kbytes	0011	4 Mbytes	0100	32 Mbytes	0101	256 Mbytes	0110	2 Gbytes	0111	16 Gbytes	1000	128 Gbytes	1001	Reserved for 1 Tbyte	1010-1111	Reserved
sz	Page Size																									
0000	8 Kbytes																									
0001	64 Kbytes																									
0010	512 Kbytes																									
0011	4 Mbytes																									
0100	32 Mbytes																									
0101	256 Mbytes																									
0110	2 Gbytes																									
0111	16 Gbytes																									
1000	128 Gbytes																									
1001	Reserved for 1 Tbyte																									
1010-1111	Reserved																									

Note that page size encodings follow the formula:

$$\text{Page Size} = 2^{((sz \times 3) + 13)} \text{ bytes}$$

14.4 Translation Storage Buffer (TSB)

The Translation Storage Buffer (TSB) is an array of Translation Table Entries managed entirely by privileged software. It serves as a cache of the software translation table, used to quickly reload the hardware translation table (TLB) in the event of a TLB miss.

14.4.1 TSB Cacheability and Consistency

The TSB exists as a data structure in memory and therefore can be cached. Indeed, the speed of the TLB miss handler relies on the TSB accesses hitting the level-2 cache at a substantial rate. This policy may result in some conflicts with normal instruction and data accesses, but the dynamic sharing of the level-2 cache resource will provide a better overall solution than that provided by a fixed partitioning.

Programming Note When software updates the TSB, it is responsible for ensuring that the store(s) used to perform the update are made visible in the memory system (for access by subsequent loads, stores, and load-stores) by use of an appropriate MEMBAR instruction.

Making a TSB update visible to fetches of instructions subsequent to the store(s) that updated the TSB may require execution of instructions such as FLUSH, DONE, or RETRY, in addition to the MEMBAR.

14.4.2 TSB Organization

The TSB is arranged as a direct-mapped cache of TTEs.

In each case, n least significant bits of the respective virtual page number are used as the offset from the TSB base address, with n equal to log base 2 of the number of TTEs in the TSB.

The TSB organization is illustrated in FIGURE 14-5. The constant n can range from 512 to an implementation-dependent number.




FIGURE 14-5 TSB Organization

IMPL. DEP. #227-U3: The maximum number of entries in a TSB is implementation-dependent in the Oracle SPARC Architecture (to a maximum of 16 million).

14.5 ASI Value, Context ID, and Endianness Selection for Translation

The selection of the context ID for a translation is the result of a two-step process:

1. The ASI is determined (conceptually by the Integer Unit) from the instruction, ASI register, trap level, the virtual processor endian mode (PSTATE.cle), and the privilege level (PSTATE.priv).
2. The context ID is determined directly from the ASI. The context ID value is read by the context ID selected by the ASI.

The ASI value and endianness (little or big) are determined, according to TABLE 14-3 through TABLE 14-4.

When using the Primary Context ID, the values stored in the Primary Context IDs are used by the Data (or Unified) MMU. The Secondary Context ID is never used for instruction accesses.

The endianness of a data access is specified by three conditions:

- The ASI specified in the opcode or ASI register
- The PSTATE current little-endian bit (.cle)
- The TTE “invert endianness” bit (ie). The TTEbit inverts the endianness that is otherwise specified for the access.

Note The D/UMMU ie bit inverts the endianness for all accesses, including alternate space loads, stores, and atomic load-stores that specify an ASI. For example,
`ldxa [%g1]#ASI_PRIMARY_LITTLE`
 will be big-endian if the ie bit = 1.
 Accesses to ASIs which are not translated by the MMU (nontranslating ASIs) are not affected by the TTE.ie bit.

TABLE 14-3 ASI Mapping for Instruction Access

Mode	TL	PSTATE .cle	Endianness	ASI Used	Resulting Address Type
Nonprivileged	0	—	Big	ASI_PRIMARY	VA
Privileged	0	—	Big	ASI_PRIMARY	VA
	1-2	—	Big	ASI_NUCLEUS	VA

TABLE 14-4 ASI Mapping for Data Accesses (1 of 2)

Access Type	Privilege Mode	TL	PSTATE .cle	TTE.ie	Endianness	ASI Used	Resulting Address Type	
Load, Store, Atomic Load-Store, or Prefetch with implicit ASI	NP	0 ¹	0	0	Big	ASI_PRIMARY	VA	
				1	Little			
		0 ¹	1	0	Little	ASI_PRIMARY_LITTLE	VA	
				1	Big			
	P	0	0	0	0	Big	ASI_PRIMARY	VA
					1	Little		
			0	1	0	Little	ASI_PRIMARY_LITTLE	VA
					1	Big		
		1-2 ¹	0	0	Big	ASI_NUCLEUS	VA	
				1	Little			
		1-2 ¹	1	0	Little	ASI_NUCLEUS_LITTLE	VA	
				1	Big			

TABLE 14-4 ASI Mapping for Data Accesses (2 of 2)

Access Type	Privilege Mode	TL	PSTATE .cle	TTE.ie	Endianness	ASI Used	Resulting Address Type
Load, Store, Atomic Load-Store, or Prefetch alternate with ASI name <i>not</i> ending in <code>_LITTLE</code>	NP	0 ¹	any	0	Big ²	Explicitly specified in instruction	VA
				1	Little ¹		
	P	0-2 ¹	any	0	Big ¹	Explicitly specified in instruction	VA
				1	Little ¹		
		0-2 ¹	any	0	Big	ASI_*REAL* ASI	RA
				1	Little		
	0-2 ¹	any	any	Big	Nontranslating ASIs	—	
	Load, Store, Atomic Load-Store, or Prefetch alternate with ASI name ending in <code>_LITTLE</code>	NP	0 ¹	any	0	Little	Explicitly specified in instruction
1					Big		
P		0-2 ¹	any	0	Little	Explicitly specified in instruction)	VA
				1	Big		
		0-2 ¹	any	0	Little	ASI_*REAL* ASI	RA
				1	Big		

1. MAX_TL = 2 for Oracle SPARC Architecture 2015 processors. Privilege mode operation is valid only for TL = 0, 1 or 2. Nonprivileged mode operation is valid only for TL = 0. See section 5.6.7 for details.

2. Accesses to nontranslating ASIs are always made in big endian mode, regardless of the setting of TTE.ie. See *ASI Values* on page 423 for information about nontranslating ASIs.

The Context ID used by the data and instruction MMUs is determined according to TABLE 14-5. The Context ID selection is not affected by the endianness of the access. For a comprehensive list of ASI values in the ASI map, see Chapter 10, *Address Space Identifiers (ASIs)*.

TABLE 14-5 IMMU, DMMU and UMMU Context ID Usage

ASI Value	Context ID Register
ASI_*NUCLEUS* (any ASI name containing the string "NUCLEUS")	Nucleus (0000 ₁₆ , hard-wired)
ASI_*PRIMARY* (any ASI name containing the string "PRIMARY")	All Primary Context IDs
ASI_*SECONDARY* (any ASI name containing the string "SECONDARY")	All Secondary Context IDs
All other ASI values	(Not applicable; no translation)

14.6 SPARC V9 "MMU Attributes"

The Oracle SPARC Architecture MMU complies completely with the SPARC V9 "MMU Attributes" as described in Appendix F.3.2.

With regard to Read, Write and Execute Permissions, SPARC V9 says "An MMU may allow zero or more of read, write and execute permissions, on a per-mapping basis. Read permission is necessary for data read accesses and atomic accesses. Write permission is necessary for data write accesses and atomic accesses. Execute permission is necessary for instruction accesses. At a minimum, an MMU must allow for 'all permissions', 'no permissions', and 'no write permission'; optionally, it can provide 'execute only' and 'write only', or any combination of 'read/write/execute' permissions."

TABLE 14-6 shows how various protection modes can be achieved, if necessary, through the presence or absence of a translation in the instruction or data MMU. Note that this behavior requires specialized TLB-miss handler code to guarantee these conditions.

TABLE 14-6 MMU SPARC V9 Appendix F.3.2 Protection Mode Compliance

Condition					Resultant Protection Mode
TTE in DMMU	TTE in IMMU	TTE in UMMU	ep Bit	Writable Attribute Bit	
Yes	No	Yes	0	0	Read-only ¹
No	Yes	N/A	1	N/A	Execute-only ¹
Yes	No	Yes	0	1	Read/Write ¹
Yes	Yes	Yes	1	0	Read-only/Execute
Yes	Yes	Yes	1	1	Read/Write/Execute
No	No	No	N/A	N/A	No Access

1. These protection modes are optional, according to SPARC V9.

14.6.1 Accessing MMU Registers

All internal MMU registers can be accessed directly by the virtual processor through defined ASIs, using LDXA and STXA instructions. Oracle SPARC Architecture-compatible processors do not require a MEMBAR #Sync, FLUSH, DONE, or RETRY instruction after a store to an MMU register for proper operation.

TABLE 14-7 lists the MMU registers and provides references to sections with more details.

TABLE 14-7 MMU Internal Registers and ASI Operations

IMMU ASI	D/UMMU ASI	VA{63:0}	Access	Register or Operation Name
21 ₁₆		8 ₁₆	RW	Primary Context ID 0 register
—	21 ₁₆	10 ₁₆	RW	Secondary Context ID 0 register
21 ₁₆		108 ₁₆	RW	Primary Context ID 1 register
—	21 ₁₆	110 ₁₆	RW	Secondary Context ID 1 register

14.6.2 Context ID Registers

The MMU architecture supports multiple primary and secondary context IDs. The address assignment of the context IDs is shown in TABLE 14-8.

TABLE 14-8 Context ID ASI Assignments

Register	ASI	Virtual Address
Primary Context ID 0	21 ₁₆	008 ₁₆
Primary Context ID 1	21 ₁₆	108 ₁₆
Secondary Context ID 0	21 ₁₆	010 ₁₆
Secondary Context ID 1	21 ₁₆	110 ₁₆

Programming Note For platforms that implement more than one primary context ID and one secondary context ID, privileged code must ensure that no more than one page translation is allowed to match at any time. An illustration of erroneous behavior is as follows:

1. An operating system constructs a mapping for virtual address *A* valid for context ID *P*;
2. it then constructs a mapping for address *A* for context ID *Q*.

By setting Primary Context ID 0 to *P* and Primary Context ID 1 to *Q*, both mappings would be active simultaneously, with conflicting translations for address *A*. Care must be taken not to construct such scenarios.

Oracle SPARC Architecture processors must prevent errors or data corruption due to multiple valid translations for a given virtual address using different contexts. TLBs may need to detect this scenario as a multiple tag hit error and cause an exception for such an access.

The Oracle SPARC Architecture supports up to two primary context IDs and two secondary context IDs, which are shared by the IMMU and D/UMMU. Primary Context ID 0 and Primary Context ID 1 are the primary context IDs, and a TLB entry for a translating primary ASI can match the contextid field with either Primary Context ID 0 or Primary Context ID 1 to produce a TLB hit.

Secondary Context ID 0 and Secondary Context ID 1 are the Secondary Context IDs, and a TLB entry for a translating secondary ASI can match the contextid field with either Secondary Context ID 0 or Secondary Context ID 1 to produce a TLB hit.

The Primary Context ID 0 and Primary Context ID 1 registers are illustrated in FIGURE 14-6, where pcontext is the context ID for the primary address space.




FIGURE 14-6 IMMU, DMMU, and UMMU Primary Context ID 0 and 1

The Secondary Context ID 0 and Secondary Context ID 1 registers is are illustrated in FIGURE 14-7, where scontextid is the context ID for the secondary address space.




FIGURE 14-7 D/UMMU Secondary Context ID 0 and 1

The Nucleus Context ID register is hardwired to zero, as illustrated in FIGURE 14-6.




FIGURE 14-8 IMMU, DMMU, and UMMU Nucleus Context ID

IMPL. DEP. #FIGURE 14-10415-S10: The size of context ID fields in MMU context registers is implementation-dependent and may range from 13 to 16 bits.

Opcode Maps

This appendix contains the Oracle SPARC Architecture 2015 instruction opcode maps.

In this appendix and in Chapter 7, *Instructions*, certain opcodes are marked with mnemonic superscripts. These superscripts and their meanings are defined in TABLE 7-1 on page 96. For preferred substitute instructions for deprecated opcodes, see the individual opcodes in Chapter 7 that are labeled “Deprecated”.

In the tables in this appendix, *reserved* (—) and shaded entries (as defined below) indicate opcodes that are not implemented in Oracle SPARC Architecture 2015 strands.

Shading	Meaning
	An attempt to execute opcode will cause an <i>illegal_instruction</i> exception.

An attempt to execute a reserved opcode behaves as defined in *Reserved Opcodes and Instruction Fields* on page 93.

TABLE A-1 op{1:0}

op {1:0}			
0	1	2	3
Branches and SETHI (See TABLE A-2)	CALL	Arithmetic & Miscellaneous (See TABLE A-3)	Loads/Stores (See TABLE A-4)

TABLE A-2 op2{2:0} (op = 0)

op2 {2:0}							
0	1	2	3	4	5	6	7
ILLTRAP (bits 29:25 = 0)	BPcc (See TABLE A-7)	Bicc ^D (See TABLE A-7)	BPr (bit 28 = 0) (See TABLE A-8)	SETHI, NOP ¹	FBPfcc (See TABLE A-7)	FBfcc ^D (See TABLE A-7)	—
— (bits 29:25 ≠ 0)			CBcond (bit 28 = 1) ¹				—

1. rd = 0, imm22 = 0

TABLE A-3 op3{5:0} (op = 10₂) (1 of 2)

		op3{5:4}			
		0	1	2	3
op3 {3:0}	0	ADD	ADDcc	TADDcc	WRY ^D (rd = 0) — (rd = 1) WRCCR (rd = 2) WRASI (rd = 3) — (rd = 4, 5) — (rd = 15, rs1 = 0, i = 1) — (rd = 15) and (rs1 ≠ 0 or i ≠ 1) — (rd = 7 – 14) WRFPRS (rd = 6) WRasr ^{PASR} (7 ≤ rd ≤ 14) — (rd = 16) — (rd = 17) — (rd = 18) WRGSR (rd = 19) WRSOFTINT_SET ^P (rd = 20) WRSOFTINT_CLR ^P (rd = 21) WRSOFTINT ^P (rd = 22) WRSTICK_CMPR ^P (rd = 25) — (rd = 26) WRPAUSE (rd = 27) WRMWAIT / MWAIT (rd = 28) — (rd = 29 - 31)
	1	AND	ANDcc	TSUBcc	SAVED ^P (fcn = 0) RESTORED ^P (fcn = 1) ALLCLEAN ^P (fcn = 2) OTHERW ^P (fcn = 3) NORMALW ^P (fcn = 4) INVALIDW ^P (fcn = 5) — (fcn ≥ 6)
	2	OR	ORcc	TADDccTV ^D	WRPR ^P (rd = 0-14 or 16) — (rd = 15 or 17-31)
	3	XOR	XORcc	TSUBccTV ^D	—
	4	SUB	SUBcc	MULSc ^D	FPop1 (see TABLE A-5)
	5	ANDN	ANDNcc	SLL (x = 0), SLLX (x = 1)	FPop2 (see TABLE A-6)
	6	ORN	ORNcc	SRL (x = 0), SRLX (x = 1)	(VIS) (See TABLE A-12)
	7	XNOR	XNORcc	SRA (x = 0), SRAX (x = 1)	

TABLE A-3 op3{5:0} (op = 10₂) (2 of 2)

		op3{5:4}			
		0	1	2	3
op3 {3:0}	8	ADDC	ADDC _{cc}	RDY ^D (rs1 = 0, i = 0) — (rs1 = 1, i = 0) RDCCR (rs1 = 2, i = 0) RDASI (rs1 = 3, i = 0) RDTICK ^{Pnpt} (rs1 = 4, i = 0) RDPC (rs1 = 5, i = 0) RDFPRS (rs1 = 6, i = 0) RDasr ^{PASR} (7 ≤ rd ≤ 14, i = 0) MEMBAR (rs1 = 15, rd = 0, i = 1, instruction bit 12 = 0) — (rs1 = 15, rd = 0, i = 1, instruction bit 12 = 1) — (i = 1, (rs1 ≠ 15 or rd ≠ 0)) — (rs1 = 15, rd = 0, i = 0) — (rs1 = 15 and rd > 0 and i = 0) — (rs1 = 16 and i = 0) — (rs1 = 17 and i = 0) — (rs1 = 18 and i = 0) RDGSR (rs1 = 19 and i = 0) — (rs1 = 20 or 21) and (i = 0)) RDSOFTINT ^P (rs1 = 22 and i = 0) RDSTICK (rs1 = 24 and i = 0) RDSTICK_CMPR ^P (rs1 = 25 and i = 0) RDCFR (rs1 = 26 and i = 0) — ((rs1 = 27 – 31) and (i = 0))	JMPL
	9	MULX	— 4-operand AES, DES, and Camelia	—	RETURN
	A	UMUL ^D	UMUL _{cc} ^D	RDPR ^P (rs1 = 1–14 or 16) — (rs1 = 15 or 17 – 31)	Tcc ((i = 0 and inst{10:5} = 0) or ((i = 1) and (inst{10:8} = 0))) (See TABLE A-7) — ((i = 0 and (inst{10:5} ≠ 0)) or (i = 1 and (inst{10:8} ≠ 0)) — (bit 29 = 1)
op3 {3:0}	B	SMUL ^D	SMUL _{cc} ^D	FLUSHW	FLUSH
	C	SUBC	SUBC _{cc}	MOV _{cc}	SAVE
	D	UDIVX	—	SDIVX	RESTORE
	E	UDIV ^D	UDIV _{cc} ^D	POPC (rs1 = 0) — (rs1 = 1,2,3) — (rs1 > 3)	DONE ^P (fcn = 0) RETRY ^P (fcn = 1) — (fcn = 2..14) — (fcn = 15) — (fcn = 16..31)
	F	SDIV ^D	SDIV _{cc} ^D	MOV _r (See TABLE A-8)	Reserved

TABLE A-4 op3{5:0} (op = 11₂)

		op3{5:4}			
		0	1	2	3
op3 {3:0}	0	LDUW	LDUWA ^{PASI}	LDF (i=1) or (i=0 and bits 12:10=0) — (i=0 and bits 12:10 ≠0)	LDF ^{PASI}
	1	LDUB	LDUBA ^{PASI}	(rd = 0) LDFSR ^D (rd = 1) LDXFSR (rd = 3) LDXEFSR — (rd = 2) — (rd ≥ 4)	Reserved
	2	LDUH	LDUHA ^{PASI}	LDQF	LDQFA ^{PASI}
	3	LDTW ^D — (rd odd)	LDTWA ^{D, PASI} LDTXA ^N — (rd odd)	LDDF (i=1) or (i=0 and bits 12:10=0) — (i=0 and bits 12:10 ≠0)	LDDFA ^{PASI} LDBLOCKF ^D LDSHORTF
	4	STW	STWA ^{PASI}	STF (i=1) or (i=0 and bits 12:10=0) — (i=0 and bits 12:10 ≠0)	STFA ^{PASI}
	5	STB	STBA ^{PASI}	(rd = 0) STFSR ^D (rd = 1) STXFSR — (rd > 1)	Reserved
	6	STH	STHA ^{PASI}	STQF	STQFA ^{PASI}
	7	STTW ^D — (rd odd)	STTWA ^{PASI} — (rd odd)	STDF (i=1) or (i=0 and bits 12:10=0) — (i=0 and bits 12:10 ≠0)	STDFA ^{PASI} STBLOCKF ^D STPARTIALF STSHORTF
	8	LDSW	LDSWA ^{PASI}	Reserved	
	9	LDSB	LDSBA ^{PASI}		
	A	LDSH	LDSHA ^{PASI}		
	B	LDX	LDXA ^{PASI}		
	C	Reserved	Reserved		
	D	LDSTUB	LDSTUBA ^{PASI}	PREFETCH (i=1) or (i=0 and bit 12=0) — (i=0 and bit 12=1) — (fcn = 5 – 15)	PREFETCHA ^{PASI} — (fcn = 5 – 15)
	E	STX	STXA ^{PASI}	Reserved	CASXA ^{PASI}
	F	SWAP ^D	SWAPA ^{D, PASI}	Reserved	Reserved

TABLE A-5 opf{8:0} (op = 10₂, op3 = 34₁₆ = FPop1)

opf{8:4}	opf{3:0}							
	0	1	2	3	4	5	6	7
00 ₁₆	—	FMOV _s	FMOV _d	FMOV _q	—	FNEG _s	FNEG _d	FNEG _q
01 ₁₆	—	—	—	—	—	—	—	—
02 ₁₆	—	—	—	—	—	—	—	—
03 ₁₆	—	—	—	—	—	—	—	—
04 ₁₆	—	FADD _s	FADD _d	FADD _q	—	FSUB _s	FSUB _d	FSUB _q
05 ₁₆	—	FNADD _s	FNADD _d	—	—	—	—	—
06 ₁₆	—	FHADD _s	FHADD _d	—	—	FHSUB _s	FHSUB _d	—
07 ₁₆	—	FNHADD _s	FNHADD _d	—	—	—	—	—
08 ₁₆	—	FsTO _x	FdTO _x	FqTO _x	FxTO _s	—	—	—
09 ₁₆	—	—	—	—	—	—	—	—
0A ₁₆	—	—	—	—	—	—	—	—
0B ₁₆	—	—	—	—	—	—	—	—
0C ₁₆	—	—	—	—	FiTO _s	—	FdTO _s	FqTO _s
0D ₁₆	—	FsTO _i	FdTO _i	FqTO _i	—	—	—	—
0E ₁₆ –1F ₁₆	—	—	—	—	—	—	—	—
	8	9	A	B	C	D	E	F
00 ₁₆	—	FABS _s	FABS _d	FABS _q	—	—	—	—
01 ₁₆	—	—	—	—	—	—	—	—
02 ₁₆	—	FSQRT _s	FSQRT _d	FSQRT _q	—	—	—	—
03 ₁₆	—	—	—	—	—	—	—	—
04 ₁₆	—	FMUL _s	FMUL _d	FMUL _q	—	FDIV _s	FDIV _d	FDIV _q
05 ₁₆	—	FNMUL _s	FNMUL _d	—	—	—	—	—
06 ₁₆	—	FsMUL _d	—	—	—	—	FdMUL _q	—
07 ₁₆	—	FNsMUL _d	—	—	—	—	—	—
08 ₁₆	FxTO _d	—	—	—	FxTO _q	—	—	—
09 ₁₆	—	—	—	—	—	—	—	—
0A ₁₆	—	—	—	—	—	—	—	—
0B ₁₆	—	—	—	—	—	—	—	—
0C ₁₆	FiTO _d	FsTO _d	—	FqTO _d	FiTO _q	FsTO _q	FdTO _q	—
0D ₁₆	—	—	—	—	—	—	—	—
0E ₁₆ –1F ₁₆	—	—	—	—	—	—	—	—

TABLE A-6 opf{8:0} (op = 10₂, op3 = 35₁₆ = FPop2)

opf{8:4}	opf{3:0}								
	0	1	2	3	4	5	6	7	8-F
00 ₁₆	—	FMOV _s (fcc0)	FMOV _d (fcc0)	FMOV _q (fcc0)	—	† ‡	† ‡	† ‡	—
01 ₁₆	—	—	—	—	—	—	—	—	—
02 ₁₆	—	—	—	—	—	FMOV _R sZ ‡	FMOV _R dZ ‡	FMOV _R qZ ‡	—
03 ₁₆	—	—	—	—	—	—	—	—	—
04 ₁₆	—	FMOV _s (fcc1)	FMOV _d (fcc1)	FMOV _q (fcc1)	—	FMOV _R sLEZ ‡	FMOV _R dLEZ ‡	FMOV _R qLEZ ‡	—
05 ₁₆	—	FCMP _s	FCMP _d	FCMP _q	—	FCMP _E s ‡	FCMP _E d ‡	FCMP _E q ‡	—
06 ₁₆	—	—	—	—	—	FMOV _R sLZ ‡	FMOV _R dLZ ‡	FMOV _R qLZ ‡	—
07 ₁₆	—	—	—	—	—	—	—	—	—
08 ₁₆	—	FMOV _s (fcc2)	FMOV _d (fcc2)	FMOV _q (fcc2)	—	†	†	†	—
09 ₁₆	—	—	—	—	—	—	—	—	—
0A ₁₆	—	—	—	—	—	FMOV _R sNZ ‡	FMOV _R dNZ ‡	FMOV _R qNZ ‡	—
0B ₁₆	—	—	—	—	—	—	—	—	—
0C ₁₆	—	FMOV _s (fcc3)	FMOV _d (fcc3)	FMOV _q (fcc3)	—	FMOV _R sGZ ‡	FMOV _R dGZ ‡	FMOV _R qGZ ‡	—
0D ₁₆	—	—	—	—	—	—	—	—	—
0E ₁₆	—	—	—	—	—	FMOV _R sGEZ ‡	FMOV _R dGEZ ‡	FMOV _R qGEZ ‡	—
0F ₁₆	—	—	—	—	—	—	—	—	—
10 ₁₆	—	FMOV _s (icc)	FMOV _d (icc)	FMOV _q (icc)	—	—	—	—	—
11	—	—	—	—	—	—	—	—	—
12 ₁₆ –15 ₁₆ (4 rows)	—	—	—	—	—	—	—	—	—
16 ₁₆ –17 ₁₆ (2 rows)	—	—	—	—	—	—	—	—	—
18 ₁₆	—	FMOV _s (xcc)	FMOV _d (xcc)	FMOV _q (xcc)	—	—	—	—	—
19 ₁₆ –1B ₁₆ (3 rows)	—	—	—	—	—	—	—	—	—
1C ₁₆ –1F ₁₆ (4 rows)	—	—	—	—	—	—	—	—	—

† Reserved variation of FMOV_R ‡ bit 13 of instruction = 0

TABLE A-7 cond{3:0} (or for CBcond, c_hi :: c_lo)

		BPcc op = 0 op2 = 1 bit 28 = 0	CBcond op = 0 op2 = 1 bit 28 = 1	Bicc op = 0 op2 = 2	FBPfcc op = 0 op2 = 5	FBfcc ^D op = 0 op2 = 6	Tcc op = 2 op3 = 3A ₁₆ bit 29 = 0
cond {3:0}	0	BPN	—	BN ^D	FBPN	FBN ^D	TN
	1	BPE	C*BE	BE ^D	FBPNE	FBNE ^D	TE
	2	BPLE	C*BLE	BLE ^D	FBPLG	FBLG ^D	TLE
	3	BPL	C*BL	BL ^D	FBPUL	FBUL ^D	TL
	4	BPLEU	C*BLEU	BLEU ^D	FBPL	FBL ^D	TLEU
	5	BPCS	C*BCS	BCS ^D	FBPUG	FBUG ^D	TCS
	6	BPNEG	C*BNEG	BNEG ^D	FBPG	FBG ^D	TNEG
	7	BPVS	C*BVS	BVS ^D	FBPU	FBU ^D	TVS
	8	BPA	—	BA ^D	FBPA	FBA ^D	TA
	9	BPNE	C*BNE	BNE ^D	FBPE	FBE ^D	TNE
	A	BPG	C*BG	BG ^D	FBPUE	FBUE ^D	TG
	B	BPGE	C*BGE	BGE ^D	FBPGE	FBGE ^D	TGE
	C	BPGU	C*BGU	BGU ^D	FBPUGE	FBUGE ^D	TGU
	D	BPCC	C*BCC	BCC ^D	FBPLE	FBLE ^D	TCC
	E	BPPOS	C*BPOS	BPOS ^D	FBPULE	FBULE ^D	TPOS
	F	BPVC	C*BVC	BVC ^D	FBPO	FBO ^D	TVC

TABLE A-8 Encoding of rcond{2:0} Instruction Field

		BPr op = 0 op2 = 3	MOVr op = 2 op3 = 2F ₁₆	FMOVr op = 2 op3 = 35 ₁₆
rcond {2:0}	0	—	—	—
	1	BRZ	MOVZR	FMOVr<s d q>Z
	2	BRLEZ	MOVRLEZ	FMOVr<s d q>LEZ
	3	BRLZ	MOVRLZ	FMOVr<s d q>LZ
	4	—	—	—
	5	BRNZ	MOVRNZ	FMOVr<s d q>NZ
	6	BRGZ	MOVRGZ	FMOVr<s d q>GZ
	7	BRGEZ	MOVRGEZ	FMOVr<s d q>GEZ

TABLE A-9 cc / opf_cc Fields (MOVcc and FMOVcc)

opf_cc			Condition Code Selected
cc2	cc1	cc0	
0	0	0	fcc0
0	0	1	fcc1
0	1	0	fcc2
0	1	1	fcc3
1	0	0	icc

TABLE A-9 cc / opf_cc Fields (MOVcc and FMOVcc)

1	0	1	—
1	1	0	xcc
1	1	1	—

TABLE A-10 cc Fields (FBPfcc, FCMP, and FCMPE)

cc1	cc0	Condition Code Selected
0	0	fcc0
0	1	fcc1
1	0	fcc2
1	1	fcc3

TABLE A-11 cc Fields (BPcc and Tcc)

cc1	cc0	Condition Code Selected
0	0	icc
0	1	—
1	0	xcc
1	1	—

TABLE A-12 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 1

		opf {8:4}							
		00	01	02	03	04	05	06	07
	0	EDGE8cc	ARRAY8	FPCMPLE16	—	FMEAN16	FPADD16	FZERO	FAND
	1	EDGE8N	ADDXC	FSL16	FMUL8x16	SUBXC	FPADD16s	FZEROS	FANDs
	2	EDGE8Lcc	ARRAY16	FPCMPNE16/ FPCMPUNE16	—	—FPADD64	FPADD32	FNOR	FXNOR
	3	EDGE8LN	ADDXCcc	FSRL16	FMUL8x16AU	SUBXCcc	FPADD32s	FNORS	FXNORS
	4	EDGE16cc	ARRAY32	FPCMPLE32/ FPCMPUEQ32	FCMPLE8	FCHKSM16	FPSUB16	FANDNOT2	FSRC1
	5	EDGE16N	—	FSL32	FMUL8x16AL	—	FPSUB16s	FANDNOT2s	FSRC1s
	6	EDGE16Lcc	UMULXHI	FPCMPNE32/ FPCMPUNE32	FMUL8SUx16	FPSUB64	FPSUB32	FNOT2	FORNOT2
	7	EDGE16LN	LZCNT	FSRL32	FMUL8ULx16	—	FPSUB32s	FNOT2s	FORNOT2s

TABLE A-13 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 2

		opf {8:4}							
		00	01	02	03	04	05	06	07
	8	EDGE32cc	ALIGNADDRESS	FPCMPGT16	FMULD8SUx16	FALIGNDATAg	FPADDS16	FANDNOT1	FSRC2
	9	EDGE32N	BMASK	FSLAS16	FMULD8ULx16	FALIGNDATAi	FPADDS16s	FANDNOT1s	FSRC2s
	A	EDGE32Lcc	ALIGNADDRESS_ LITTLE	FPCMPPEQ16/ FPCMPUEQ16	FPACK32	—	FPADDS32	FNOT1	FORNOT1
	B	EDGE32LN	CMASK8 ((bits 29:25 = 0) and (bits 18:14 = 0))	FSRA16	FPACK16	FPMERGE	—FPADDS32s	FNOT1s	FORNOT1s

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 3

		opf {8:4}							
		00	01	02	03	04	05	06	07
	C	—	—	FPCMPGT32	FCMPGT8	BSHUFFLE	FPSUBS16	FXOR	FOR
	D	—	CMASK16 ((bits 29:25 = 0) and (bits 18:14 = 0)) — ((bits 29:25 ≠ 0) or (bits 18:14 ≠ 0))	FSLAS32	FPCMPFIX	FEXPAND	FPSUBS16s	FXORS	FORs
	E	—	—	FPCMPPEQ32/ FPCMPUEQ32	PDIST ^D	—	FPSUBS32	FNAND	FONE
	F	—	CMASK32 ((bits 29:25 = 0) and (bits 18:14 = 0))	FSRA32	PDISTN	—	FPSUBS32S	FNANDs	FONEs

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 4

		opf {8:4}							
		08	09	0A	0B	0C	0D	0E	0F
opf {3:0}	0	—	—	—	—	—	—	—	—
	1	SIAM	—	—	—	—	—	—	—
	2	—	—	—	—	Reserved	Reserved	—	—
	3	—	—	—	—	—	—	—	—
	4	—	—	—	—	—	—	—	—
	5	—	—	—	Reserved	—	—	—	—
	6	—	—	—	—	Reserved	Reserved	—	—
	7	—	—	Reserved	—	—	—	—	—

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 5

		opf {8:4}							
		08	09	0A	0B	0C	0D	0E	0F
opf {3:0}	8	—	—	Reserved	—	—	—	—	—
	9	—	—	Reserved	—	—	—	—	—
	A	—	—	Reserved	—	Reserved	Reserved	—	—
	B	—	—	Reserved	Reserved	—	—	—	—
	C	—	—	—	Reserved	—	—	—	—
	D	—	—	—	Reserved	—	—	—	—
	E	—	—	—	—	Reserved	Reserved	—	—
	F	—	—	—	—	—	—	—	—

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 6

		opf {8:4}					
		10	11	12	13	14	15
opf {3:0}	0	—	MOVdTOx	FPCMPULE8	3-op AES_KEXPAND0 ^N	1-op MD5 ^N	—
	1	—	MOV _s TOuw	—	3-op AES_KEXPAND2 ^N	1-op SHA1 ^N	FLCMPs
	2	—	—	FPCMPNE8/ FPCMPUNE8	—	1-op SHA256 ^N	FLCMPd
	3	—	MOV _s TOsw	FPADDUS16	—	1-op SHA512 ^N	FPSUBUS16
	4	—	—	FPADD8	2-op DES_IP ^N	—	FPSUB8
	5	—	XMULX	—	2-op DES_IIP ^N	—	—
	6	—	XMULXHI	FPADDUS8	2-op DES_KEXPAND ^N	—	FPSUBS8
	7	—	—	—	—	3-op CRC32C ^N	—
7	—	—	FPADDUS8	—	3-op CRC32C ^N	FPSUBUS8	

opf {3:0}

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 7

		opf {8:4}					
		10	11	12	13	14	15
	8	—	MOVxTOd	FPCMPUGT8	—	MPMUL ^N (rd=0, rs1=0) XMPMUL ^N (rd=1,rs1=0) — (rd>1, rs1=0)	—
	9	—	MOVwTOs	—	—	MONTMUL ^N (rd=0, rs1=0) XMONTMUL ^N (rd=1,rs1=0) — (rd>1, rs1=0)	—
	A	—	FPMIN8	FPCMPEQ8/ FPCMPUEQ8	—	MONTSQR ^N (rd=0, rs1=0) XMONTSQR ^N (rd=1,rs1=0) — (rd>1, rs1=0)	FPMINU8
	B	—	FPMIN16	FPCMPUGT16	—	—	FPMINU16
	C	—	FPMIN32	FPCMPUGT32	3-op CAMELIA_FL ^N	—	FPMINU32
	D	—	FPMAX8	—	3-op CAMELIA_FLI ^N	—	FPMAXU8
	E	—	FPMAX16	FPCMPULE16	—	—	FPMAXU16
	F	—	FPMAX32	FPCMPULE32	—	—	FPMAXU32

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 8

		opf {8:4}						
		16	17	18	19	1A	1B	1C–1F
	0		—	—	—	—		
	1		—	—	—	—		
	2		—	—	—	—		
	3		—	—	—	—		
	4		—	—	—	—		
	5		—	—	—	—		
	6		—	—	—	—		
	7		—	—	—	—		

TABLE A-14 opf{8:0} for VIS opcodes (op = 10₂, op3 = 36₁₆) — part 9

		opf {8:4}						
		16	17	18	19	1A	1B	1C–1F
	8		—	—	—	—		
	9		—	—	—	—		
	A		—	—	—	—		
	B		—	—	—	—		
	C		—	—	—	—		
	D		—	—	—	—		
	E		—	—	—	—		
	F		—	—	—	—		

opf {3:0}

—

Reserved

Reserved

TABLE A-15 op5{3:0} (op = 10₂, op3 = 37₁₆ = FMAf)

		op5{1:0} (size)			
		0	1	2	3
op5{3:2} (var)	0	FPMADDX	FMADDs	FMADDd	—
	1	FPMADDXHI	FMSUBs	FMSUBd	—
	2	—	FNMSUBs	FNMSUBd	—
	3	—	FNMADDs	FNMADDd	—

Note: This chapter is undergoing final review; please check back later for a copy of Oracle SPARC Architecture 2015 containing the final version of this chapter.

Implementation Dependencies

This appendix summarizes implementation dependencies in the SPARC V9 standard. In SPARC V9, the notation “**IMPL. DEP. #*nn***” identifies the definition of an implementation dependency; the notation “(impl. dep. #*nn*)” identifies a reference to an implementation dependency. These dependencies are described by their number *nn* in TABLE B-1 on page 499.

The appendix contains these sections:

- **Definition of an Implementation Dependency** on page 497.
- **Hardware Characteristics** on page 498.
- **Implementation Dependency Categories** on page 498.
- **List of Implementation Dependencies** on page 498.

B.1 Definition of an Implementation Dependency

The SPARC V9 architecture is a *model* that specifies unambiguously the behavior observed by *software* on SPARC V9 systems. Therefore, it does not necessarily describe the operation of the *hardware* of any actual implementation.

An implementation is *not* required to execute every instruction in hardware. An attempt to execute a SPARC V9 instruction that is not implemented in hardware generates a trap. Whether an instruction is implemented directly by hardware, simulated by software, or emulated by firmware is implementation dependent.

The two levels of SPARC V9 compliance are described in *Oracle SPARC Architecture 2015 Compliance with SPARC V9 Architecture* on page 16.

Some elements of the architecture are defined to be implementation dependent. These elements include certain registers and operations that may vary from implementation to implementation; they are explicitly identified as such in this appendix.

Implementation elements (such as instructions or registers) that appear in an implementation but are not defined in this document (or its updates) are not considered to be SPARC V9 elements of that implementation.

B.2 Hardware Characteristics

Hardware characteristics that do not affect the behavior observed by software on SPARC V9 systems are not considered architectural implementation dependencies. A hardware characteristic may be relevant to the user system design (for example, the speed of execution of an instruction) or may be transparent to the user (for example, the method used for achieving cache consistency). The SPARC International document, *Implementation Characteristics of Current SPARC V9-based Products, Revision 9.x*, provides a useful list of these hardware characteristics, along with the list of implementation-dependent design features of SPARC V9-compliant implementations.

In general, hardware characteristics deal with

- Instruction execution speed
- Whether instructions are implemented in hardware
- The nature and degree of concurrency of the various hardware units constituting a SPARC V9 implementation

B.3 Implementation Dependency Categories

Many of the implementation dependencies can be grouped into four categories, abbreviated by their first letters throughout this appendix:

- **Value (v)**
The semantics of an architectural feature are well defined, except that a value associated with the feature may differ across implementations. A typical example is the number of implemented register windows (impl. dep. #2-V8).
- **Assigned Value (a)**
The semantics of an architectural feature are well defined, except that a value associated with the feature may differ across implementations and the actual value is assigned by SPARC International. Typical examples are the `impl` field of the Version register (VER) (impl. dep. #13-V8) and the `FSR.ver` field (impl. dep. #19-V8).
- **Functional Choice (f)**
The SPARC V9 architecture allows implementors to choose among several possible semantics related to an architectural function. A typical example is the treatment of a catastrophic error exception, which may cause either a deferred or a disrupting trap (impl. dep. #31-V8-Cs10).
- **Total Unit (t)**
The existence of the architectural unit or function is recognized, but details are left to each implementation. Examples include the handling of I/O registers (impl. dep. #7-V8) and some alternate address spaces (impl. dep. #29-V8).

B.4 List of Implementation Dependencies

TABLE B-1 provides a complete list of the SPARC V9 implementation dependencies. The Page column lists the page for the context in which the dependency is defined; bold face indicates the main page on which the implementation dependency is described.

TABLE B-1 SPARC V9 Implementation Dependencies (1 of 7)

Nbr	Category	Description	Page
1-V8	f	Software emulation of instructions Whether an instruction complies with Oracle SPARC Architecture 2015 by being implemented directly by hardware, simulated by software, or emulated by firmware is implementation dependent.	16
2-V8	v	Number of IU registers An Oracle SPARC Architecture implementation may contain from 72 to 640 general-purpose 64-bit R registers. This corresponds to a grouping of the registers into <i>MAXPGL</i> + 1 sets of global R registers plus a circular stack of <i>N_REG_WINDOWS</i> sets of 16 registers each, known as register windows. The number of register windows present (<i>N_REG_WINDOWS</i>) is implementation dependent, within the range of 3 to 32 (inclusive).	17, 34
3-V8	f	Incorrect IEEE Std 754-1985 results An implementation may indicate that a floating-point instruction did not produce a correct IEEE Std 754-1985 result by generating an <i>fp_exception_other</i> exception with <i>FSR.ftt</i> = <i>unfinished_FPop</i> . In this case, software running in a higher privilege mode shall emulate any functionality not present in the hardware.	92
4, 5		<i>Reserved.</i>	
6-V8	f	I/O registers privileged status Whether I/O registers can be accessed by nonprivileged code is implementation dependent.	19
7-V8	t	I/O register definitions The contents and addresses of I/O registers are implementation dependent.	19
8-V8-	t	RDAsr/WRAsr target registers	20, 48,
■ Cs20		Ancillary state registers (ASRs) in the range 0–27 that are not defined in Oracle SPARC Architecture 2015 are reserved for future architectural use.	310, 373
9-V8-	f	RDAsr/WRAsr privileged status	20, 48,
■ Cs20		The privilege level required to execute each of the implementation-dependent read/write ancillary state register instructions (for ASRs 28–31) is implementation dependent.	310, 373
10-V8–12-V8		<i>Reserved.</i>	
13-V8	a	(this implementation dependency applies to execution modes with greater privileges)	
14-V8–15-V8		<i>Reserved.</i>	
16-V8-Cu3		<i>Reserved.</i>	
17-V8		<i>Reserved.</i>	

TABLE B-1 SPARC V9 Implementation Dependencies (2 of 7)

Nbr	Category	Description	Page
18-V8- Ms10	f	<p>Nonstandard IEEE 754-1985 results</p> <p>When FSR.ns = 1, the FPU produces implementation-dependent results that may not correspond to IEEE Standard 754-1985.</p> <p>a: When FSR.ns = 1 and a floating-point <i>source operand</i> is subnormal, an implementation may treat the subnormal operand as if it were a floating-point zero value of the same sign.</p> <p>The cases in which this replacement is performed are implementation dependent. However, if it occurs,</p> <p>(1) it should <i>not</i> apply to FLCMP, FABS, FMOV, or FNEG instructions and</p> <p>(2) FADD, FSUB, FLCMPE, and FCMP should give identical treatment to subnormal source operands.</p> <p>Treating a subnormal source operand as zero may generate an IEEE 754 floating-point “inexact”, “division by zero”, or “invalid” condition (see <i>Current Exception (cexc)</i> on page 46). Whether the generated condition(s) trigger an <i>fp_exception_ieee_754</i> exception or not depends on the setting of FSR.tem.</p> <p>b: When a floating-point operation generates a subnormal <i>result</i> value, an Oracle SPARC Architecture implementation may either write the result as a subnormal value or replace the subnormal result by a floating-point zero value of the same sign and generate IEEE 754 floating-point “inexact” and “underflow” conditions. Whether these generated conditions trigger an <i>fp_exception_ieee_754</i> exception or not depends on the setting of FSR.tem.</p> <p>c: If an FPop generates an <i>intermediate</i> result value, the intermediate value is subnormal, and FSR.ns = 1, it is implementation dependent whether (1) the operation continues, using the subnormal value (possibly with some loss of accuracy), or (2) the virtual processor replaces the subnormal intermediate value with a floating-point zero value of the same sign, generates IEEE 754 floating-point “inexact” and “underflow” conditions, completes the instruction, and writes a final result (possibly with some loss of accuracy). Whether generated IEEE conditions trigger an <i>fp_exception_ieee_754</i> exception or not depends on the setting of FSR.tem.</p>	394
19-V8	a	<p>FPU version, FSR.ver</p> <p>Bits 19:17 of the FSR, FSR.ver, identify one or more implementations of the FPU architecture.</p>	43
20-V8–21-V8		<i>Reserved.</i>	
22-V8	f	<p>FPU tem, cexc, and aexc</p> <p>An Oracle SPARC Architecture implementation implements the tem, cexc, and aexc fields in hardware, conformant to IEEE Std 754-1985.</p>	48
23-V8		<i>Reserved.</i>	
24-V8		<i>Reserved.</i>	
25-V8	f	<p>RDPR of FQ with nonexistent FQ</p> <p>An Oracle SPARC Architecture implementation does not contain a floating-point queue (FQ). Therefore, FSR.ftt = 4 (sequence_error) does not occur, and an attempt to read the FQ with the RDPR instruction causes an <i>illegal_instruction</i> exception.</p>	45, 314
26-V8–28-V8		<i>Reserved.</i>	
29-V8	t	<p>Address space identifier (ASI) definitions</p> <p>In SPARC V9, many ASIs were defined to be implementation dependent. Some of those ASIs have been allocated for standard uses in the Oracle SPARC Architecture. Others remain implementation dependent in the Oracle SPARC Architecture. See <i>ASI Assignments</i> on page 424 and <i>Block Load and Store ASIs</i> on page 436 for details.</p>	84
30-V8- Cu3	f	<p>ASI address decoding</p> <p>In SPARC V9, an implementation could choose to decode only a subset of the 8-bit ASI specifier. In Oracle SPARC Architecture implementations, all 8 bits of each ASI specifier must be decoded. Refer to Chapter 10, <i>Address Space Identifiers (ASIs)</i>, of this specification for details.</p>	84

TABLE B-1 SPARC V9 Implementation Dependencies (3 of 7)

Nbr	Category	Description	Page
31-V8-Cs10	f	This implementation dependency is no longer used in the Oracle SPARC Architecture, since “catastrophic” errors are now handled using normal error-reporting mechanisms.	—
32-V8-MS10	t	Restartable deferred traps Whether any restartable deferred traps (and associated deferred-trap queues) are present is implementation dependent.	447
33-V8-Cs10	f	Trap precision In an Oracle SPARC Architecture implementation, all exceptions that occur as the result of program execution are precise.	449
34-V8	f	Interrupt clearing a: The method by which an interrupt is removed is now defined in the Oracle SPARC Architecture (see <i>Clearing the Software Interrupt Register</i> on page 468). b: How quickly a virtual processor responds to an interrupt request, like all timing-related issues, is implementation dependent.	468
35-V8-Cs20	t	Implementation-dependent traps Trap type (TT) values 060 ₁₆ –07F ₁₆ were reserved for <i>implementation_dependent_exception_n</i> exceptions in SPARC V9 but are now all defined as standard Oracle SPARC Architecture exceptions.	451
36-V8	f	Trap priorities The relative priorities of traps defined in the Oracle SPARC Architecture are fixed. However, the absolute priorities of those traps are implementation dependent (because a future version of the architecture may define new traps). The priorities (both absolute and relative) of any new traps are implementation dependent.	458
41-V8		<i>Reserved.</i>	
42-V8-Cs10	t, f, v	FLUSH instruction FLUSH is implemented in hardware in all Oracle SPARC Architecture 2015 implementations, so never causes a trap as an unimplemented instruction.	
43-V8		<i>Reserved.</i>	
44-V8-Cs10	f	Data access FPU trap a: If a load floating-point instruction generates an exception that causes a non-precise trap, it is implementation dependent whether the contents of the destination floating-point register(s) or floating-point state register are undefined or are guaranteed to remain unchanged. b: If a load floating-point alternate instruction generates an exception that causes a non-precise trap, it is implementation dependent whether the contents of the destination floating-point register(s) are undefined or are guaranteed to remain unchanged.	247, 264 250
45-V8–46-V8		<i>Reserved.</i>	
47-V8-Cs20	t	RDAsr RDAsr instructions with rd in the range 28–31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For an RDAsr instruction with rs1 in the range 28–31, the following are implementation dependent: <ul style="list-style-type: none"> the interpretation of bits 13:0 and 29:25 in the instruction whether the instruction is nonprivileged or privileged (impl. dep. #9-V8-Cs20) whether an attempt to execute the instruction causes an <i>illegal_instruction</i> exception 	311
48-V8-Cs20	t	WRAsr WRAsr instructions with rd of 16-18, 28, 29, or 31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For a WRAsr instruction using one of those rd values, the following are implementation dependent: <ul style="list-style-type: none"> the interpretation of bits 18:0 in the instruction the operation(s) performed (for example, xor) to generate the value written to the ASR whether the instruction is nonprivileged or privileged (impl. dep. #9-V8-Cs20) whether an attempt to execute the instruction causes an <i>illegal_instruction</i> exception 	374

TABLE B-1 SPARC V9 Implementation Dependencies (4 of 7)

Nbr	Category	Description	Page
49-V8-54-V8		<i>Reserved.</i>	
55-V8-Cs10	f	Tininess detection In SPARC V9, it is implementation-dependent whether “tininess” (an IEEE 754 term) is detected before or after rounding. In all Oracle SPARC Architecture implementations, tininess is detected before rounding.	48
56-100		<i>Reserved.</i>	
101-V9-CS10	v	Maximum trap level (MAXPTL) The architectural parameter <i>MAXPTL</i> is a constant for each implementation; its legal values are from 2 to 6 (supporting from 2 to 6 levels of saved trap state). In a typical implementation <i>MAXPTL</i> = <i>MAXPGL</i> (see impl. dep. #401-S10). Architecturally, <i>MAXPTL</i> must be ≥ 2 .	74, 75
102-V9	f	Clean windows trap An implementation may choose either to implement automatic “cleaning” of register windows in hardware or to generate a <i>clean_window</i> trap, when needed, for window(s) to be cleaned by software.	460
103-V9-Ms10	f	Prefetch instructions The following aspects of the PREFETCH and PREFETCHA instructions are implementation dependent:	
		a: the attributes of the block of memory prefetched: its size (minimum = 64 bytes) and its alignment (minimum = 64-byte alignment)	304
		b: whether each defined prefetch variant is implemented (1) as a NOP, (2) with its full semantics, or (3) with common-case prefetching semantics	304, 306
		c: whether and how variants 16, 18, 19 and 24–31 are implemented; if not implemented, a variant must execute as a NOP	308C
		The following aspects of the PREFETCH and PREFETCHA instructions used to be (but are no longer) implementation dependent:	
		d: while in nonprivileged mode (<i>PSTATE.priv</i> = 0), an attempt to reference an ASI in the range $0_{16}..7F_{16}$ by a PREFETCHA instruction executes as a NOP; specifically, it does not cause a <i>privileged_action</i> exception.	—
		e: PREFETCH and PREFETCHA have no observable effect in privileged code	—
		g: while in privileged mode (<i>PSTATE.priv</i> = 1), an attempt to reference an ASI in the range $30_{16}..7F_{16}$ by a PREFETCHA instruction executes as a NOP (specifically, it does not cause a <i>privileged_action</i> exception)	—
105-V9	f	TICK register a: If an accurate count cannot always be returned when TICK is read, any inaccuracy should be small, bounded, and documented. b: An implementation may implement fewer than 63 bits in <i>TICK.counter</i> ; however, the counter as implemented must be able to count for at least 10 years without overflowing. Any upper bits not implemented must read as 0.	53
106-V9Cs10	f	IMPDEP2A instructions The IMPDEP2A instructions were defined to be completely implementation dependent in SPARC V9. The opcodes that have not been used in this space are now just documented as reserved opcodes.	
107-V9	f	Unimplemented LDTW[A] trap a: It is implementation dependent whether LDTW is implemented in hardware. If not, an attempt to execute an LDTW instruction will cause an <i>unimplemented_LDTW</i> exception. b: It is implementation dependent whether LDTWA is implemented in hardware. If not, an attempt to execute an LDTWA instruction will cause an <i>unimplemented_LDTW</i> exception.	257 259

TABLE B-1 SPARC V9 Implementation Dependencies (5 of 7)

Nbr	Category	Description	Page
108- V9	f	<p>Unimplemented STTW[A] trap</p> <p>a: It is implementation dependent whether STTW is implemented in hardware. If not, an attempt to execute an STTW instruction will cause an <i>unimplemented_STTW</i> exception.</p> <p>b: It is implementation dependent whether STDA is implemented in hardware. If not, an attempt to execute an STTWA instruction will cause an <i>unimplemented_STTW</i> exception.</p>	352 354
109- V9- Cs10	f	<p>LDDF[A]_mem_address_not_aligned</p> <p>a: LDDF requires only word alignment. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute a valid (i = 1 or instruction bits 12:5 = 0) LDDF instruction may cause an <i>LDDF_mem_address_not_aligned</i> exception. In this case, the trap handler software shall emulate the LDDF instruction and return. (In an Oracle SPARC Architecture processor, the <i>LDDF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the LDDF instruction)</p> <p>b: LDDFA requires only word alignment. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute a valid (i = 1 or instruction bits 12:5 = 0) LDDFA instruction may cause an <i>LDDF_mem_address_not_aligned</i> exception. In this case, the trap handler software shall emulate the LDDFA instruction and return. (In an Oracle SPARC Architecture processor, the <i>LDDF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the LDDFA instruction)</p>	79, 79, 337, 463 340
110- V9- Cs10	f	<p>STDF[A]_mem_address_not_aligned</p> <p>a: STDF requires only word alignment in memory. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute a valid (i = 1 or instruction bits 12:5 = 0) STDF instruction may cause an <i>STDF_mem_address_not_aligned</i> exception. In this case, the trap handler software must emulate the STDF instruction and return. (In an Oracle SPARC Architecture processor, the <i>STDF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the STDF instruction)</p> <p>b: STDFA requires only word alignment in memory. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute a valid (i = 1 or instruction bits 12:5 = 0) STDFA instruction may cause an <i>STDF_mem_address_not_aligned</i> exception. In this case, the trap handler software must emulate the STDFA instruction and return. (In an Oracle SPARC Architecture processor, the <i>STDF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the STDFA instruction)</p>	79, 471, 464 473

TABLE B-1 SPARC V9 Implementation Dependencies (6 of 7)

Nbr	Category	Description	Page
111- V9- Cs10	f	<p>LDQF[A]_mem_address_not_aligned</p> <p>a: LDQF requires only word alignment. However, if the effective address is word-aligned but not quadword-aligned, an attempt to execute an LDQF instruction may cause an <i>LDQF_mem_address_not_aligned</i> exception. In this case, the trap handler software must emulate the LDQF instruction and return. (In an Oracle SPARC Architecture processor, the <i>LDQF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the LDQF instruction) (this exception does not occur in hardware on Oracle SPARC Architecture 2015 implementations, because they do not implement the LDQF instruction in hardware)</p> <p>b: LDQFA requires only word alignment. However, if the effective address is word-aligned but not quadword-aligned, an attempt to execute an LDQFA instruction may cause an <i>LDQF_mem_address_not_aligned</i> exception. In this case, the trap handler software must emulate the LDQF instruction and return. (In an Oracle SPARC Architecture processor, the <i>LDQF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the LDQFA instruction) (this exception does not occur in hardware on Oracle SPARC Architecture 2015 implementations, because they do not implement the LDQFA instruction in hardware)</p>	80, 79, 337, 464 340
112- V9- Cs10	f	<p>STQF[A]_mem_address_not_aligned</p> <p>a: STQF requires only word alignment in memory. However, if the effective address is word aligned but not quadword aligned, an attempt to execute an STQF instruction may cause an <i>STQF_mem_address_not_aligned</i> exception. In this case, the trap handler software must emulate the STQF instruction and return. (In an Oracle SPARC Architecture processor, the <i>STQF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the STQF instruction) (this exception does not occur in hardware on Oracle SPARC Architecture 2015 implementations, because they do not implement the STQF instruction in hardware)</p> <p>b: STQFA requires only word alignment in memory. However, if the effective address is word aligned but not quadword aligned, an attempt to execute an STQFA instruction may cause an <i>STQF_mem_address_not_aligned</i> exception. In this case, the trap handler software must emulate the STQFA instruction and return. (In an Oracle SPARC Architecture processor, the <i>STQF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the STQFA instruction) (this exception does not occur in hardware on Oracle SPARC Architecture 2015 implementations, because they do not implement the STQFA instruction in hardware)</p>	80, 471, 464 473
113- V9- Ms10	f	<p>Implemented memory models</p> <p>Whether memory models represented by $PSTATE.mm = 10_2$ or 11_2 are supported in an Oracle SPARC Architecture processor is implementation dependent. If the 10_2 model is supported, then when $PSTATE.mm = 10_2$ the implementation must correctly execute software that adheres to the RMO model described in <i>The SPARC Architecture Manual-Version 9</i>. If the 11_2 model is supported, its definition is implementation dependent.</p>	71, 415
118- V9	f	<p>Identifying I/O locations</p> <p>The manner in which I/O locations are identified is implementation dependent.</p>	409
119- Ms10	f	<p>Unimplemented values for PSTATE.mm</p> <p>The effect of an attempt to write an unsupported memory model designation into $PSTATE.mm$ is implementation dependent; however, it should never result in a value of $PSTATE.mm$ value greater than the one that was written. In the case of an Oracle SPARC Architecture implementation that only supports the TSO memory model, $PSTATE.mm$ always reads as zero and attempts to write to it are ignored.</p>	71, 416

TABLE B-1 SPARC V9 Implementation Dependencies (7 of 7)

Nbr	Category	Description	Page
120-V9	f	Coherence and atomicity of memory operations The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent.	409
121-V9	f	Implementation-dependent memory model An implementation may choose to identify certain addresses and use an implementation-dependent memory model for references to them.	409
122-V9	f	FLUSH latency The latency between the execution of FLUSH on one virtual processor and the point at which the modified instructions have replaced outdated instructions in a multiprocessor is implementation dependent.	171, 421
123-V9	f	Input/output (I/O) semantics The semantic effect of accessing I/O registers is implementation dependent.	19
124-V9	v	Implicit ASI when TL > 0 In SPARC V9, when TL > 0, the implicit ASI for instruction fetches, loads, and stores is implementation dependent. In Oracle SPARC Architecture 2015 implementations, when TL > 0, the implicit ASI for instruction fetches is ASI_NUCLEUS; loads and stores will use ASI_NUCLEUS if PSTATE.cle = 0 or ASI_NUCLEUS_LITTLE if PSTATE.cle = 1.	411
125-V9-Cs10	f	Address masking (1) When PSTATE.am = 1, only the less-significant 32 bits of the PC register are stored in the specified destination register(s) in CALL, JMPL, and RDPC instructions, while the more-significant 32 bits of the destination registers(s) are set to 0. (2) When PSTATE.am = 1, during a trap, only the less-significant 32 bits of the PC and NPC are stored (respectively) to TPC[TL] and TNPC[TL]; the more-significant 32 bits of TPC[TL] and TNPC[TL] are set to 0.	72, 72, 130, 238, 311, 459
126-V9-Ms10-Cs40		Register Windows State registers width (this implementation dependency was replaced by implementation dependencies 600-Ms50, 602-Ms50, and 601-Ms50 in Oracle SPARC Architecture 2015)	
127-199		<i>Reserved.</i>	—

TABLE B-2 provides a list of implementation dependencies that, in addition to those in TABLE B-1, apply to Oracle SPARC Architecture processors. Bold face indicates the main page on which the implementation dependency is described. See Appendix C in the Extensions Documents for further information.

TABLE B-2 Oracle SPARC Architecture Implementation Dependencies (1 of 6)

Nbr	Description	Page
200-201	<i>Reserved.</i>	—
203-U3-Cs10	Dispatch Control register (DCR) bits 13:6 and 1 <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	
204-U3-Cs10	DCR bits 5:3 and 0 <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	
205-U3-Cs10	Instruction Trap Register <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	
206-U3-Cs10	SHUTDOWN instruction	
208-U3	Ordering of errors captured in instruction execution The order in which errors are captured in instruction execution is implementation dependent. Ordering may be in program order or in order of detection.	—

TABLE B-2 Oracle SPARC Architecture Implementation Dependencies (2 of 6)

Nbr	Description	Page
209-U3	Software intervention after instruction-induced error Precision of the trap to signal an instruction-induced error of which recovery requires software intervention is implementation dependent.	—
211-U3	Error logging registers' information The information that the error logging registers preserves beyond the reset induced by an ERROR signal is implementation dependent.	—
212-U3- Cs10	<i>Trap with fatal error</i> <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
213-U3	AFSR .priv The existence of the AFSR .priv bit is implementation dependent. If AFSR .priv is implemented, it is implementation dependent whether the logged AFSR .priv indicates the privileged state upon the detection of an error or upon the execution of an instruction that induces the error. For the former implementation to be effective, operating software must provide error barriers appropriately.	—
226-U3		
227-U3	TSB number of entries The maximum number of entries in a TSB is implementation dependent in the Oracle SPARC Architecture (to a maximum of 16 million).	479
228-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
229-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i> TSB Base address generation Whether the implementation generates the TSB Base address by exclusive-ORing the TSB Base register and a TSB register or by taking the <code>tsb_base</code> field directly from a TSB register is implementation dependent in Oracle SPARC Architecture. This implementation dependency existed for UltraSPARC III/IV, only to maintain compatibility with the TLB miss handling software of UltraSPARC I/II.	—
230	<i>Reserved.</i>	—
230-U3- Cs20	<i>This implementation dependency no longer applies, in Oracle SPARC Architecture 2015</i>	—
232-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
233-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
235-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
236-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i> t	—
239-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
240-U3- Cs10	<i>Reserved.</i>	—
243-U3	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
244-U3- Cs10	Data Watchpoint Reliability Data Watchpoint traps are completely implementation-dependent in Oracle SPARC Architecture processors.	—
245-U3- Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—

TABLE B-2 Oracle SPARC Architecture Implementation Dependencies (3 of 6)

Nbr	Description	Page
248-U3	Conditions for <i>fp_exception_other</i> with unfinished_FPop The conditions under which an <i>fp_exception_other</i> exception with floating-point trap type of unfinished_FPop can occur are implementation dependent. An implementation may cause <i>fp_exception_other</i> with unfinished_FPop under a different (but specified) set of conditions.	45
249-U3-Cs10	Data Watchpoint for Partial Store Instruction For an STPARTIAL instruction, the following aspects of data watchpoints are implementation dependent: (a) whether data watchpoint logic examines the byte store mask in R[rs2] or it conservatively behaves as if every Partial Store always stores all 8 bytes, and (b) whether data watchpoint logic examines individual bits in the Virtual (Physical) Data Watchpoint Mask in DCUCR to determine which bytes are being watched or (when the Watchpoint Mask is nonzero) it conservatively behaves as if all 8 bytes are being watched.	349
250-U3-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2007.</i>	—
251	<i>Reserved.</i>	
252-U3-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
253-U3-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
255-U3-Cs10	LDDFA with ASI E0₁₆ or E1₁₆ and misaligned destination register number If an LDDFA opcode is used with an ASI of E0 ₁₆ or E1 ₁₆ (Block Store Commit ASI, an illegal combination with LDDFA) and a destination register number rd is specified which is not a multiple of 8 (“misaligned” rd), an Oracle SPARC Architecture virtual processor generates an <i>illegal_instruction</i> exception.	250
256-U3	LDDFA with ASI E0₁₆ or E1₁₆ and misaligned memory address If an LDDFA opcode is used with an ASI of E0 ₁₆ or E1 ₁₆ (Block Store Commit ASI, an illegal combination with LDDFA) and a memory address is specified with less than 64-byte alignment, the virtual processor generates an exception. It is implementation dependent whether the exception generated is <i>DAE_invalid_asi</i> , <i>mem_address_not_aligned</i> , or <i>LDDF_mem_address_not_aligned</i> .	250
257-U3	LDDFA with ASI C0₁₆–C5₁₆ or C8₁₆–CD₁₆ and misaligned memory address If an LDDFA opcode is used with an ASI of C0 ₁₆ –C5 ₁₆ or C8 ₁₆ –CD ₁₆ (Partial Store ASIs, which are an illegal combination with LDDFA) and a memory address is specified with less than 8-byte alignment, the virtual processor generates an exception. It is implementation dependent whether the exception generated is <i>DAE_invalid_asi</i> , <i>mem_address_not_aligned</i> , or <i>LDDF_mem_address_not_aligned</i> .	250
259–299	<i>Reserved.</i>	—
300-U4-Cs10	Attempted access to ASI registers with LDTWA If an LDTWA instruction referencing a non-memory ASI is executed, it generates a <i>DAE_invalid_asi</i> exception.	260
301-U4-Cs10	Attempted access to ASI registers with STTWA If an STTWA instruction referencing a non-memory ASI is executed, it generates a <i>DAE_invalid_asi</i> exception.	355
302-U4-Cs10	Scratchpad registers An Oracle SPARC Architecture processor includes eight privileged Scratchpad registers (64 bits each, read/write accessible).	438
303-U4-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
305-U4-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—

TABLE B-2 Oracle SPARC Architecture Implementation Dependencies (4 of 6)

Nbr	Description	Page
306-U4-Cs10	Trap type generated upon attempted access to noncacheable page with LDTXA When an LDTXA instruction attempts access from an address that is not mapped to cacheable memory space, a <i>DAE_nc_page</i> exception is generated.	263
307-U4-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
308-U3-Cs10	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
309-U4-Cs10	<i>Reserved.</i>	—
310-U4	Large page sizes Which, if any, of the following optional page sizes are supported by the MMU in an Oracle SPARC Architecture implementation is implementation dependent: 512 KBytes, 32 MBytes, 256 MBytes, 2 GBytes, and 16 GBytes.	471
311–319	<i>Reserved.</i>	
327–399	<i>Reserved</i>	
400-S10	Global Level register (GL) implementation Although GL is defined as a 4-bit register, an implementation may implement any subset of those bits sufficient to encode the values from 0 to <i>MAXPGL</i> for that implementation. If any bits of GL are not implemented, they read as zero and writes to them are ignored.	75
401-S10	Maximum Global Level (<i>MAXPGL</i>) The architectural parameter <i>MAXPGL</i> is a constant for each implementation; its legal values are from 2 to 15 (supporting from 3 to 16 sets of global registers). In a typical implementation <i>MAXPGL</i> = <i>MAXPTL</i> (see impl. dep. #101-V9-CS10). Architecturally, <i>MAXPTL</i> must be ≥ 2 .	74, 75
403-S10	Setting of “dirty” bits in FPRS A “dirty” bit (du or dl) in the FPRS register must be set to ‘1’ if any of its corresponding F registers is actually modified. If an instruction that normally writes to an F register is executed and causes an <i>fp_disabled</i> exception, FPRS.du and FPRS.dl are unchanged. Beyond that, the specific conditions under which a dirty bit is set are implementation dependent.	54, 54
404-S10	Scratchpad registers 4 through 7 The degree to which Scratchpad registers 4–7 are accessible to privileged software is implementation dependent. Each may be (1) fully accessible, (2) accessible, with access much slower than to scratchpad register 0–3, or (3) inaccessible (cause a <i>DAE_invalid_asi</i> exception).	438
405-S10	Virtual address range An Oracle SPARC Architecture implementation may support a full 64-bit virtual address space or a more limited range of virtual addresses. In an implementation that does not support a full 64-bit virtual address space, the supported range of virtual addresses is restricted to two equal-sized ranges at the extreme upper and lower ends of 64-bit addresses; that is, for <i>n</i> -bit virtual addresses, the valid address ranges are 0 to $2^{n-1} - 1$ and $2^{64} - 2^{n-1}$ to $2^{64} - 1$. (see also impl. dep. #451-S20)	18
409-S10	FLUSH instruction and memory consistency The implementation of the FLUSH instruction is implementation dependent. If the implementation automatically maintains consistency between instruction and data memory, (1) the FLUSH address is ignored and (2) the FLUSH instruction cannot cause any data access exceptions, because its effective address operand is not translated or used by the MMU. On the other hand, if the implementation does <i>not</i> maintain consistency between instruction and data memory, the FLUSH address is used to access the MMU and the FLUSH instruction can cause data access exceptions.	172

TABLE B-2 Oracle SPARC Architecture Implementation Dependencies (5 of 6)

Nbr	Description	Page
410-S10	<p data-bbox="483 222 704 243">Block Load behavior</p> <p data-bbox="483 247 1341 302">The following aspects of the behavior of block load (LDBLOCKF^D) instructions are implementation dependent:</p> <ul data-bbox="483 306 1409 747" style="list-style-type: none"> <li data-bbox="483 306 1409 361">• What memory ordering model is used by LDBLOCKF^D (LDBLOCKF^D is not required to follow TSO memory ordering) <li data-bbox="483 365 1409 441">• Whether LDBLOCKF^D follows memory ordering with respect to stores (including block stores), including whether the virtual processor detects read-after-write and write-after-read hazards to overlapping addresses <li data-bbox="483 445 1409 499">• Whether LDBLOCKF^D appears to execute out of order, or follow LoadLoad ordering (with respect to older loads, younger loads, and other LDBLOCKFs) <li data-bbox="483 504 1409 558">• Whether LDBLOCKF^D follows register-dependency interlocks, as do ordinary load instructions <li data-bbox="483 562 1409 651">• Whether the MMU ignores the side-effect bit (TTE.e) for LDBLOCKF^D accesses (in which case, LDBLOCKFs behave as if TTE.e = 0) <li data-bbox="483 693 1409 747">• Whether <i>VA_watchpoint</i> exceptions are recognized on accesses to all 64 bytes of a LDBLOCKF^D (the recommended behavior), or only on accesses to the first eight bytes 	244
411-S10	<p data-bbox="483 764 704 785">Block Store behavior</p> <p data-bbox="483 789 1341 844">The following aspects of the behavior of block store (STBLOCKF^D) instructions are implementation dependent:</p> <ul data-bbox="483 848 1409 1230" style="list-style-type: none"> <li data-bbox="483 848 1409 903">• The memory ordering model that STBLOCKF^D follows (other than as constrained by the rules outlined on page 338). <li data-bbox="483 907 1409 961">• Whether <i>VA_watchpoint</i> exceptions are recognized on accesses to all 64 bytes of a STBLOCKF^D (the recommended behavior), or only on accesses to the first eight bytes. <li data-bbox="483 966 1409 1020">• Whether STBLOCKFs to non-cacheable pages execute in strict program order or not. If not, a STBLOCKF^D to a non-cacheable page causes a <i>DAE_nc_page</i> exception. <li data-bbox="483 1024 1409 1079">• Whether STBLOCKF^D follows register dependency interlocks (as ordinary stores do). <li data-bbox="483 1083 1409 1138">• Whether a non-Commit STBLOCKF^D forces the data to be written to memory and invalidates copies in all caches present (as the Commit variants of STBLOCKF do). <li data-bbox="483 1142 1409 1197">• Whether the MMU ignores the side-effect bit (TTE.e) for STBLOCKF^D accesses (in which case, STBLOCKFs behave as if TTE.e = 0) <li data-bbox="483 1201 1409 1230">• Any other restrictions on the behavior of STBLOCKF^D, as described in implementation-specific documentation. 	338, 339
412-S10	<p data-bbox="483 1247 695 1268">MEMBAR behavior</p> <p data-bbox="483 1272 1321 1331">An Oracle SPARC Architecture implementation may define the operation of each MEMBAR variant in any manner that provides the required semantics.</p>	270
413-S10	<p data-bbox="483 1348 867 1369">Load Twin Extended Word behavior</p> <p data-bbox="483 1373 1386 1453">It is implementation dependent whether <i>VA_watchpoint</i> exceptions are recognized on accesses to all 16 bytes of a LDTXA instruction (the recommended behavior) or only on accesses to the first 8 bytes.</p>	263
414	<i>Reserved.</i>	—
415-S10	<p data-bbox="483 1507 737 1528">Size of ContextID fields</p> <p data-bbox="483 1533 1398 1591">The size of context ID fields in MMU context registers is implementation-dependent and may range from 13 to 16 bits.</p>	483
417-S10	<p data-bbox="483 1608 1166 1629">Behavior of DONE and RETRY when TSTATE[TL].pstate.am = 1</p> <p data-bbox="483 1633 1409 1749">If (1) TSTATE[TL].pstate.am = 1 and (2) a DONE or RETRY instruction is executed (which sets PSTATE.am to '1' by restoring the value from TSTATE[TL].pstate.am to PSTATE.am), it is implementation dependent whether the DONE or RETRY instruction masks (zeroes) the more-significant 32 bits of the values it places into PC and NPC.</p>	73, 147318

TABLE B-2 Oracle SPARC Architecture Implementation Dependencies (6 of 6)

Nbr	Description	Page
442-S10	STICK register a: If an accurate count cannot always be returned when STICK is read, any inaccuracy should be small, bounded, and documented. b: <i>Implementation dependency 442-S10(b) no longer applies, as of Oracle SPARC Architecture 2015.</i> c: The bit number m of the least significant implemented bit of the counter field of the STICK register is implementation-dependent, but must be in range 4 to 0, inclusively (that is, $4 \geq m \geq 0$). Any low-order bits not implemented must read as zero.	57
443-S10	(this implementation dependency is not in use)	—
444–449	<i>Reserved for UltraSPARC Architecture 2005</i>	
450-S20	Availability of <i>control_transfer_instruction</i> exception feature Availability of the <i>control_transfer_instruction</i> exception feature is implementation dependent. If not implemented, trap type 074 ₁₆ is unused, PSTATE.tct always reads as zero, and writes to PSTATE.tct are ignored.	70,460
451-S20	Width of Virtual Addresses supported The width of the virtual address supported is implementation dependent. If fewer than 64 bits are supported, then the unsupported bits must have the same value as the most significant supported bit. For example, if the model supports 48 virtual address bits, then bits 63:48 must have the same value as bit 47. (see also impl. dep. #405-S10)	471,
452-S20	Width of Real Addresses supported The number of real address (RA) bits supported is implementation dependent. A minimum of 40 bits and maximum of 56 bits can be provided for real addresses (RA). See implementation-specific documentation for details.	471
453-S20	Unified vs. Split Instruction and Data MMUs It is implementation dependent whether there is a unified MMU (UMMU) or a separate IMMU (for instruction accesses) and DMMU (for data accesses). The Oracle SPARC Architecture supports both configurations.	471
453-S20	Unified vs. Split Instruction and Data MMUs It is implementation dependent whether there is a unified MMU (UMMU) or a separate IMMU (for instruction accesses) and DMMU (for data accesses). The Oracle SPARC Architecture supports both configurations.	471
454-499	<i>Reserved for UltraSPARC Architecture 2007</i>	
502-S30	Maximum Value of PAUSE register (<i>MAX_PAUSE_COUNT</i>) The maximum number of nanoseconds that a virtual processor may be paused by the PAUSE instruction, <i>MAX_PAUSE_COUNT</i> , is $2^B - 1$. B , the number of bits used to implement the <i>pause_count</i> field, is implementation-dependent but B must be ≥ 13 . Therefore $MAX_PAUSE_COUNT \geq 2^{13} - 1$ (8191) nanoseconds. If a PAUSE instruction requests a pause of more than <i>MAX_PAUSE_COUNT</i> nanoseconds, the virtual processor must saturate the pause count to its maximum value, that is, write a value of <i>MAX_PAUSE_COUNT</i> to the PAUSE register.	60
504–549	<i>Reserved for Oracle SPARC Architecture 2011</i>	
553–599	<i>Reserved for Oracle SPARC Architecture 2015</i>	
600 and up	<i>Reserved for future use</i>	

Assembly Language Syntax

This appendix supports Chapter 7, *Instructions*. Each instruction description in Chapter 7 includes a table that describes the suggested assembly language format for that instruction. This appendix describes the notation used in those assembly language syntax descriptions and lists some synthetic instructions provided by Oracle SPARC Architecture assemblers for the convenience of assembly language programmers.

The appendix contains these sections:

- **Notation Used** on page 511.
- **Syntax Design** on page 516.
- **Synthetic Instructions** on page 516.

C.1 Notation Used

The notations defined here are also used in the assembly language syntax descriptions in Chapter 7, *Instructions*.

Items in *typewriter font* are literals to be written exactly as they appear. Items in *italic font* are metasyms that are to be replaced by numeric or symbolic values in actual SPARC V9 assembly language code. For example, "*imm_asi*" would be replaced by a number in the range 0 to 255 (the value of the *imm_asi* bits in the binary instruction) or by a symbol bound to such a number.

Subscripts on metasyms further identify the placement of the operand in the generated binary instruction. For example, *reg_{rs2}* is a *reg* (register name) whose binary value will be placed in the *rs2* field of the resulting instruction.

C.1.1 Register Names

reg. A *reg* is an integer register name. It can have any of the following values:¹

```
%r0-%r31
%g0-%g7 (global registers; same as %r0-%r7)
%o0-%o7 (out registers; same as %r8-%r15)
%l0-%l7 (local registers; same as %r16-%r23)
%i0-%i7 (in registers; same as %r24-%r31)
%fp      (frame pointer; conventionally same as %i6)
%sp      (stack pointer; conventionally same as %o6)
```

Subscripts identify the placement of the operand in the binary instruction as one of the following:

¹ In actual usage, the *%sp*, *%fp*, *%gn*, *%on*, *%ln*, and *%in* forms are preferred over *%rn*.

reg_{rs1} (rs1 field)
reg_{rs2} (rs2 field)
reg_{rd} (rd field)

freg. An *freg* is a floating-point register name. It may have the following values:

%f0, %f1, %f2, ... %f31
%f32, %f34, ... %f60, %f62 (even-numbered only, from *%f32* to *%f62*)
%d0, %d2, %d4, ... %d60, %d62 (*%dn*, where *n mod 2 = 0*, only)
%q0, %q4, %q8, ... %q56, %q60 (*%qn*, where *n mod 4 = 0*, only)

See *Floating-Point Registers* on page 38 for a detailed description of how the single-precision, double-precision, and quad-precision floating-point registers overlap.

Subscripts further identify the placement of the operand in the binary instruction as one of the following:

freg_{rs1} (rs1 field)
freg_{rs2} (rs2 field)
freg_{rs3} (rs3 field)
freg_{rd} (rd field)

asr_reg. An *asr_reg* is an Ancillary State Register name. It may have one of the following values:

%asr16–%asr31

Subscripts further identify the placement of the operand in the binary instruction as one of the following:

asr_reg_{rs1} (rs1 field)
asr_reg_{rd} (rd field)

i_or_x_cc. An *i_or_x_cc* specifies a set of integer condition codes, those based on either the 32-bit result of an operation (*icc*) or on the full 64-bit result (*xcc*). It may have either of the following values:

%icc
%xcc

fccn. An *fccn* specifies a set of floating-point condition codes. It can have any of the following values:

%fcc0
%fcc1
%fcc2
%fcc3

C.1.2 Special Symbol Names

Certain special symbols appear in the syntax table in typewriter font. They must be written exactly as they are shown, including the leading percent sign (%).

The symbol names and the registers or operators to which they refer are as follows:

<i>%asi</i>	Address Space Identifier (ASI) register
<i>%canrestore</i>	Restorable Windows register
<i>%cansave</i>	Savable Windows register
<i>%ccr</i>	Condition Codes register
<i>%cleanwin</i>	Clean Windows register
<i>%cwp</i>	Current Window Pointer (CWP) register

<code>%fprs</code>	Floating-Point Registers State (FPRS) register
<code>%fsr</code>	Floating-Point State register
<code>%gsr</code>	General Status Register (GSR)
<code>%otherwin</code>	Other Windows (OTHERWIN) register
<code>%pc</code>	Program Counter (PC) register
<code>%pil</code>	Processor Interrupt Level register
<code>%pstate</code>	Processor State register
<code>%softint</code>	Soft Interrupt register
<code>%softint_clr</code>	Soft Interrupt register (clear selected bits)
<code>%softint_set</code>	Soft Interrupt register (set selected bits)
<code>%stick †</code>	System Timer (STICK) register
<code>%stick_cmp †</code>	System Timer Compare (STICK_CMPR) register
<code>%tba</code>	Trap Base Address (TBA) register
<code>%tick</code>	Cycle count (TICK) register
<code>%tl</code>	Trap Level (TL) register
<code>%tnpc</code>	Trap Next Program Counter (TNPC) register
<code>%tpc</code>	Trap Program Counter (TPC) register
<code>%tstate</code>	Trap State (TSTATE) register
<code>%tt</code>	Trap Type (TT) register
<code>%wstate</code>	Window State register
<code>%y</code>	Y register

† The original assembly language names for `%stick` and `%stick_cmp` were, respectively, `%sys_tick` and `%sys_tick_cmp`, which are now deprecated. Over time, assemblers will support the new `%stick` and `%stick_cmp` names for these registers (which are consistent with `%tick`). In the meantime, some existing assemblers may only recognize the original names.

The following special symbol names are prefix unary operators that perform the functions described, on an argument that is a constant, symbol, or expression that evaluates to a constant offset from a symbol:

<code>%hh</code>	Extracts bits 63:42 (high 22 bits of upper word) of its operand
<code>%hm</code>	Extracts bits 41:32 (low-order 10 bits of upper word) of its operand
<code>%hi</code> or <code>%lm</code>	Extracts bits 31:10 (high-order 22 bits of low-order word) of its operand
<code>%lo</code>	Extracts bits 9:0 (low-order 10 bits) of its operand

For example, the value of `"%lo(symbol)"` is the least-significant 10 bits of `symbol`.

Certain predefined value names appear in the syntax table in `typewriter` font. They must be written exactly as they are shown, including the leading sharp sign (#). The value names and the constant values to which they are bound are listed in TABLE C-1.

TABLE C-1 Value Names and Values (1 of 2)

Value Name in Assembly Language	Value	Comments
<i>for PREFETCH instruction "fcr" field</i>		
<code>#n_reads</code>	0	
<code>#one_read</code>	1	
<code>#n_writes</code>	2	
<code>#one_write</code>	3	
<code>#page</code>	4	
<code>#unified</code>	17 (11 ₁₆)	
<code>#n_reads_strong</code>	20 (14 ₁₆)	
<code>#one_read_strong</code>	21 (15 ₁₆)	

TABLE C-1 Value Names and Values (2 of 2)

Value Name in Assembly Language	Value	Comments
#n_writes_strong	22 (16 ₁₆)	
#one_write_strong	23 (17 ₁₆)	
<i>for MEMBAR instruction "mmask" field</i>		
#LoadLoad	01 ₁₆	
#StoreLoad	02 ₁₆	
#LoadStore	04 ₁₆	
<i>for MEMBAR instruction "cmask" field</i>		
#StoreStore	08 ₁₆	
#Lookaside ^D	10 ₁₆	Use of #Lookaside is deprecated and only supported for legacy software. New software should use a slightly more restrictive MEMBAR operation (such as #StoreLoad) instead.
#MemIssue	20 ₁₆	
#Sync	40 ₁₆	

C.1.3 Values

Some instructions use operand values as follows:

<i>const4</i>	A constant that can be represented in 4 bits
<i>const22</i>	A constant that can be represented in 22 bits
<i>imm_asi</i>	An alternate address space identifier (0–255)
<i>siam_mode</i>	A 3-bit mode value for the SIAM instruction
<i>simm7</i>	A signed immediate constant that can be represented in 7 bits
<i>simm8</i>	A signed immediate constant that can be represented in 8 bits
<i>simm10</i>	A signed immediate constant that can be represented in 10 bits
<i>simm11</i>	A signed immediate constant that can be represented in 11 bits
<i>simm13</i>	A signed immediate constant that can be represented in 13 bits
<i>value</i>	Any 64-bit value
<i>shcnt32</i>	A shift count from 0–31
<i>shcnt64</i>	A shift count from 0–63

C.1.4 Labels

A label is a sequence of characters that comprises alphabetic letters (a–z, A–Z [with upper and lower case distinct]), underscores (_), dollar signs (\$), periods (.), and decimal digits (0–9). A label may contain decimal digits, but it may not begin with one. A local label contains digits only.

C.1.5 Other Operand Syntax

Some instructions allow several operand syntaxes, as follows:

reg_plus_imm Can be any of the following:

<i>reg_{rs1}</i>	(equivalent to <i>reg_{rs1}</i> + %g0)
<i>reg_{rs1}</i> + <i>simm13</i>	

$reg_{rs1} - imm13$
 $imm13$ (equivalent to $\%g0 + imm13$)
 $imm13 + reg_{rs1}$ (equivalent to $reg_{rs1} + imm13$)

address Can be any of the following:

reg_{rs1} (equivalent to $reg_{rs1} + \%g0$)
 $reg_{rs1} + imm13$
 $reg_{rs1} - imm13$
 $imm13$ (equivalent to $\%g0 + imm13$)
 $imm13 + reg_{rs1}$ (equivalent to $reg_{rs1} + imm13$)
 $reg_{rs1} + reg_{rs2}$

membar_mask Is the following:

const7 A constant that can be represented in 7 bits. Typically, this is an expression involving the logical OR of some combination of #Lookaside^D, #MemIssue, #Sync, #StoreStore, #LoadStore, #StoreLoad, and #LoadLoad (see TABLE 7-13 and TABLE 7-14 on page 270 for a complete list of mnemonics).

prefetch_fcn (*prefetch function*) Can be any of the following:

0–31
Predefined constants (the values of which fall in the 0-31 range) useful as *prefetch_fcn* values can be found in TABLE C-1 on page 513.

regaddr (*register-only address*) Can be any of the following:

reg_{rs1} (equivalent to $reg_{rs1} + \%g0$)
 $reg_{rs1} + reg_{rs2}$

reg_or_imm (*register or immediate value*) Can be either of:

reg_{rs2}
 $imm13$

reg_or_imm5 (*register or immediate value*) Can be either of:

reg_{rs2}
 $imm5$

reg_or_imm10 (*register or immediate value*) Can be either of:

reg_{rs2}
 $imm10$

reg_or_imm11 (*register or immediate value*) Can be either of:

reg_{rs2}
 $imm11$

reg_or_shcnt (*register or shift count value*) Can be any of:

reg_{rs2}

shcnt32
shcnt64

software_trap_number Can be any of the following:

reg_{rs1} (equivalent to $reg_{rs1} + \%g0$)
 $reg_{rs1} + reg_{rs2}$
 $reg_{rs1} + simm8$
 $reg_{rs1} - simm8$
 $simm8$ (equivalent to $\%g0 + simm8$)
 $simm8 + reg_{rs1}$ (equivalent to $reg_{rs1} + simm8$)

The resulting operand value (software trap number) must be in the range 0–255, inclusive.

C.1.6 Comments

Two types of comments are accepted by the SPARC V9 assembler: C-style “/* . . . */” comments, which may span multiple lines, and “! . . .” comments, which extend from the “!” to the end of the line.

C.2 Syntax Design

The SPARC V9 assembly language syntax is designed so that the following statements are true:

- The destination operand (if any) is consistently specified as the last (rightmost) operand in an assembly language instruction.
- A reference to the *contents* of a memory location (for example, in a load, store, or load-store instruction) is always indicated by square brackets ([]); a reference to the *address* of a memory location (such as in a JMPL, CALL, or SETHI) is specified directly, without square brackets.

The follow additional syntax constraints have been adopted for Oracle SPARC Architecture:

- Instruction mnemonics should be limited to a maximum of 15 characters.

C.3 Synthetic Instructions

TABLE C-2 describes the mapping of a set of synthetic (or “pseudo”) instructions to actual instructions. These synthetic instructions are provided by the SPARC V9 assembler for the convenience of assembly language programmers.

Note: Synthetic instructions should not be confused with “pseudo ops,” which typically provide information to the assembler but do not generate instructions. Synthetic instructions always generate instructions; they provide more mnemonic syntax for standard SPARC V9 instructions.

TABLE C-2 Mapping Synthetic to SPARC V9 Instructions (1 of 3)

Synthetic Instruction	SPARC V9 Instruction(s)	Comment
cmp <i>reg_{rs1}, reg_{or_imm}</i>	subcc <i>reg_{rs1}, reg_{or_imm}, %g0</i>	Compare.
jmp <i>address</i>	jmp1 <i>address, %g0</i>	
call <i>address</i>	jmp1 <i>address, %o7</i>	

TABLE C-2 Mapping Synthetic to SPARC V9 Instructions (2 of 3)

Synthetic Instruction	SPARC V9 Instruction(s)	Comment
call <i>address, reg_{rd}</i>	jmp1 <i>address, reg_{rd}</i>	
iprefetch <i>label</i>	bn, a, pt %xcc, <i>label</i>	Originally envisioned as an encoding for an "instruction prefetch" operation, but functions as a NOP on all Oracle SPARC Architecture implementations. (See PREFETCH function 17 on page 303 for an alternative method of prefetching instructions.)
tst <i>reg_{rs1}</i>	orcc %g0, <i>reg_{rs1}</i> , %g0	Test.
ret	jmp1 %i7+8, %g0	Return from subroutine.
retl	jmp1 %o7+8, %g0	Return from leaf subroutine.
restore	restore %g0, %g0, %g0	Trivial RESTORE.
save	save %g0, %g0, %g0	Trivial SAVE. (Warning: trivial SAVE should only be used in kernel code!)
setuw <i>value, reg_{rd}</i>	sethi %hi(<i>value</i>), <i>reg_{rd}</i> — or — or %g0, <i>value</i> , <i>reg_{rd}</i> — or —	(When ((<i>value</i> &3FF ₁₆) == 0).) (When $0 \leq \textit{value} \leq 4095$).
	sethi %hi(<i>value</i>), <i>reg_{rd}</i> ; or <i>reg_{rd}</i> , %lo(<i>value</i>), <i>reg_{rd}</i>	(Otherwise) Warning: do not use setuw in the delay slot of a DCTI.
set <i>value, reg_{rd}</i>		synonym for setuw.
setsw <i>value, reg_{rd}</i>	sethi %hi(<i>value</i>), <i>reg_{rd}</i> — or — or %g0, <i>value</i> , <i>reg_{rd}</i> — or — sethi %hi(<i>value</i>), <i>reg_{rd}</i> sra <i>reg_{rd}</i> , %g0, <i>reg_{rd}</i> — or — sethi %hi(<i>value</i>), <i>reg_{rd}</i> ; or <i>reg_{rd}</i> , %lo(<i>value</i>), <i>reg_{rd}</i> — or — sethi %hi(<i>value</i>), <i>reg_{rd}</i> ; or <i>reg_{rd}</i> , %lo(<i>value</i>), <i>reg_{rd}</i> sra <i>reg_{rd}</i> , %g0, <i>reg_{rd}</i>	(When (<i>value</i> >= 0) and ((<i>value</i> & 3FF ₁₆) == 0).) (When $4096 \leq \textit{value} \leq 4095$). (Otherwise, if (<i>value</i> < 0) and ((<i>value</i> & 3FF ₁₆) = = 0)) (Otherwise, if <i>value</i> = 0) (Otherwise, if <i>value</i> < 0) Warning: do not use setsw in the delay slot of a CTI.
setx <i>value, reg, reg_{rd}</i>	sethi %hh(<i>value</i>), <i>reg</i> or <i>reg</i> , %hm(<i>value</i>), <i>reg</i> sllx <i>reg</i> , 32, <i>reg</i> sethi %hi(<i>value</i>), <i>reg_{rd}</i> or <i>reg_{rd}</i> , <i>reg</i> , <i>reg_{rd}</i> or <i>reg_{rd}</i> , %lo(<i>value</i>), <i>reg_{rd}</i>	Create 64-bit constant. ("reg" is used as a temporary register.) Note: setx optimizations are possible but not enumerated here. The worst case is shown. Warning: do not use setx in the delay slot of a CTI.

TABLE C-2 Mapping Synthetic to SPARC V9 Instructions (3 of 3)

Synthetic Instruction	SPARC V9 Instruction(s)	Comment	
signx	<i>reg_{rs1}</i> , <i>reg_{rd}</i>	sra <i>reg_{rs1}</i> , %g0, <i>reg_{rd}</i>	Sign-extend 32-bit value to 64 bits.
signx	<i>reg_{rd}</i>	sra <i>reg_{rd}</i> , %g0, <i>reg_{rd}</i>	
not	<i>reg_{rs1}</i> , <i>reg_{rd}</i>	xnor <i>reg_{rs1}</i> , %g0, <i>reg_{rd}</i>	One's complement.
not	<i>reg_{rd}</i>	xnor <i>reg_{rd}</i> , %g0, <i>reg_{rd}</i>	One's complement.
neg	<i>reg_{rs2}</i> , <i>reg_{rd}</i>	sub %g0, <i>reg_{rs2}</i> , <i>reg_{rd}</i>	Two's complement.
neg	<i>reg_{rd}</i>	sub %g0, <i>reg_{rd}</i> , <i>reg_{rd}</i>	Two's complement.
cas	[<i>reg_{rs1}</i>], <i>reg_{rs2}</i> , <i>reg_{rd}</i>	casa [<i>reg_{rs1}</i>]#ASI_P, <i>reg_{rs2}</i> , <i>reg_{rd}</i>	Compare and swap.
casl	[<i>reg_{rs1}</i>], <i>reg_{rs2}</i> , <i>reg_{rd}</i>	casa [<i>reg_{rs1}</i>]#ASI_P_L, <i>reg_{rs2}</i> , <i>reg_{rd}</i>	Compare and swap, little-endian.
casx	[<i>reg_{rs1}</i>], <i>reg_{rs2}</i> , <i>reg_{rd}</i>	casxa [<i>reg_{rs1}</i>]#ASI_P, <i>reg_{rs2}</i> , <i>reg_{rd}</i>	Compare and swap extended.
casxl	[<i>reg_{rs1}</i>], <i>reg_{rs2}</i> , <i>reg_{rd}</i>	casxa [<i>reg_{rs1}</i>]#ASI_P_L, <i>reg_{rs2}</i> , <i>reg_{rd}</i>	Compare and swap extended, little-endian.
inc	<i>reg_{rd}</i>	add <i>reg_{rd}</i> , 1, <i>reg_{rd}</i>	Increment by 1.
inc	<i>const13</i> , <i>reg_{rd}</i>	add <i>reg_{rd}</i> , <i>const13</i> , <i>reg_{rd}</i>	Increment by <i>const13</i> .
inccc	<i>reg_{rd}</i>	addcc <i>reg_{rd}</i> , 1, <i>reg_{rd}</i>	Increment by 1; set icc & xcc.
inccc	<i>const13</i> , <i>reg_{rd}</i>	addcc <i>reg_{rd}</i> , <i>const13</i> , <i>reg_{rd}</i>	Incr by <i>const13</i> ; set icc & xcc.
dec	<i>reg_{rd}</i>	sub <i>reg_{rd}</i> , 1, <i>reg_{rd}</i>	Decrement by 1.
dec	<i>const13</i> , <i>reg_{rd}</i>	sub <i>reg_{rd}</i> , <i>const13</i> , <i>reg_{rd}</i>	Decrement by <i>const13</i> .
deccc	<i>reg_{rd}</i>	subcc <i>reg_{rd}</i> , 1, <i>reg_{rd}</i>	Decrement by 1; set icc & xcc.
deccc	<i>const13</i> , <i>reg_{rd}</i>	subcc <i>reg_{rd}</i> , <i>const13</i> , <i>reg_{rd}</i>	Decr by <i>const13</i> ; set icc & xcc.
btst	<i>reg_or_imm</i> , <i>reg_{rs1}</i>	andcc <i>reg_{rs1}</i> , <i>reg_or_imm</i> , %g0	Bit test.
bset	<i>reg_or_imm</i> , <i>reg_{rd}</i>	or <i>reg_{rd}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	Bit set.
bclr	<i>reg_or_imm</i> , <i>reg_{rd}</i>	andn <i>reg_{rd}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	Bit clear.
btog	<i>reg_or_imm</i> , <i>reg_{rd}</i>	xor <i>reg_{rd}</i> , <i>reg_or_imm</i> , <i>reg_{rd}</i>	Bit toggle.
clr	<i>reg_{rd}</i>	or %g0, %g0, <i>reg_{rd}</i>	Clear (zero) register.
clrb	[<i>address</i>]	stb %g0, [<i>address</i>]	Clear byte.
clrh	[<i>address</i>]	sth %g0, [<i>address</i>]	Clear half-word.
clr	[<i>address</i>]	stw %g0, [<i>address</i>]	Clear word.
clrx	[<i>address</i>]	stx %g0, [<i>address</i>]	Clear extended word.
clruw	<i>reg_{rs1}</i> , <i>reg_{rd}</i>	srl <i>reg_{rs1}</i> , %g0, <i>reg_{rd}</i>	Copy and clear upper word.
clruw	<i>reg_{rd}</i>	srl <i>reg_{rd}</i> , %g0, <i>reg_{rd}</i>	Clear upper word.
mov	<i>reg_or_imm</i> , <i>reg_{rd}</i>	or %g0, <i>reg_or_imm</i> , <i>reg_{rd}</i>	
mov	%y, <i>reg_{rd}</i>	rd %y, <i>reg_{rd}</i>	
mov	%asrn, <i>reg_{rd}</i>	rd %asrn, <i>reg_{rd}</i>	
mov	<i>reg_or_imm</i> , %y	wr %g0, <i>reg_or_imm</i> , %y	
mov	<i>reg_or_imm</i> , %asrn	wr %g0, <i>reg_or_imm</i> , %asrn	

Index

A

- a (annul) instruction field
 - branch instructions, 123, 126, 128, 157, 160
- accesses
 - cacheable, 408
 - I/O, 408
 - restricted ASI, 411
 - with side effects, 408, 416
- accrued exception (aexc) field of FSR register, 46, 450, 500
- ADD instruction, 110
- ADDC instruction, 110
- ADDcc instruction, 110, 332
- ADDCCc instruction, 110
- address
 - operand syntax, 515
 - space identifier (ASI), 423
- address mask (am) field of PSTATE register
 - description, 71
- address space, 5, 13
- address space identifier (ASI), 5, 407
 - accessing MMU registers, 482
 - appended to memory address, 17, 77
 - architecturally specified, 411
 - changed in UA
 - ASI_REAL, 438
 - ASI_REAL_IO, 438
 - ASI_REAL_IO_LITTLE, 438
 - ASI_REAL_LITTLE, 438
 - definition, 5, 8
 - encoding address space information, 79
 - explicit, 83
 - explicitly specified in instruction, 84
 - implicit, *See* implicit ASIs
 - nontranslating, 9, 260, 355
 - nontranslating ASI, 424
 - with prefetch instructions, 304
 - real ASI, 424
 - restricted, 411, 423
 - privileged, 411
 - restriction indicator, 51
 - SPARC V9 address, 409
 - translating ASI, 424
 - unrestricted, 411, 423
- address space identifier (ASI) register
 - for load/store alternate instructions, 51
 - address for explicit ASI, 83
 - and LDDA instruction, 248, 259
 - and LDSTUBA instruction, 256
 - load integer from alternate space instructions, 240
 - with prefetch instructions, 304
 - for register-immediate addressing, 411
 - restoring saved state, 147, 318
 - saving state, 443
 - and STDA instruction, 354
 - store floating-point into alternate space instructions, 342
 - store integer to alternate space instructions, 335
 - and SWAPA instruction, 360
 - after trap, 21
 - and TSTATE register, 68
 - and write state register instructions, 374
- addressing modes, 14
- ADDX instruction, 111, 112, 131, 133, 143, 213, 268
- ADDX instruction (SPARC V8), 110
- ADDXC instruction, 112, 131, 133, 143, 213, 268
- ADDXcc instruction, 111, 358
- ADDXcc instruction (SPARC V8), 110
- alias
 - floating-point registers, 38
- aliased, 5
- ALIGNADDRESS instruction, 117
- ALIGNADDRESS_LITTLE instruction, 117
- alignment
 - data (load/store), 18, 79, 409
 - doubleword, 18, 79, 409
 - extended-word, 79
 - halfword, 18, 79, 409
 - instructions, 18, 79, 409
 - integer registers, 250, 257, 260
 - memory, 409, 463
 - quadword, 18, 79, 409
 - word, 18, 79, 409
- ALLCLEAN instruction, 118
- alternate space instructions, 19, 51
- ancillary state registers (ASRs)
 - access, 48
 - assembly language syntax, 512
 - I/O register access, 19
 - possible registers included, 311, 374
 - privileged, 20, 499
 - reading/writing implementation-dependent processor

- registers, 20, 499
- writing to, 373
- AND instruction, 118
- ANDcc instruction, 118
- ANDN instruction, 118
- ANDNcc instruction, 118
- annul bit
 - in branch instructions, 128
 - in conditional branches, 158
- annulled branches, 128
- application program, 5, 48
- architectural direction note, 4
- architecture, meaning for SPARC V9, 13
- arithmetic overflow, 51
- ARRAY16 instruction, 120
- ARRAY32 instruction, 120
- ARRAY8 instruction, 120
- ASI, 5
 - invalid, and *DAE_invalid_asi*, 461
- ASI register, 49
- ASI, *See* address space identifier (ASI)
- ASI_AIUP, 425, 433
- ASI_AIUPL, 425, 433
- ASI_AIUS, 425, 433
- ASI_AIUS_L, 262
- ASI_AIUSL, 425, 433
- ASI_AS_IF_USER*, 72, 475
- ASI_AS_IF_USER_NONFAULT_LITTLE, 412
- ASI_AS_IF_USER_PRIMARY, 425, 433, 461, 474
- ASI_AS_IF_USER_PRIMARY_LITTLE, 412, 425, 433, 461
- ASI_AS_IF_USER_SECONDARY, 412, 425, 433, 461, 474
- ASI_AS_IF_USER_SECONDARY_LITTLE, 412, 425, 433, 461
- ASI_AS_IF_USER_SECONDARY_NOFAULT_LITTLE, 412
- ASI_BLK_AIUP, 425, 433
- ASI_BLK_AIUPL, 425, 433
- ASI_BLK_AIUS, 425, 433
- ASI_BLK_AIUSL, 425, 433
- ASI_BLK_COMMIT_P, 430
- ASI_BLK_COMMIT_S, 430
- ASI_BLK_P, 430
- ASI_BLK_PL, 431
- ASI_BLK_S, 431
- ASI_BLK_SL, 431
- ASI_BLOCK_AS_IF_USER_PRIMARY, 425, 433
- ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE, 425, 433
- ASI_BLOCK_AS_IF_USER_SECONDARY, 425, 433
- ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE, 425, 433
- ASI_BLOCK_COMMIT_PRIMARY, 430
- ASI_BLOCK_COMMIT_SECONDARY, 430
- ASI_BLOCK_PRIMARY, 430
- ASI_BLOCK_PRIMARY_LITTLE, 431
- ASI_BLOCK_SECONDARY, 431
- ASI_BLOCK_SECONDARY_LITTLE, 431
- ASI_FL16_P, 429
- ASI_FL16_PL, 430
- ASI_FL16_PRIMARY, 429
- ASI_FL16_PRIMARY_LITTLE, 430
- ASI_FL16_S, 430
- ASI_FL16_SECONDARY, 430
- ASI_FL16_SECONDARY_LITTLE, 430
- ASI_FL16_SL, 430
- ASI_FL8_P, 429
- ASI_FL8_PL, 430
- ASI_FL8_PRIMARY, 429
- ASI_FL8_PRIMARY_LITTLE, 430
- ASI_FL8_S, 429
- ASI_FL8_SECONDARY, 429
- ASI_FL8_SECONDARY_LITTLE, 430
- ASI_FL8_SL, 430
- ASI_MAIUP, 425
- ASI_MAIUS, 425
- ASI_MMU, 426
- ASI_MMU_CONTEXTID, 426
- ASI_MONITOR_AS_IF_USER_PRIMARY, 425
- ASI_MONITOR_AS_IF_USER_SECONDARY, 425
- ASI_MONITOR_PRIMARY, 428
- ASI_MONITOR_SECONDARY, 428
- ASI_MP, 428
- ASI_MS, 428
- ASI_N, 425
- ASI_NL, 425
- ASI_NUCLEUS, 83, 84, 425, 474
- ASI_NUCLEUS_LITTLE, 84, 425
- ASI_P, 428
- ASI_PHY_BYPASS_EC_WITH_EBIT_L, 438
- ASI_PHYS_BYPASS_EC_WITH_EBIT, 438
- ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE, 438
- ASI_PHYS_USE_EC, 438
- ASI_PHYS_USE_EC_L, 438
- ASI_PHYS_USE_EC_LITTLE, 438
- ASI_PIC, 429
- ASI_PL, 428
- ASI_PNF, 428
- ASI_PNFL, 428
- ASI_PRIMARY, 83, 411, 412, 428
- ASI_PRIMARY_LITTLE, 84, 411, 428, 480
- ASI_PRIMARY_NO_FAULT, 409, 421, 428
- ASI_PRIMARY_NO_FAULT_LITTLE, 409, 421, 428, 461
- ASI_PRIMARY_NOFAULT_LITTLE, 412
- ASI_PST16_P, 347, 429
- ASI_PST16_PL, 347, 429
- ASI_PST16_PRIMARY, 429
- ASI_PST16_PRIMARY_LITTLE, 429
- ASI_PST16_S, 347, 429
- ASI_PST16_SECONDARY, 429
- ASI_PST16_SECONDARY_LITTLE, 429
- ASI_PST16_SL, 347
- ASI_PST32_P, 347, 429
- ASI_PST32_PL, 347, 429
- ASI_PST32_PRIMARY, 429
- ASI_PST32_PRIMARY_LITTLE, 429
- ASI_PST32_S, 347, 429
- ASI_PST32_SECONDARY, 429
- ASI_PST32_SECONDARY_LITTLE, 429
- ASI_PST32_SL, 347, 429
- ASI_PST8_P, 429
- ASI_PST8_PL, 429
- ASI_PST8_PRIMARY, 429
- ASI_PST8_PRIMARY_LITTLE, 429

ASI_PST8_S, 429
 ASI_PST8_SECONDARY, 429
 ASI_PST8_SECONDARY_LITTLE, 429
 ASI_PST8_SL, 347, 429
 ASI_QUAD_LDD_REAL (deprecated), 427
 ASI_QUAD_LDD_REAL_LITTLE (deprecated), 428
 ASI_REAL, 425, 433, 438
 ASI_REAL_IO, 425, 434, 438
 ASI_REAL_IO_L, 425
 ASI_REAL_IO_LITTLE, 425, 434, 438
 ASI_REAL_L, 425
 ASI_REAL_LITTLE, 425, 434, 438
 ASI_S, 428
 ASI_SECONDARY, 428, 474
 ASI_SECONDARY_LITTLE, 428
 ASI_SECONDARY_NO_FAULT, 421, 428, 461
 ASI_SECONDARY_NO_FAULT_LITTLE, 421, 428, 461
 ASI_SECONDARY_NOFAULT, 412
 ASI_SL, 428
 ASI_SNF, 428
 ASI_SNFL, 428
 ASI_ST_BLKINIT_AS_IF_USER_PRIMARY, 426
 ASI_ST_BLKINIT_AS_IF_USER_PRIMARY_LITTLE, 427
 ASI_ST_BLKINIT_AS_IF_USER_SECONDARY, 426
 ASI_ST_BLKINIT_AS_IF_USER_SECONDARY_LITTLE, 4
 27
 ASI_ST_BLKINIT_NUCLEUS, 427
 ASI_ST_BLKINIT_NUCLEUS_LITTLE, 428
 ASI_ST_BLKINIT_PRIMARY, 430
 ASI_ST_BLKINIT_PRIMARY_LITTLE, 430
 ASI_ST_BLKINIT_REAL, 427
 ASI_ST_BLKINIT_REAL_LITTLE, 428
 ASI_ST_BLKINIT_SECONDARY, 430
 ASI_ST_BLKINIT_SECONDARY_LITTLE, 430
 ASI_STBI_AIUP, 426
 ASI_STBI_AIUPL, 427
 ASI_STBI_AIUS, 426
 ASI_STBI_AIUS_L, 427
 ASI_STBI_N, 427
 ASI_STBI_NL, 428
 ASI_STBI_P, 430
 ASI_STBI_PL, 430
 ASI_STBI_PLM, 431
 ASI_STBI_PM, 431
 ASI_STBI_RL, 428
 ASI_STBI_S, 430
 ASI_STBI_SL, 430
 ASI_STBI_SLM, 431
 ASI_STBI_SM, 431
 ASI_TWINK_AIUP, 262, 426, 435
 ASI_TWINK_AIUP_L, 262, 427, 435
 ASI_TWINK_AIUS, 262, 435
 ASI_TWINK_AIUS_L, 427, 435
 ASI_TWINK_AS_IF_USER_PRIMARY, 426, 435
 ASI_TWINK_AS_IF_USER_PRIMARY_LITTLE, 427, 435
 ASI_TWINK_AS_IF_USER_SECONDARY, 426, 435
 ASI_TWINK_AS_IF_USER_SECONDARY_LITTLE, 427, 435
 ASI_TWINK_N, 262, 427
 ASI_TWINK_NL, 262, 428, 435
 ASI_TWINK_NUCLEUS, 427, 435

ASI_TWINK_NUCLEUS[_L], 410
 ASI_TWINK_NUCLEUS_LITTLE, 428, 435
 ASI_TWINK_P, 262, 430
 ASI_TWINK_PL, 262, 430
 ASI_TWINK_PRIMARY, 430, 436
 ASI_TWINK_PRIMARY_LITTLE, 430, 436
 ASI_TWINK_R, 427, 435
 ASI_TWINK_REAL, 262, 427, 435
 ASI_TWINK_REAL[_L], 410
 ASI_TWINK_REAL_L, 428, 435
 ASI_TWINK_REAL_LITTLE, 428, 435
 ASI_TWINK_RL, 435
 ASI_TWINK_S, 262, 430
 ASI_TWINK_SECONDARY, 430, 436
 ASI_TWINK_SECONDARY_LITTLE, 430, 436
 ASI_TWINK_SL, 262, 430
 ASR, 5
asr_reg, 512
 atomic
 memory operations, 262, 418, 419
 store doubleword instruction, 352, 354
 store instructions, 334, 335
 atomic load-store instructions
 compare and swap, 134
 load-store unsigned byte, 255, 360
 load-store unsigned byte to alternate space, 256
 simultaneously addressing doublewords, 359
 swap R register with alternate space memory, 360
 swap R register with memory, 134, 359
 atomicity, 409, 505

B

BA instruction, 123, 491
 BCC instruction, 123, 491
 bclrg synthetic instruction, 518
 BCS instruction, 123, 491
 BE instruction, 123, 491
 Berkeley RISCs, 15
 BG instruction, 123, 491
 BGE instruction, 123, 491
 BGU instruction, 123, 491
 Bicc instructions, 123, 485
 big-endian, 5
 big-endian byte order, 18, 70, 80
 binary compatibility, 15
 BL instruction, 123, 491
 BLD, *See* LDBLOCKF instruction
 BLE instruction, 123, 491
 BLEU instruction, 123, 491
 block load instructions, 38, 243, 436
 block store instructions, 38, 337, 436
 with commit, 249, 338, 436
 blocked byte formatting, 120
 BMASK instruction, 125
 BN instruction, 123, 491
 BNE instruction, 123, 491
 BNEG instruction, 123, 491
 BP instructions, 491
 BPA instruction, 126, 491

- BPCC instruction, 126, 491
- BPcc instructions, 51, 126, 492
- BPCS instruction, 126, 491
- BPE instruction, 126, 491
- BPG instruction, 126, 491
- BPGE instruction, 126, 491
- BPGU instruction, 126, 491
- BPL instruction, 126, 491
- BPLE instruction, 126, 491
- BPLEU instruction, 126, 491
- BPn instruction, 126, 491
- BPNE instruction, 126, 491
- BPNEG instruction, 126, 491
- BPOS instruction, 123, 491
- BPPOS instruction, 126, 491
- BPr instructions, 128, 491
- BPVC instruction, 126, 491
- BPVS instruction, 126, 491
- branch
 - annulled, 128
 - delayed, 77
 - elimination, 89, 90
 - fcc-conditional, 157, 160
 - icc-conditional, 123
 - instructions
 - on floating-point condition codes, 157
 - on floating-point condition codes with prediction, 159
 - on integer condition codes with prediction (BPcc), 126
 - on integer condition codes, *See* Bicc instructions
 - when contents of integer register match condition, 128
 - prediction bit, 128
 - unconditional, 123, 127, 157, 160
 - with prediction, 14
- BRGEZ instruction, 128
- BRGZ instruction, 128
- BRLEZ instruction, 128
- BRLZ instruction, 128
- BRNZ instruction, 128
- BRZ instruction, 128
- bset synthetic instruction, 518
- BSHUFFLE instruction, 125
- BST, *See* STBLOCKF instruction
- btog synthetic instruction, 518
- btst synthetic instruction, 518
- BVC instruction, 123, 491
- BVS instruction, 123, 491
- byte, 5
 - addressing, 83
 - data format, 23
 - order, 18
 - order, big-endian, 18
 - order, little-endian, 18
- byte order
 - big-endian, 70
 - implicit, 70
 - in trap handlers, 449
 - little-endian, 70

C

- C*BCC instruction, 491
- C*BCS instruction, 491
- C*BE instruction, 491
- C*BG instruction, 491
- C*BGE instruction, 491
- C*BGU instruction, 491
- C*BL instruction, 491
- C*BLE instruction, 491
- C*BLEU instruction, 491
- C*BNE instruction, 491
- C*BNEG instruction, 491
- C*BPOS instruction, 491
- C*BVC instruction, 491
- C*BVS instruction, 491
- C8BL instruction, 136
- cache
 - coherency protocol, 408
 - data, 414
 - instruction, 414
 - miss, 308
 - nonconsistent instruction cache, 414
- cacheable accesses, 408
- caching, TSB, 479
- CALL instruction
 - description, 130
 - displacement, 20
 - does not change CWP, 35
 - and JMPL instruction, 238
 - writing address into R[15], 37
- call synthetic instruction, 516, 517
- CANRESTORE (restorable windows) register, 64
 - and *clean_window* exception, 90
 - and CLEANWIN register, 64, 66, 465
 - counting windows, 66
 - decremented by RESTORE instruction, 315
 - decremented by SAVED instruction, 324
 - detecting window underflow, 36
 - if registered window was spilled, 315
 - incremented by SAVE instruction, 322
 - modified by NORMALW instruction, 295
 - modified by OTHERW instruction, 297
 - range of values, 62
 - RESTORE instruction, 90
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - window underflow, 465
- CANSAVE (savable windows) register, 63
 - decremented by SAVE instruction, 322
 - detecting window overflow, 36
 - FLUSHW instruction, 174
 - if equals zero, 90
 - incremented by RESTORE, 315
 - incremented by SAVED instruction, 324
 - range of values, 62
 - SAVE instruction, 466
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - window overflow, 465
- CAS synthetic instruction, 419

- CASA instruction, **134**
 - 32-bit compare-and-swap, 418
 - alternate space addressing, 18
 - and *DAE_nc_page* exception, 461
 - atomic operation, 255
 - hardware primitives for mutual exclusion of CASXA, 418
 - in multiprocessor system, 256, 359, 360
 - R register use, 79
 - word access (memory), 79
- casn* synthetic instructions, 518
- CASX synthetic instruction, 418, 419
- CASXA instruction, **134**
 - 64-bit compare-and-swap, 418
 - alternate space addressing, 18
 - and *DAE_nc_page* exception, 461
 - atomic operation, 256
 - doubleword access (memory), 79
 - hardware primitives for mutual exclusion of CASA, 418
 - in multiprocessor system, 255, 256, 359, 360
 - R register use, 79
- catastrophic error exception, **443**
- CBcond instructions
 - opcode maps, 491
- cc0 instruction field
 - branch instructions, 126, 160
 - floating point compare instructions, 162
 - move instructions, 281, 491
- cc1 instruction field
 - branch instructions, 126, 160
 - floating point compare instructions, 162
 - move instructions, 281, 491
- cc2 instruction field
 - compare and branch instructions, 136
 - move instructions, 281, 491
- CCR (condition codes register), **5**
- CCR (condition codes) register, **50**
 - 32-bit operation (icc) bit of condition field, **50, 51**
 - 64-bit operation (xcc) bit of condition field, **50, 51**
 - ADD instructions, 110
 - ASR for, 49
 - carry (c) bit of condition fields, **51**
 - icc field, *See* CCR.icc field
 - negative (n) bit of condition fields, **50**
 - overflow bit (v) in condition fields, **51**
 - restored by RETRY instruction, 147, 318
 - saved after trap, 443
 - saving after trap, 21
 - TSTATE register, 68
 - write instructions, 374
 - xcc field, *See* CCR.xcc field
 - zero (z) bit of condition fields, **50**
- CCR.icc field
 - add instructions, 110, 362
 - bit setting for signed division, 326
 - bit setting for signed/unsigned multiply, 333, 371
 - bit setting for unsigned division, 370
 - branch instructions, 123, 127, 282
 - integer subtraction instructions, 357
 - logical operation instructions, 119, 296, 390
 - Tcc instruction, 365
- CCR.xcc field
 - add instructions, 110, 362
 - bit setting for signed/unsigned divide, 326, 370
 - bit setting for signed/unsigned multiply, 333, 371
 - branch instructions, 127, 282
 - logical operation instructions, 119, 296, 390
 - subtract instructions, 357
 - Tcc instruction, 365
- CFR register, **59**
 - RDCFR instruction, 310
- checksum, for 16-bit TCP/UDP, 161
- clean register window, 322, 460
- clean window, **5**
 - and window traps, 66, 465
 - CLEANWIN register, 66
 - definition, **465**
 - number is zero, 90
 - trap handling, 466
- clean_window* exception, 64, 90, 323, **460, 465, 502**
- CLEANWIN (clean windows) register, **64**
 - CANSAVE instruction, 90
 - clean window counting, 64
 - incremented by trap handler, 466
 - range of values, 62, 65
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - specifying number of available clean windows, 465
 - value calculation, 66
- clock cycle, counts for virtual processor, 52
- clock tick registers, *See* TICK *and* STICK registers
- clock-tick register (TICK), 463
- c1rn* synthetic instructions, 518
- CMASK16 instruction, **139**
- CMASK32 instruction, **139**
- CMASK8 instruction, **139**
- cmp synthetic instruction, 357, 516
- code
 - self-modifying, 419
- coherence, **5**
 - between processors, 505
 - data cache, 414
 - domain, 408
 - memory, 409
 - unit, memory, 410
- compare and branch instructions, **136**
- compare and swap instructions, **134, 136**
- comparison instruction, 85, 357
- Compatibility Feature register, *See* CFR
- compatibility note, **4**
- completed (memory operation), **5**
- compliance
 - SPARC V9, 481
- compliant SPARC V9 implementation, **16**
- cond instruction field
 - branch instructions, 123, 126, 157, 160
 - floating point move instructions, 180
 - move instructions, 281
- condition codes
 - adding, 362
 - effect of compare-and-swap instructions, 135

- extended integer (xcc), 51
- floating-point, 157
- icc field, 50
- integer, 50
- results of integer operation (icc), 51
- subtracting, 357, 367
- trapping on, 365
- xcc field, 50
- condition codes register, *See* CCR register
- conditional branches, 123, 157, 160
- conditional move instructions, 21
- conforming SPARC V9 implementation, **16**
- consistency
 - between instruction and data spaces, 419
 - processor, 414, 417
 - processor self-consistency, 416
 - sequential, 409, 415
 - strong, 415
- const22 instruction field of ILLTRAP instruction, 236
- constants, generating, 327
- context, 5
 - nucleus, 173
 - selection for translation, 479
- context identifier, **410**
- Context register
 - determination of, 479
 - Nucleus, 483
 - Primary, 483
 - Secondary, 483
- control transfer
 - pseudo-control-transfer via WRPR to PSTATE.am, 73
- control_transfer_instruction* exception, **460**
 - CALL/JMPL instructions, 130, 238
 - DONE/RETRY instructions, 148, 318, 365
 - RETURN, 320
 - and Tcc instruction, 366
 - with branch instructions, 124, 127, 128, 138, 158, 160
- control-transfer instructions, 20
- control-transfer instructions (CTIs), 20, 147, 318
- conventions
 - font, 2
 - notational, 2
- conversion
 - between floating-point formats instructions, 232
 - floating-point to integer instructions, 231, 393
 - integer to floating-point instructions, 168, 235
 - planar to packed, 216
- copyback, 5
- CPI, 5
- CPU, pipeline draining, 67
- cpu_mondo* exception, **461**
- CRC32C instruction, **141**
- cross-call, 5
- CTI, **6**, 11
- current exception (cexc) field of FSR register, **46**, 92, 500
- current window, **6**
- current window pointer register, *See* CWP register
- current_little_endian (cle) field of PSTATE register, **70**, 411
- CWBCC instruction, 136
- CWBCS instruction, 136
- CWBE instruction, 136
- CWBG instruction, 136
- CWBGE instruction, 136
- CWBGU instruction, 136
- CWBL instruction, 136
- CWBLE instruction, 136
- CWBLEU instruction, 136
- CWBNE instruction, 136
- CWBNEG instruction, 136
- CWBPOS instruction, 136
- CWBVC instruction, 136
- CWBVS instruction, 136
- CWP (current window pointer) register
 - and instructions
 - CALL and JMPL instructions, 35
 - FLUSHW instruction, 174
 - RDPR instruction, 313
 - RESTORE instruction, 90, 315
 - SAVE instruction, 90, 315, 322
 - WRPR instruction, 376
 - and traps
 - after spill trap, 466
 - after spill/fill trap, 21
 - on window trap, 466
 - saved by hardware, 443
- CWP (current window pointer) register, **63**
 - clean windows, 64
 - definition, 6
 - incremented/decremented, 35, 315, 322
 - overlapping windows, 35
 - range of values, 62, 63
 - restored during RETRY, 147, 318
 - specifying windows for use without cleaning, 465
 - and TSTATE register, 68
- CXBCC instruction, 136
- CXBCS instruction, 136
- CXBE instruction, 136
- CXBG instruction, 136
- CXBGE instruction, 136
- CXBGU instruction, 136
- CXBLE instruction, 136
- CXBLEU instruction, 136
- CXBNE instruction, 136
- CXBNEG instruction, 136
- CXBPOS instruction, 136
- CXBVC instruction, 136
- CXBVS instruction, 136
- cycle, **6**

D

- D superscript on instruction name, **96**
- d10_hi instruction field
 - compare and branch instructions, 136
- d10_lo instruction field
 - compare and branch instructions, 136
- d16hi instruction field
 - branch instructions, 128
- d16lo instruction field
 - branch instructions, 128

- DAE_invalid_ASI* exception
 - with load instructions and ASIs, 250, 435, 436, 437, 438
 - with store instructions and ASIs, 250, 435, 436, 437, 438
- DAE_invalid_asi* exception, 135, **461**
 - accessing noncacheable page, 418
 - SWAP/SWAPA, 361
 - with load alternate instructions, 135, 241, 249, 252, 253, 256, 260, 335, 336, 343, 347, 355, 360
 - with load instructions, 250, 256, 261
 - with load instructions and ASIs, 250
 - with nonfaulting load, 421
 - with store instructions, 256, 336, 344, 355
- DAE_invalid_ASI* exception (replacing SPARC V9 *data_access_exception*), **461**
- DAE_invalid_ASIn* exception
 - with load instructions and ASIs, 436
- DAE_nc_page* exception, **461**
 - accessing noncacheable page, 418
 - with compare-and-swap instructions, 135
 - with load instructions, 244, 255, 256, 263
 - with store instructions, 256
 - with SWAP instructions, 361
- DAE_nc_page* exception (replacing SPARC V9 *data_access_exception*), **461**
- DAE_nfo_page* exception, 135, **461**
 - and TTE.nfo, 476
 - with FLUSH instructions, 173
 - with LDTXA instructions, 263
 - with load alternate instructions, 241
 - with load instructions, 239, 247, 250, 252, 254, 255, 256, 258, 261, 265, 353
 - with nonfaulting load, 421
 - with store instructions, 256, 336, 339, 341, 344, 345, 349, 350, 355, 356
 - with SWAP instruction, 359
 - with SWAP instructions, 361
- DAE_nfo_page* exception (replacing SPARC V9 *data_access_exception*), **461**
- DAE_noncacheable_page* exception
 - with LDTXA instructions, 263
- DAE_privilege_violation* exception, 135, **461**
 - and TTE.p, 478
 - with load alternate instructions, 241
 - with load instructions, 245, 250, 252, 256, 258, 261, 339, 353
 - with store instructions, 256, 336, 341, 344, 345, 349, 350, 355, 356, 359
 - with SWAP instructions, 239, 361
- DAE_privilege_violation* exception (replacing SPARC V9 *data_access_exception*), **461**
- DAE_side_effect_page*
 - with nonfaulting loads, 409
- DAE_side_effect_page* exception, **462**
 - with load alternate instructions, 241
 - with load instructions, 245, 250, 261
 - with nonfaulting load, 421
- DAE_side_effect_page* exception (replacing SPARC V9 *data_access_exception*), **461**
- data
 - access, 6
 - cache coherence, 414
 - conversion between SIMD formats, 29
 - flow order constraints
 - memory reference instructions, **413**
 - register reference instructions, **413**
 - formats
 - byte, **23**
 - doubleword, **23**
 - halfword, **23**
 - Int16 SIMD, **30**
 - Int32 SIMD, **30**
 - quadword, **23**
 - tagged word, **23**
 - UInt8 SIMD, **30**
 - word, **23**
 - memory, 420
 - types
 - floating-point, 23
 - signed integer, 23
 - unsigned integer, 23
 - width, 23
 - watchpoint exception, 348
- Data Cache Unit Control register, *See* DCUCR
- data_access_exception* exception (SPARC V9), **461**
- DCTI couple, 89
- DCTI instructions, 6
 - behavior, 77
 - RETURN instruction effects, 320
- dec synthetic instructions, 518
- deccc synthetic instructions, 518
- deferred trap, **446**
 - distinguishing from disrupting trap, 447
 - floating-point, 314
 - restartable
 - implementation dependency, 447
 - software actions, 447
- delay instruction
 - and annul field of branch instruction, 157
 - annulling, 20
 - conditional branches, 160
 - DONE instruction, 147
 - executed after branch taken, 128
 - following delayed control transfer, 20
 - RETRY instruction, 318
 - RETURN instruction, 320
 - unconditional branches, 160
 - with conditional branch, 127
- delayed branch, 77
- delayed control transfer, 128
- delayed CTI, *See* DCTI
- denormalized number, 6
- deprecated, 6
- deprecated exceptions
 - tag_overflow*, **464**
- deprecated instructions
 - FBA, 157
 - FBE, 157
 - FBG, 157
 - FBGE, 157
 - FBL, 157

- FBLE, 157
- FBLG, 157
- FBN, 157
- FBNE, 157
- FBO, 157
- FBU, 157
- FBUE, 157
- FBUGE, 157
- FBUL, 157
- FBULE, 157
- LDFSR, 251
- LDTW, 257
- LDTWA, 259
- RDY, 49, 50, **310**
- SDIV, 50, 325
- SDIVcc, 50, 325
- SMUL, 50, 333
- SMULcc, 50, 333
- STFSR, 345
- STTW, 352
- STTWA, 354
- SWAP, 359
- SWAPA, 360
- TADDccTV, 363
- TSUBccTV, 368
- UDIV, 50, 369
- UDIVcc, 50, 369
- UMUL, 50, 371
- UMULcc, 50, 371
- WRY, 49, 50, 373
- dev_mondo* exception, **462**
- disp19 instruction field
 - branch instructions, 126, 160
- disp22 instruction field
 - branch instructions, 123, 157
- disp30 instruction field
 - word displacement (CALL), 130
- disrupting trap, **447**
- disrupting_performance_event* exception, **462**
- divide instructions, 20, 291, 325, 369
- division_by_zero* exception, 85, 291, **462**
- division-by-zero bits of FSR.aexc/FSR.cexc fields, 48
- DMMU
 - context register usage, 481
 - Secondary Context register, 483
- DONE instruction, **147**
 - effect on TNPC register, 68
 - effect on TSTATE register, 68
 - generating *illegal_instruction* exception, 463
 - modifying CCR.xcc condition codes, 51
 - return from trap, 443
 - return from trap handler with different GL value, 76
 - target address, 20
- doubleword, **6**
 - addressing, 82
 - alignment, 18, 79, 409
 - data format, **23**
 - definition, 6

E

- EDGE16 instruction, 149
- EDGE16L instruction, 149
- EDGE16LN instruction, 150
- EDGE16N instruction, 150
- EDGE32 instruction, 149
- EDGE32L instruction, 149
- EDGE32LN instruction, 150
- EDGE32N instruction, 150
- EDGE8 instruction, 149
- EDGE8L instruction, 149
- EDGE8LN instruction, 150
- EDGE8N instruction, 150
- emulating multiple unsigned condition codes, 90
- enable floating-point
 - See FPRS register, *fef* field
 - See PSTATE register, *pef* field
- even parity, **6**
- exception, **6**
- exceptions
 - See also individual exceptions
 - catastrophic error, 443
 - causing traps, 443
 - clean_window*, **460**, 502
 - control_transfer_instruction*, **460**
 - cpu_mondo*, **461**
 - DAE_invalid_asl*, **461**
 - DAE_invalid_ASI* (replacing SPARC V9 *data_access_exception*), **461**
 - DAE_nc_page*, **461**
 - DAE_nc_page* (replacing SPARC V9 *data_access_exception*), **461**
 - DAE_nfo_page*, **461**
 - DAE_nfo_page* (replacing SPARC V9 *data_access_exception*), **461**
 - DAE_privilege_violation*, **461**
 - DAE_privilege_violation* (replacing SPARC V9 *data_access_exception*), **461**
 - DAE_side_effect_page*, **462**
 - DAE_side_effect_page* (replacing SPARC V9 *data_access_exception*), **461**
 - data_access_exception* (SPARC V9), **461**
 - definition, 443
 - dev_mondo*, **462**
 - disrupting_performance_event*, **462**
 - division_by_zero*, **462**
 - fill_n_normal*, **462**
 - fill_n_other*, **462**
 - fp_disabled*
 - and GSR, 54
 - fp_disabled*, **462**
 - fp_exception_ieee_754*, **462**
 - fp_exception_other*, **462**
 - htrap_instruction*, **462**
 - IAE_privilege_violation*, **462**
 - illegal_instruction*, **462**
 - instruction_access_exception* (SPARC V9), **463**
 - instruction_VA_watchpoint*, **463**
 - interrupt_level_14*
 - and SOFTINT.int_level, 55

- interrupt_level_14*, 463
- interrupt_level_15*
 - and SOFTINT.int_level, 55
- interrupt_level_n*
 - and SOFTINT register, 55
 - and SOFTINT.int_level, 55
- interrupt_level_n*, 448, 463
- LDDF_mem_address_not_aligned*, 463
- LDQF_mem_address_not_aligned*, 464
- mem_address_not_aligned*, 463
- nonresumable_error*, 463
- pending, 21
- privileged_action*, 463
- privileged_opcode*
 - and access to register-window PR state registers, 62, 66, 74, 75
 - and access to SOFTINT, 55
 - and access to SOFTINT_CLR, 56
 - and access to SOFTINT_SET, 56
- privileged_opcode*, 464
- resumable_error*, 464
- spill_n_normal*, 323, 464
- spill_n_other*, 323, 464
- STDF_mem_address_not_aligned*, 464
- STQF_mem_address_not_aligned*, 464
- tag_overflow* (deprecated), 464
- trap_instruction*, 464
- unimplemented_LDTW*, 464
- unimplemented_STTW*, 464
- VA_watchpoint*, 464
- execute unit, 412
- execute_state
 - trap processing, 458
- explicit ASI, 6, 83, 424
- extended word, 6
 - addressing, 82

F

- F registers, 6, 17, 92, 391, 450
- FABSd instruction, 152, 489, 490
- FABSq instruction, 152, 152, 489, 490
- FABSs instruction, 152
- FADD, 153
- FADDd instruction, 153
- FADDq instruction, 153, 153, 164
- FADDs instruction, 153
- FALIGNDATAg instruction, 154
- FALIGNDATAi instruction, 155
- FAND instruction (now FANDd), 227
- FANDd instruction, 227
- FANDNOT1 instruction (now FANDNOT1d), 227
- FANDNOT1d instruction, 227
- FANDNOT1S instruction, 227
- FANDNOT2 instruction (now FANDNOT2d), 227
- FANDNOT2d instruction, 227
- FANDNOT2S instruction, 227
- FANDS instruction, 227
- FBA instruction, 157, 491
- FBE instruction, 157, 491

- FBfcc instructions, 42, 157, 462, 485, 491
- FBG instruction, 157, 491
- FBGE instruction, 157, 491
- FBL instruction, 157, 491
- FBLE instruction, 157, 491
- FBLG instruction, 157, 491
- FBN instruction, 157, 491
- FBNE instruction, 157, 491
- FBO instruction, 157, 491
- FBPA instruction, 159, 160, 491
- FBPE instruction, 159, 491
- FBPfcc instructions, 42, 159, 485, 491, 492
- FBPG instruction, 159, 491
- FBPGE instruction, 159, 491
- FBPL instruction, 159, 491
- FBPLE instruction, 159, 491
- FBPLG instruction, 159, 491
- FBPN instruction, 159, 160, 491
- FBPNE instruction, 159, 491
- FBPO instruction, 159, 491
- FBPU instruction, 159, 491
- FBPUE instruction, 159, 491
- FBPUG instruction, 159, 491
- FBPUGE instruction, 159, 491
- FBPUL instruction, 159, 491
- FBPULE instruction, 159, 491
- FBU instruction, 157, 491
- FBUE instruction, 157, 491
- FBUG instruction, 157, 491
- FBUGE instruction, 157, 491
- FBUL instruction, 157, 491
- FBULE instruction, 157, 491
- fcc-conditional branches, 157, 160
- fccn*, 6
- FCHKSM16 instruction, 161
- FCMP instructions, 492
- FCMP* instructions, 42, 162
- FCMPd instruction, 162, 490
- FCMPE instructions, 492
- FCMPE* instructions, 42, 162
- FCMPEd instruction, 162, 490
- FCMPEq instruction, 162, 163, 490
- FCMPEQ16 instruction, 207
- FCMPEQ16, *See* FPCMPEQ16 instruction
- FCMPEQ32 instruction, 207
- FCMPEQ32, *See* FPCMPEQ32 instruction
- FCMPEs instruction, 162, 490
- FCMPGT16 instruction, 207, 207
- FCMPGT16, *See* FPCMPTGT16 instruction
- FCMPGT32 instruction, 207
- FCMPGT32, *See* FPCMPTGT32 instruction
- FCMPLE16 instruction, 207
- FCMPLE16, *See* FPCMPL16 instruction
- FCMPLE32 instruction, 207
- FCMPLE32, *See* FPCMPL32 instruction
- FCMPLE8 instruction, 207
- FCMPNE16 instruction, 207
- FCMPNE16, *See* FPCMPTNE16 instruction
- FCMPNE32 instruction, 207
- FCMPNE32, *See* FPCMPTNE32 instruction

FCMPq instruction, **162**, 163, 490
 FCMPs instruction, **162**, 490
 FCMPUEQ16 instruction, **210**
 FCMPUEQ32 instruction, **210**
 FCMPUNE16 instruction, **210**
 FCMPUNE32 instruction, **210**
 fcn instruction field
 DONE instruction, 147
 PREFETCH, 303
 RETRY instruction, 318
 FDIVd instruction, **164**
 FDIVq instruction, **164**
 FDIVs instructions, 164
 FdMULq instruction, **190**, 190
 FdTOi instruction, **231**, 393
 FdTOq instruction, **232**, 232
 FdTOs instruction, **232**
 FdTOx instruction, **231**, 490
 feature set, **100**
 fef field of FPRS register, **53**
 and access to GSR, 54
 and *fp_disabled* exception, 462
 branch operations, 158, 160
 byte permutation, 125
 comparison operations, 163, 170, 209, 212
 component distance, 284, 285, 299, 300
 data formatting operations, 165, 197, 216
 data movement operations, 282
 enabling FPU, 71
 FCHKSM16 instruction, 161
 floating-point operations, 140, 152, 153, 164, 168, 176, 179,
 182, 185, 190, 191, 193, 195, 230, 231, 232, 234, 235, 246,
 248, 251, 253, 264
 integer arithmetic operations, 186, 203, 206, 215, 218, 221,
 224, 229
 integer SIMD operations, 178
 logical operations, 225, 226, 227
 memory operations, 244
 read operations, 311, 330, 339
 special addressing operations, 117, 154, 156, 340, 342, 345,
 348, 350, 356, 374
 fef, *See* FPRS register, fef field
 FEXPAND instruction, 165
 FEXPAND operation, 165
 FHADDd instruction, **166**, **194**
 FHADDs instruction, **166**
 FHSUBd instruction, **167**
 FHSUBs instruction, **167**
 fill handler, 315
 fill register window, 462
 overflow/underflow, 36
 RESTORE instruction, 66, 315, 465
 RESTORED instruction, 91, 466
 RETRY instruction, 466
 selection of, 465
 trap handling, 465, 466
 trap vectors, 315
 window state, 66
 fill_n_normal exception, 316, 321, 462, **462**
 fill_n_other exception, 316, 321, **462**
 FiTOd instruction, **168**
 FiTOq instruction, **168**, 168, 235
 FiTOs instruction, **168**
 fixed-point scaling, 186
 FLCMPD instruction, 169
 FLCMPS instruction, 169
 floating point
 absolute value instructions, **152**
 add instructions, **153**
 compare instructions, 42, **162**, 162
 condition code bits, 157
 condition codes (fcc) fields of FSR register, 45, 157, 160,
 162
 data type, 23
 deferred-trap queue (FQ), 314
 divide instructions, **164**
 exception, **6**
 exception, encoding type, 44
 FPRS register, 374
 FSR condition codes, 42
 lexicographic compares, 169
 move instructions, **179**
 multiply instructions, **190**
 multiply-add
 See also FMA instructions
 multiply-add/subtract, 175, 191, 195
 negate instructions, **193**
 operate (FPop) instructions, **6**, **21**, 44, 46, 92, 251
 registers
 destination F, 391
 FPRS, *See* FPRS register
 FSR, *See* FSR register
 programming, 41
 rounding direction, 43
 square root instructions, **230**
 subtract instructions, **234**
 trap types, **6**
 IEEE_754_exception, 44, **45**, 46, 48, 391, 392
 invalid_fp_register, 234
 unfinished_FPop, 44, **45**, 48, 153, 164, 190, 230, 232, 234,
 391
 results after recovery, 45
 unimplemented_FPop, 48, 185, 392
 traps
 deferred, 314
 precise, 314
 floating-point condition codes (fcc) fields of FSR register, 450
 floating-point operate (FPop) instructions, 462
 floating-point trap types
 IEEE_754_exception, 450, 462
 floating-point unit (FPU), **7**, **17**
 FLUSH instruction, 172
 memory ordering control, 270
 FLUSH instruction
 memory/instruction synchronization, 171
 FLUSH instruction, **171**, 420
 data access, 6
 immediacy of effect, 173
 in multiprocessor system, 171
 in self-modifying code, 172

- latency, 505
- flush instruction memory, *See* FLUSH instruction
- flush register windows instruction, 174
- FLUSHW instruction, 174, 464
 - effect, 21
 - management by window traps, 66, 465
 - spill exception, 91, 174, 466
- FMA instructions
 - fused, 175
- FMADDd instruction, 175
- FMADDs instruction, 175
- FMEAN16 instruction, 177
- FMOVcc instructions
 - conditionally moving floating-point register contents, 51
 - conditions for copying floating-point register contents, 89
 - copying a register, 42
 - encoding of `opf<84>` bits, 490
 - encoding of `opf_cc` instruction field, 491
 - encoding of `rcond` instruction field, 491
 - floating-point moves, 180
 - FPop instruction, 92
 - used to avoid branches, 183, 282
- FMOVccd instruction, 490
- FMOVccq instruction, 490
- FMOVd instruction, 179, 489, 490
- FMOVdfcc instructions, 180
- FMOVdGEZ instruction, 184
- FMOVdGZ instruction, 184
- FMOVdicc instructions, 180
- FMOVdLEZ instruction, 184
- FMOVdLZ instruction, 184
- FMOVdNZ instruction, 184
- FMOVdZ instruction, 184
- FMOVq instruction, 179, 179, 489, 490
- FMOVQfcc instructions, 180, 182
- FMOVqGEZ instruction, 184
- FMOVqGZ instruction, 184
- FMOVQicc instructions, 180, 182
- FMOVqLEZ instruction, 184
- FMOVqLZ instruction, 184
- FMOVqNZ instruction, 184
- FMOVqZ instruction, 184
- FMOVr instructions, 92, 491
- FMOVRq instructions, 185
- FMOVRsGZ instruction, 184
- FMOVRsLEZ instruction, 184
- FMOVRsLZ instruction, 184
- FMOVRsNZ instruction, 184
- FMOVRsZ instruction, 184
- FMOVs instruction, 179
- FMOVSc instructions, 182
- FMOVsfcc instructions, 180
- FMOVsGEZ instruction, 184
- FMOVsicc instructions, 180
- FMOVsxcc instructions, 180
- FMOVxcc instructions, 180, 182
- FMSUBd instruction, 175
- FMSUBs instruction, 175
- FMUL8SUx16 instruction, 186, 188
- FMUL8ULx16 instruction, 186, 188
- FMUL8x16 instruction, 186
- FMUL8x16AL instruction, 186, 187
- FMUL8x16AU instruction, 186, 187
- FMULd instruction, 190
- FMULD8SUx16 instruction, 186, 189
- FMULD8ULx16 instruction, 186, 189
- FMULq instruction, 190, 190
- FMULs instruction, 190
- FNADD instructions, fused, 191
- FNADDd instruction, 191
- FNADDs instruction, 191
- FNAND instruction (now FNANDd), 227
- FNANDd instruction, 227
- FNANDS instruction, 227
- FNEG instructions, 193
- FNEGd instruction, 193, 489, 490
- FNEGq instruction, 193, 193, 489, 490
- FNEGs instruction, 193
- FNHADDs instruction, 194
- FNMADDd instruction, 175
- FNMADDs instruction, 175
- FNMSUBd instruction, 175
- FNMSUBs instruction, 175
- FNMUL instructions
 - fused, 195
- FNMULd instruction, 195
- FNMULs instruction, 195
- FNOR instruction (now FNORd), 227
- FNORd instruction, 227
- FNORS instruction, 227
- FNOT1 instruction (now FNOT1d), 226
- FNOT1d instruction, 226
- FNOT1S instruction, 226
- FNOT2 instruction (now FNOT2d), 226
- FNOT2d instruction, 226
- FNOT2S instruction, 226
- FNsMULd instruction, 195
- FONE instruction (now FONEd), 225
- FONEd instruction, 225
- FONES instruction, 225
- FOR instruction (FORd), 227
- FORd instruction, 227
- formats, instruction, 78
- FORNOT1 instruction (now FORNOT1d), 227
- FORNOT1d instruction, 227
- FORNOT1S instruction, 227
- FORNOT2 instruction (now FORNOT2d), 227
- FORNOT2d instruction, 227
- FORNOT2S instruction, 227
- FORS instruction, 227
- fp_disabled* exception, 462
 - absolute value instructions, 152, 153, 234
 - add and halve instructions, 166, 167, 194
 - and GSR, 54
 - checksum instruction, 161
 - FPop instructions, 92
 - FPRS.fef disabled, 53
 - MOVfTOi instruction, 284
 - MOViTOf instruction, 285
 - pixel component distance, 300

- PSTATE.pef not set, 53, 54, 508
- with branch instructions, 158, 160
- with compare instructions, 209, 212
- with conversion instructions, 168, 231, 232, 235
- with floating-point arithmetic instructions, 164, 176, 178, 190, 192, 196, 203, 206, 215, 218, 221, 229
- with FMOV instructions, 179
- with load instructions, 250, 254
- with move instructions, 183, 185, 282
- with negate instructions, 193
- with store instructions, 340, 341, 342, 344, 345, 348, 350, 356, 374
- fp_exception* exception, 46
- fp_exception_ieee_754* "invalid" exception, 231
- fp_exception_ieee_754* exception, 462
 - and tem bit of FSR, 43
 - cause encoded in FSR.ftt, 44
 - FSR.aexc, 46
 - FSR.cexc, 47
 - FSR.ftt, 46
 - generated by FCMP or FCMPE, 42
 - and IEEE 754 overflow/underflow conditions, 46, 47
 - trap handler, 392
 - when FSR.ns = 1, 394, 500
 - when FSR.tem = 0, 450
 - when FSR.tem = 1, 450
 - with floating-point add and halve instructions, 166, 167, 194
 - with floating-point arithmetic instructions, 153, 164, 176, 190, 192, 196, 234
- fp_exception_other* exception, 48, 462
 - cause encoded in FSR.ftt, 44
 - FSUBq instruction, 234
 - incorrect IEEE Std 754-1985 result, 92, 499
 - never generated by FLCMP instructions, 169
 - supervisor handling, 392
 - trap type of unfinished_FPop, 45
 - when quad FPop unimplemented in hardware, 46
 - with floating-point arithmetic instructions, 164, 176, 190, 192, 196
- FPACK instruction, 54
- FPACK instructions, 197–200
- FPACK16 instruction, 197, 198
- FPACK16 operation, 198
- FPACK32 instruction, 197, 199
- FPACK32 operation, 199
- FPACKFIX instruction, 197, 200
- FPACKFIX operation, 200
- FPADD16 instruction, 201
- FPADD16S instruction, 201
- FPADD32 instruction, 201
- FPADD32S instruction, 201
- FPADD64 instruction, 201
- FPADD8 instruction, 201
- FPADDS16 instruction, 204
- FPADDS16S instruction, 204
- FPADDS32 instruction, 204
- FPADDS32S instruction, 204
- FPADDS8 instruction, 204
- FPADDUS16 instruction, 204
- FPADDUS8 instruction, 204
- FPCMPEQ16 instruction, 208
- FPCMPEQ32 instruction, 208
- FPCMPEQ8 instruction, 207
- FPCMPGT instruction, 208
- FPCMPLE16 instruction, 208
- FPCMPLE32 instruction, 208
- FPCMPNE16 instruction, 208
- FPCMPNE32 instruction, 208
- FPCMPNE8 instruction, 207
- FPCMPUEQ8 instruction, 210
- FPCMPUGT16 instruction, 210
- FPCMPUGT32 instruction, 210
- FPCMPUGT8 instruction, 210
- FPCMPULE16 instruction, 210
- FPCMPULE32 instruction, 210
- FPCMPULE8 instruction, 210
- FPCMPUNE8 instruction, 210
- FPMAX16 instruction, 214
- FPMAX32 instruction, 214
- FPMAX8 instruction, 214
- FPMAXU16 instruction, 214
- FPMAXU32 instruction, 214
- FPMAXU8 instruction, 214
- FPMERGE instruction, 216
- FPMIN16 instruction, 217
- FPMIN32 instruction, 217
- FPMIN8 instruction, 217
- FPMINU16 instruction, 217
- FPMINU32 instruction, 217
- FPMINU8 instruction, 217
- FPop, 7
- FPop, *See* floating-point operate (FPop) instructions
- FPRS register
 - See also* floating-point registers state (FPRS) register
 - FPRS register, 53
 - ASR summary, 49
 - definition, 7
 - fef field, 92, 449
 - RDFPRS instruction, 311
 - FPRS register fields
 - dl (dirty lower fp registers), 54
 - du (dirty upper fp registers, 53
 - fef, 53
 - fef, *See also* fef field of FPRS register
- FPSUB16 instruction, 219
- FPSUB16s instruction, 219
- FPSUB32 instruction, 219
- FPSUB32s instruction, 219
- FPSUB64 instruction, 219
- FPSUB8 instruction, 219
- FPSUBS16 instruction, 222
- FPSUBS16S instruction, 222
- FPSUBS32 instruction, 222
- FPSUBS32S instruction, 222
- FPSUBS8 instruction, 222
- FPSUBUS16 instruction, 222
- FPSUBUS8 instruction, 222
- FPU, 6, 7
- FqTOd instruction, 232, 232

FqTOi instruction, **231**, 231, 393
 FqTOs instruction, **232**, 232
 FqTOx instruction, **231**, 231, 489, 490
freg, **512**
 FsMULd instruction, **190**
 FSQRTd instruction, **230**
 FSQRTq instruction, **230**, 230
 FSQRTs instruction, **230**
 FSR (floating-point state) register
 fields
 aexc (accrued exception), 44, 45, **46**, 46, 391
 aexc (accrued exceptions), 176, 191, 195
 in user-mode trap handler, 392
 -- dza (division by zero) bit of aexc, 48
 -- nxa (rounding) bit of aexc, 48
 cexc (current exception), 43, 44, 45, **46**, 46, 47, 391, 462
 cexc (current exceptions), 176, 191, 195
 in user-mode trap handler, 392
 -- dzc (division by zero) bit of cexc, 48
 -- nxc (rounding) bit of cexc, 48
 fcc (condition codes), **42**, 45, 392, 512
 fcc*n*, **42**
 ftt (floating-point trap type), 42, **44**, 46, 92, 264, 345, 356, 462
 in user-mode trap handler, 392
 modified by LDXEFSR instruction, 42
 not modified by LDFSR/LDXFSR instructions, 42
 ns (nonstandard mode), 42, 251, 264
 qne (queue not empty), 42, 251, 264
 in user-mode trap handler, 392
 rd (rounding), **43**
 tem (trap enable mask), **43**, 46, 47, 176, 191, 195, 393, 462
 ver, **43**
 ver (version), 42, 264
 FSR (floating-point state) register, **42**
 after floating-point trap, 391
 compliance with IEEE Std 754-1985, 48
 LDFSR instruction, 251
 reading/writing, 42
 values in ftt field, 44
 writing to memory, 345, 356
 FSRC1 instruction (now FSRC1d), 226
 FSRC1d instruction, **226**
 FSRC1S instruction, 226
 FSRC2 instruction (now FSRC2d), 226
 FSRC2d instruction, **226**
 FSRC2S instruction, 226
 FsTOd instruction, **232**
 FsTOi instruction, **231**, 393
 FsTOq instruction, **232**, 232
 FsTOx instruction, **231**, 489, 490
 FSUBd instruction, **234**
 FSUBq instruction, **234**, 234
 FSUBs instruction, **234**
 functional choice, implementation-dependent, 498
 fused FMA instructions, 175
 fused FNADD instructions, **191**
 fused FNMUL instructions, **195**
 FXNOR instruction (now FXNORd), 227
 FXNORd instruction, **227**

FXNORS instruction, 227
 FXOR instruction (now FXORd), 227
 FXORd instruction, **227**
 FXORS instruction, 227
 FxTOd instruction, **235**, 490
 FxTOq instruction, **235**, 490
 FxTOs instruction, **235**, 490
 FZERO instruction (now FZEROd), 225
 FZEROd instruction, **225**
 FZEROS instruction, 225

G

general status register, *See* GSR (general status) register
 generating constants, 327
 GL register, **75**
 access, 75
 during trap processing, 458
 function, 75
 reading with RDPR instruction, 313, 376
 relationship to TL, 75
 restored during RETRY, 147, 318
 SPARC V9 compatibility, 73
 and TSTATE register, 68
 value restored from TSTATE[TL], 76
 writing to, 75
 global level register, *See* GL register
 global registers, 14, 17, 32, **33**, 33, 499
 graphics status register, *See* GSR (general status) register
 GSR (general status) register
 fields
 align, **54**
 im (interval mode) field, **54**
 irnd (rounding), **54**
 mask, **54**
 scale, **54**
 GSR (general status) register
 ASR summary, 49
 mask values with CMASK instructions, 139

H

halfword, 7
 alignment, 18, **79**, 409
 data format, **23**
 hardware
 dependency, **498**
 traps, 451
 hardware trap stack, 21
htrap_instruction exception, 366, **462**
 hyperprivileged, 7

I

i (integer) instruction field
 arithmetic instructions, 291, 296, 325, 333, 369, 371
 floating point load instructions, 246, 248, 251, 264
 flush memory instruction, 171
 flush register instruction, 174
 jump-and-link instruction, 238

- load instructions, 239, 255, 256, 257, 259
- logical operation instructions, 118, 296, 390
- move instructions, 281, 283
- POPC, 301
- PREFETCH, 303
- RETURN, 320
- I/O
 - access, 408
 - memory, 407
 - memory-mapped, 408
- IAE_privilege_violation* exception, 462
 - and TTE.p, 478
- IAE_unauth_access* exception, 478
- IEEE 754, 7
- IEEE Std 754-1985, 7, 13, 43, 45, 47, 48, 92, 391, 499
- IEEE_754_exception floating-point trap type, 7, 44, 45, 46, 48, 391, 392, 450, 462
- IEEE-754 exception, 7
- IER register (SPARC V8), 374
- illegal_instruction* exception, 135, 174, 462
 - attempt to write in nonprivileged mode, 57
 - DONE/RETRY, 148, 319, 320
 - ILLTRAP, 236
 - not implemented in hardware, 109
 - POPC, 302
 - PREFETCH, 309
 - RETURN, 321
 - with BPr instruction, 128
 - with branch instructions, 127, 129
 - with CASA and CASXA instructions, 134, 296
 - with CBcond instructions, 138
 - with DONE instruction, 148
 - with FMOV instructions, 179
 - with FMOVcc instructions, 183
 - with FMOVr instructions, 185
 - with load instructions, 38, 245, 247, 257, 260, 265, 437
 - with move instructions, 282, 283, 284, 285
 - with read hyperprivileged register instructions, 313
 - with read instructions, 311, 313, 377, 501
 - with store instructions, 250, 341, 345, 352, 353, 354, 355, 356
 - with Tcc instructions, 366
 - with TPC register, 67
 - with TSTATE register, 68
 - with write instructions, 374, 377
 - write to ASR 5, 53
 - write to STICK register, 57
 - write to TICK register, 52
- ILLTRAP instruction, 236, 462
- imm_asi* instruction field
 - explicit ASI, providing, 83
 - floating point load instructions, 248
 - load instructions, 256, 257, 259
 - PREFETCH, 303
- immediate CTI, 77
- I-MMU
 - and instruction prefetching, 409
- IMMU
 - context register usage, 481
- IMPDEP1 instructions, 92
- IMPDEP2 instructions, 92
- IMPDEP2A instructions, 502
- IMPDEP2B instructions, 92
- implementation, 7
- implementation dependency, 497
- implementation dependent, 7
- implementation note, 3, 4
- implementation-dependent functional choice, 498
- implicit ASI, 7, 83, 424
- implicit ASI memory access
 - LDFSR, 251
 - LDSTUB, 255
 - load fp instructions, 246, 264
 - load integer doubleword instructions, 257
 - load integer instructions, 239
 - STD, 352
 - STFSR, 345
 - store floating-point instructions, 340, 356
 - store integer instructions, 334
 - SWAP, 359
- implicit byte order, 70
- in* registers, 33, 35, 322
- inccc* synthetic instructions, 518
- inexact accrued (*nx*) bit of *aexc* field of FSR register, 392
- inexact current (*nx*) bit of *cexc* field of FSR register, 392
- inexact mask (*nxm*) field of FSR.tem, 47
- inexact quotient, 325, 369
- infinity, 393
- initiated, 7
- input/output (I/O) locations
 - access by nonprivileged code, 499
 - behavior, 408
 - contents and addresses, 499
 - identifying, 504
 - order, 408
 - semantics, 505
 - value semantics, 408
- instruction fields, 7
 - See also* individual instruction fields
 - definition, 7
- instruction group, 7
- instruction MMU, *See* I-MMU
- instruction prefetch buffer, invalidation, 172
- instruction set architecture (ISA), 7, 7, 14
- instruction_access_exception* exception (SPARC V9), 463
- instruction_va_watchpoint* exception, 463
- instructions
 - 32-bit wide, 13
 - alignment, 79
 - alignment, 18, 117, 409
 - arithmetic, integer
 - addition, 110, 111, 362
 - division, 20, 291, 325, 369
 - multiplication, 20, 291, 333, 371
 - subtraction, 357, 358, 367
 - tagged, 20
 - array addressing, 120
 - atomic
 - CASA/CASXA, 134
 - load twin extended word from alternate space, 262

- load-store, 79, 134, 136, 255, 256, 359, 360
- load-store unsigned byte, 255, 256
- successful loads, 239, 240, 258, 260
- successful stores, 334, 335
- branch
 - branch if contents of integer register match condition, **128**
 - branch on floating-point condition codes, **157, 159**
 - branch on integer condition codes, **123, 126**
- cache, 414
- causing illegal instruction, 236
- compare and swap, **134, 136**
- comparison, 85, 357
- conditional move, 21
- control-transfer (CTIs), 20, 147, 318
- conversion
 - convert between floating-point formats, **232**
 - convert floating-point to integer, **231**
 - convert integer to floating-point, **168, 235**
 - floating-point to integer, 393
- count of number of bits, 301
- crypto
 - AES, 112
 - Camellia, 131
 - DES, 143
 - MD5 hash operation, 268
 - MONTMUL, 272
 - MONTSQR, 276
 - MPMUL, 286
 - SHA1, SHA256, SHA512 hash operations, 328
 - XMONTMUL, 378
 - XMONTSQR, 381
 - XMPMUL, 384
- edge handling, 150
- fetches, 79
- floating point
 - compare, 42, 162
 - floating-point add, **153**
 - floating-point divide, **164**
 - floating-point load, 79, 246
 - floating-point load from alternate space, 248
 - floating-point load state register, 246, 264
 - floating-point move, 179, **180**, 184
 - floating-point operate (FPop), 21, 251
 - floating-point square root, **230**
 - floating-point store, 79, 340
 - floating-point store to alternate space, 342
 - floating-point subtract, **234**
 - operate (FPop), 44, 46
 - short floating-point load, 253
 - short floating-point store, 350
 - status of floating-point load, 251
- flush instruction memory, **171**
- flush register windows, **174**
- formats, **78**
- integer multiply-add, 213
- jump and link, 20, **238**
- loads
 - block load, 243
 - floating point, *See* instructions: floating point
 - integer, 79
 - simultaneously addressing doublewords, 359
 - unsigned byte, 134, **255**
 - unsigned byte to alternate space, 256
- logical operations
 - 64-bit/32-bit, 226, 227
 - AND, **118**
 - logical 1-operand ops on F registers, 225
 - logical 2-operand ops on F registers, 226
 - logical 3-operand ops on F registers, 227
 - logical XOR, 390
 - OR, **296**
- memory, 420
- moves
 - floating point, *See* instructions: floating point
 - move integer register, **280, 283**
 - on condition, 14
- MOVxTOd, **284, 285**
- MWAITMWAIT, 61
- MWAITMWAIT, 292
- ordering MEMBAR, 85
- PAUSE, **298**
- permuting bytes specified by GSR.mask, 125
- pixel component distance, **299, 299, 300, 300**
- pixel formatting (PACK), 197
- prefetch data, **303**
- read privileged register, **313**
- read state register, 20, **310**
- register window management, 21
- reordering, 412
- reserved, 93
- reserved* fields, 108
- RETRY
 - and restartable deferred traps, 447
- RETURN vs. RESTORE, 320
- sequencing MEMBAR, 85
- set high bits of low word, 327
- set interval arithmetic mode, 330
- setting GSR.mask field, 125
- shift, 19
- shift, **331**
- shift count, 331
- SIMD, **11**
- simultaneous addressing of doublewords, 360
- stores
 - block store, 337
 - floating point, *See* instructions: floating point
 - integer, 79, 334
 - integer (except doubleword), 334
 - integer into alternate space, 335
 - partial, 347
 - unsigned byte, 134
 - unsigned byte to alternate space, 256
 - unsigned bytes, 255
- swap R register, **359, 360**
- synthetic (for assembly language programmers), 516–518
- tagged addition, 362
- test-and-set, 418
- timing, 109
- trap on integer condition codes, **364**

- write privileged register, 376
- write state register, 374
- WRMWAITMWAIT, 61
- XMULX, 389
- XMULXHI, 389
- integer unit (IU)
 - condition codes, 51
 - definition, 7
 - description, 16
- interrupt
 - enable (ie) field of PSTATE register, 448, 449
 - level, 74
 - request, 7, 21, 443
- interrupt_level_14* exception, 55, 463
 - and SOFTINT.int_level, 55
- interrupt_level_15* exception
 - and SOFTINT.int_level, 55
- interrupt_level_n* exception, 448, 463
 - and SOFTINT register, 55
 - and SOFTINT.int_level, 55
- inter-strand operation, 7
- intra-strand operation, 7
- invalid accrued (nva) bit of aexc field of FSR register, 48
- invalid ASI
 - and *DAE_invalid_asi*, 461
- invalid current (nvc) bit of cexc field of FSR register, 48, 393
- invalid mask (nvm) field of FSR.tem, 47, 393
- invalid_exception* exception, 231
- invalid_fp_register floating-point trap type, 185, 193, 234
- INVALW instruction, 237
- iprefetch synthetic instruction, 517
- ISA, 7
- ISA, *See* instruction set architecture
- issue unit, 412, 412
- issued, 8
- italic font, in assembly language syntax, 511
- IU, 8
- ixc synthetic instructions, 518

J

- jmp synthetic instruction, 516
- JMPL instruction, 238
 - computing target address, 20
 - does not change CWP, 35
 - mem_address_not_aligned* exception, 463
 - reexecuting trapped instruction, 320
- jump and link, *See* JMPL instruction

L

- LD instruction (SPARC V8), 239
- LDBLOCKF instruction, 243, 436
- LDBLOCKF instruction, *DAE_nc_page* exception, 461
- LDD instruction (SPARC V8 and V9), 257
- LDDA instruction, 436
- LDDA instruction (SPARC V8 and V9), 260
- LDDF instruction, 79, 246, 463
- LDDF_mem_address_not_aligned* exception, 246, 463
 - address not doubleword aligned, 503
 - address not quadword aligned, 504
 - LDDF/LDDFA instruction, 79
 - load instruction with partial store ASI and misaligned address, 250
 - with load instructions, 248, 437
 - with store instructions, 437
- LDDF_mem_not_aligned* exception, 41
- LDDFA instruction, 248, 349
 - alignment, 79
 - ASIs for fp load operations, 437
 - behavior with block store with Commit ASIs, 249
 - behavior with partial store ASIs, 247-??, 250, 250-??, 264-??, 437-??
 - causing *LDDF_mem_address_not_aligned* exception, 79, 463
 - for block load operations, 436
 - used with ASIs, 436
- LDF instruction, 41, 246
- LDFFA instruction, 41, 248
- LDFSR instruction, 42, 44, 251, 463
- LDQF instruction, 246, 464
- LDQF_mem_address_not_aligned* exception, 246, 464
 - address not quadword aligned, 504
 - LDQF/LDQFA instruction, 80
- LDQF_mem_not_aligned* exception, 41
- LDQFA instruction, 248
- LDSB instruction, 239
- LDSBA instruction, 240
- LDSH instruction, 239
- LDSHA instruction, 240
- LDSHORTF instruction, 253
- LDSTUB instruction, 79, 255, 256, 418, 419
 - and *DAE_nc_page* exception, 461
 - hardware primitives for mutual exclusion of LDSTUB, 418
- LDSTUBA instruction, 255, 256
 - alternate space addressing, 18
 - and *DAE_nc_page* exception, 461
 - hardware primitives for mutual exclusion of LDSTUBA, 418
- LDSW instruction, 239
- LDSWA instruction, 240
- LDTW instruction, 38, 79
- LDTW instruction (deprecated), 257
- LDTWA instruction, 38, 79
- LDTWA instruction (deprecated), 259
- LDTX instruction, 434
- LDTX instruction, *DAE_nc_page* exception, 461
- LDTXA instruction, 80, 82, 262, 435
 - access alignment, 79
 - access size, 79
- LDUB instruction, 239
- LDUBA instruction, 240
- LDUH instruction, 239
- LDUHA instruction, 240
- LDUW instruction, 239
- LDUWA instruction, 240
- LDX instruction, 239
- LDXA instruction, 240, 260, 416
- LDXEFSR instruction, 42, 44, 46, 79, 264

LDXFSR instruction, 42, 44, 251, **264**, **324**, 463
leading zero count (LZCNT) instruction, 266
leading zero count instruction, 266
leaf procedure
 modifying windowed registers, 91
little-endian byte order, 8, 18, 70
load
 block, *See* block load instructions
 floating-point from alternate space instructions, **248**
 floating-point instructions, **246**, 251
 floating-point state register instructions, **246**, **264**
 from alternate space, 19, 51, 83
 instructions, 8
 instructions accessing memory, 79
 nonfaulting, 412
 short floating-point, *See* short floating-point load instructions
load short floating-point instructions, 253
LoadLoad MEMBAR relationship, 269
LoadLoad MEMBAR relationship, 419
LoadLoad predefined constant, 515
loads
 nonfaulting, 421
load-store alignment, 18, **79**, 409
load-store instructions
 compare and swap, **134**, **136**
 definition, 8
 load-store unsigned byte, 134, **255**, 359, 360
 load-store unsigned byte to alternate space, **256**
 memory access, 17
 swap R register with alternate space memory, **360**
 swap R register with memory, 134, **359**
LoadStore MEMBAR relationship, 269, 419
LoadStore predefined constant, 515
local registers, 33, 35, 315
logical XOR instructions, **390**
Lookaside predefined constant, 515
LSTPARTIALF instruction, 437
LZCNT instruction, 266

M

MAXPGL, 17, 32, 33, **74**, **75**, 75, 508
MAXPTL
 and MAXPGL, 75
 instances of TNPC register, 67
 instances of TPC register, 67
 instances of TSTATE register, 68
 instances of TT register, 69
may (keyword), 8
mem_address_not_aligned exception, 135, **463**
 generated by virtual processor, 250
 JMWPL instruction, 238
 LDTXA, 435, 436
 load instruction with partial store ASI and misaligned address, 250
 RETURN, 320, 321
 with CASA instruction, 135
 with load instructions, 79–80, 239, 240, 246, 248, 250, 251, 253, 254, 257, 258, 260, 261, 264, 356, 436, 437
 with store instructions, 79–80, 250, 334, 335, 336, 340, 342, 344, 345, 348, 350, 353, 355, 436, 437
 with swap instructions (deprecated), 359, 361
MEMBAR
 #Sync
 semantics, 271
instruction
 atomic operation ordering, 419
 FLUSH instruction, 171, 420
 functions, **269**, **418–420**
 memory ordering, 270
 memory synchronization, 85
 side-effect accesses, 409
 STBAR instruction, 270
mask encodings
 #LoadLoad, 269, 419
 #LoadStore, 269, 419
 #Lookaside, 420
 #Lookaside (deprecated), 270
 #MemIssue, 270, 420
 #StoreLoad, 269, 419
 #StoreStore, 269, 419
 #Sync, 270, 420
predefined constants
 #LoadLoad, 515
 #LoadStore, 515
 #Lookaside (deprecated), 515
 #MemIssue, 515
 #StoreLoad, 515
 #StoreStore, 515
 #Sync, 515
MEMBAR
 #Lookaside, 416
 #StoreLoad, 416
membar_mask, **515**
MemIssue predefined constant, 515
memory
 access instructions, 17, 79
 alignment, 409
 atomic operations, 418
 atomicity, 505
 cached, 407
 coherence, 409, 505
 coherency unit, 410
 data, 420
 instruction, 420
 location, 407
 models, **407**
 ordering unit, 410
 real, 408
 reference instructions, data flow order constraints, 413
 synchronization, 270
 virtual address, 407
 virtual address 0, 421
Memory Management Unit
 definition, 8
Memory Management Unit, *See* MMU
memory model
 mode control, 415
 partial store order (PSO), **415**

- relaxed memory order (RMO), 271, 415
- sequential consistency, 415
- strong, 415
- total store order (TSO), 271, 415, 416
- weak, 415
- memory model (mm) field of PSTATE register, 71
- memory order
 - pending transactions, 414
 - program order, 412
- memory version mismatch, 8
- memory_model (mm) field of PSTATE register, 415
- memory-mapped I/O, 408
- metrics
 - for architectural performance, 440
 - for implementation performance, 440
 - See also* performance monitoring hardware
- microcore, 8
- MMU
 - accessing registers, 482
 - contexts, 474
 - definition, 8
 - page sizes, 471
 - SPARC V9 compliance, 481
- mode
 - nonprivileged, 15
 - privileged, 16, 66, 411
- MONTMUL instruction, 272, 378
- MONTSQR instruction, 276, 381
- motion estimation, 299, 300
- MOVA instruction, 280
- MOVCC instruction, 280
- MOVcc instructions, 280
 - conditionally moving integer register contents, 51
 - conditions for copying integer register contents, 89
 - copying a register, 42
 - encoding of cond field, 491
 - encoding of opf_cc instruction field, 491
 - used to avoid branches, 183, 282
- MOVCS instruction, 280
- move conditionally by BSHUFFLE instruction, 140
- move floating-point register if condition is true, 180
- move floating-point register if contents of integer register satisfy condition, 184
- MOVE instruction, 280
- move integer register if condition is satisfied
 - instructions, 280
- move integer register if contents of integer register satisfies condition instructions, 283
- move on condition instructions, 14
- MOVFA instruction, 281
- MOVFE instruction, 281
- MOVFG instruction, 281
- MOVFGE instruction, 281
- MOVFL instruction, 281
- MOVFLE instruction, 281
- MOVFLG instruction, 281
- MOVFN instruction, 281
- MOVFNE instruction, 281
- MOVFO instruction, 281
- MOVFU instruction, 281
- MOVFUE instruction, 281
- MOVFUG instruction, 281
- MOVFUGE instruction, 281
- MOVFUL instruction, 281
- MOVFULE instruction, 281
- MOVG instruction, 280
- MOVGE instruction, 280
- MOVGU instruction, 280
- MOVL instruction, 280
- MOVLE instruction, 280
- MOVLEU instruction, 280
- MOVN instruction, 280
- movn synthetic instructions, 518
- MOVNE instruction, 280
- MOVNEG instruction, 280
- MOVPOS instruction, 280
- MOVr instructions, 90, 283, 491
- MOVrGEZ instruction, 283
- MOVrGZ instruction, 283
- MOVrLEZ instruction, 283
- MOVrLZ instruction, 283
- MOVrNZ instruction, 283
- MOVrZ instruction, 283
- MOVVC instruction, 280
- MOVVS instruction, 280
- MOVxTOd instruction, 284, 285
- MPMUL instruction, 286, 384
- multiple unsigned condition codes, emulating, 90
- multiplication, very large numbers, 372
- multiply instructions, 291, 333, 371
- multiply-add instructions (fused), 175
- multiply-add, *See* FMA instructions
- multiply-subtract instructions (fused), 175
- multiply-subtract, *See* FMA instructions
- multiprocessor synchronization instructions, 134, 359, 360
- multiprocessor system, 8, 171, 307, 359, 360, 414, 505
- MULX instruction, 291
- must (keyword), 8
- MWAIT instruction, 61, 292

N

- N superscript on instruction name, 96
- N_REG_WINDOWS, 9
 - integer unit registers, 17, 499
 - RESTORE instruction, 315
 - SAVE instruction, 322
 - value of, 32, 62
- NaN (not-a-number)
 - conversion to integer, 393
 - converting floating-point to integer, 231
 - signalling, 42, 162, 232
- neg synthetic instructions, 518
- negative add instructions (fused), 191, 195
- negative infinity, 393
- negative multiply-add instructions (fused), 175
- negative multiply-subtract instructions (fused), 175
- nested traps, 14
- next program counter register, *See* NPC register
- NFO, 8

- noncacheable
 - accesses, 408
- nonfaulting load, 8, 412
- nonfaulting loads
 - behavior, 421
 - use by optimizer, 421
- non-faulting-only page
 - illegal access to, 461
- non-faulting-only page, illegal access to
 - and TTE.nfo, 476
- nonleaf routine, 238
- nonprivileged, 8
 - mode, 5, 8, 15, 16, 44
 - software, 53
- nonresumable_error* exception, 463
- nonstandard floating-point, *See* floating-point status register (FSR) NS field
- nontranslating ASI, 9, 260, 355, 424
- nonvirtual memory, 308
- NOP instruction, 123, 157, 160, 294, 304, 365
- normal traps, 451
- NORMALW instruction, 295
- not synthetic instructions, 518
- note
 - architectural direction, 4
 - compatibility, 4
 - general, 3
 - implementation, 3
 - programming, 3
- NPC (next program counter) register, 53
 - and PSTATE.tct, 70
 - control flow alteration, 11
 - definition, 8
 - DONE instruction, 147
 - instruction execution, 77
 - relation to TNPC register, 67
 - RETURN instruction, 318
 - saving after trap, 21
- nucleus context, 173
- Nucleus Context register, 483
- nucleus software, 9
- NUMA, 9
- nvm (invalid mask) field of FSR.tem, 47, 393
- NWIN, *See* N_REG_WINDOWS
- nxm (inexact mask) field of FSR.tem, 47

O

- octlet, 9
- odd parity, 9
- ofm (overflow mask) field of FSR.tem, 47
- op3 instruction field, 134
 - arithmetic instructions, 110, 127, 128, 291, 325, 333, 369, 371
 - floating point load instructions, 246, 248, 251, 264
 - flush instructions, 171, 174
 - jump-and-link instruction, 238
 - load instructions, 239, 255, 256, 257, 259
 - logical operation instructions, 118, 296, 390
 - PREFETCH, 303

- RETURN, 320
- opcode
 - definition, 9
- opf instruction field
 - floating point arithmetic instructions, 153, 164, 190, 230
 - floating point compare instructions, 162
 - floating point conversion instructions, 231, 232, 235
 - floating point instructions, 152
 - floating point integer conversion, 168
 - floating point move instructions, 179
 - floating point negate instructions, 193
- opf_cc instruction field
 - floating point move instructions, 180
 - move instructions, 491
- opf_low instruction field, 180
- optional, 9
- OR instruction, 296
- ORcc instruction, 296
- ordering MEMBAR instructions, 85
- ordering unit, memory, 410
- ORN instruction, 296
- ORNcc instruction, 296
- OTHERW instruction, 297
- OTHERWIN (other windows) register, 65
 - FLUSHW instruction, 174
 - keeping consistent state, 66
 - modified by OTHERW instruction, 297
 - partitioned, 66
 - range of values, 62
 - rd designation for WRPR instruction, 376
 - rs1 designation for RDPR instruction, 313
 - SAVE instruction, 322
 - zeroed by INVALIDW instruction, 237
 - zeroed by NORMALW instruction, 295
- OTHERWIN register trap vectors
 - fill/spill traps, 465
 - handling spill/fill traps, 465
 - selecting spill/fill vectors, 465
- out register #7, 37
- out registers, 33, 35, 322
- overflow
 - bits
 - (v) in condition fields of CCR, 85
 - accrued (ofa) in aexc field of FSR register, 48
 - current (ofc) in cexc field of FSR register, 48
 - causing spill trap, 465
 - tagged add/subtract instructions, 85
 - overflow mask (ofm) field of FSR.tem, 47

P

- p (predict) instruction field of branch instructions, 126, 128, 160
- P superscript on instruction name, 96
- packed-to-planar conversion, 216
- packing instructions, *See* FPACK instructions
- page fault, 308
- page table entry (PTE), *See* translation table entry (TTE)
- parity, even, 6
- parity, odd, 9

- partial store instructions, 347, 437
- partial store order (PSO) memory model, 415, 415
- partitioned
 - addition with saturation, 204
 - additions, 201, 219
 - shift instructions, 228
 - subtraction with saturation, 222
 - subtracts, 219
- P_{ASI} superscript on instruction name, 96
- P_{ASR} superscript on instruction name, 96
- PAUSE instruction, 298
- PAUSE register, 60, 61
- PC (program counter) register, 9, 49, 53
 - after instruction execution, 77
 - and PSTATE.tct, 70
 - CALL instruction, 130
 - changed by NOP instruction, 294
 - copied by JMPL instruction, 238
 - saving after trap, 21
 - set by DONE instruction, 147
 - set by RETRY instruction, 318
 - Trap Program Counter register, 67
- P_{dis} superscript on instruction or register name, 96
- PDIST instruction, 299
- PDISTN instruction, 300
- pef field of PSTATE register
 - and access to GSR, 54
 - and *fp_disabled* exception, 462
 - and FPop instructions, 92
 - branch operations, 158, 160
 - byte permutation, 125
 - comparison operations, 163, 170, 209, 212
 - component distance, 284, 285, 299, 300
 - data formatting operations, 165, 197, 216
 - data movement operations, 282
 - enabling FPU, 53
 - FCHKSM16 instruction, 161
 - floating-point operations, 140, 152, 153, 164, 168, 176, 179, 182, 185, 190, 191, 193, 195, 230, 231, 232, 234, 235, 246, 248, 251, 253, 264
 - integer arithmetic operations, 186, 203, 206, 215, 218, 219, 221, 224, 229
 - integer SIMD operations, 178
 - logical operations, 225, 226, 227
 - memory operations, 244
 - read operations, 311, 330, 339
 - special addressing operations, 117, 154, 156, 340, 342, 345, 348, 350, 356, 374
 - trap control, 449
- pef, *See* PSTATE, pef field
- performance monitoring hardware
 - accuracy requirements, 440
 - classes of data reported, 440
 - counters and controls, 441
 - high-level requirements, 439
 - kinds of user needs, 439
 - See also* instruction sampling
- physical processor, 9
- PIL (processor interrupt level) register, 74
 - interrupt conditioning, 448
 - interrupt request level, 450
 - interrupt_level_n*, 463
 - specification of register to read, 313
 - specification of register to write, 376
 - trap processing control, 449
- pipeline, 9
- pipeline draining of CPU, 67
- pipeline draining of virtual processor, 63
- pixel instructions
 - compare, 207, 210
 - component distance, 299, 299
 - formatting, 197
- planar-to-packed conversion, 216
- POPC instruction, 301
- positive infinity, 393
- precise floating-point traps, 314
- precise trap, 446
 - conditions for, 446
 - software actions, 446
 - vs. disrupting trap, 447
- predefined constants
 - LoadLoad, 515
 - lookaside (deprecated), 515
 - MemIssue, 515
 - StoreLoad, 515
 - StoreStore, 515
 - Sync, 515
- predict bit, 128
- prefetch
 - for one read, 307
 - for one write, 307
 - for several reads, 307
 - for several writes, 307
 - page, 308
 - to nearest unified cache, 308
- prefetch data instruction, 303
- PREFETCH instruction, 79, 303, 502
 - prefetch_fcn*, 515
- PREFETCHA instruction, 303, 502
 - and invalid ASI or VA, 461
- prefetchable, 9
- Primary Context ID 0, 482
- Primary Context ID 1, 482
- Primary Context register, 483
- priority of traps, 449, 458
- privileged_action* exception
 - read from TICK register when access disabled, 52
- privilege violation, and *DAE_privilege_violation* exception, 461
- privileged, 9
 - mode, 16, 66
 - registers, 66
 - software, 15, 36, 44, 71, 84, 174, 451, 502
- privileged (priv) field of PSTATE register, 73, 135, 148, 240, 244, 248, 249, 256, 260, 335, 339, 343, 354, 360, 411, 463, 464
- privileged mode, 9
- privileged_action* exception, 135, 463
 - accessing restricted ASIs, 411
 - read from TICK register when access disabled, 52, 311
 - restricted ASI access attempt, 84, 423

- TICK register access attempt, 51
- with CASA instruction, 135
- with load alternate instructions, 240, 244, 249, 256, 260, 335, 339, 343, 354, 360
- with load instructions, 248
- with RDasr instructions, 312
- with read instructions, 312
- with store instructions, 344
- with swap instructions, 361
- privileged_opcode* exception, 52, 464
 - DONE instruction, 148
 - RETRY instruction, 319
 - SAVED instruction, 324
 - with DONE instruction, 148, 313, 318, 377
 - with write instructions, 377
 - WRPR in nonprivileged mode, 52
- processor, 9
 - execute unit, 412
 - issue unit, 412, 412
 - privilege-mode transition diagram, 445
 - reorder unit, 412
 - self-consistency, 412
- processor cluster, *See* processor module
- processor consistency, 414, 417
- processor interrupt level register, *See* PIL register
- processor self-consistency, 412, 416
- processor state register, *See* PSTATE register
- processor states
 - execute_state*, 458
- program counter register, *See* PC register
- program counters, saving, 443
- program order, 412, 412, 413
- programming note, 3
- PSO, *See* partial store order (PSO) memory model
- PSR register (SPARC V8), 374
- PSTATE register
 - entering privileged execution mode, 443
 - restored by RETRY instruction, 147, 318
 - saved after trap, 443
 - saving after trap, 21
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - and TSTATE register, 68
- PSTATE register fields
 - ag
 - unimplemented, 73
 - am
 - CALL instruction, 130
 - description, 71
 - masked/unmasked address, 147, 238, 318, 320
 - cle
 - and implicit ASIs, 83
 - and PSTATE.tle, 70
 - description, 70
 - ie
 - description, 73
 - enabling disrupting traps, 448
 - interrupt conditioning, 448
 - masking disrupting trap, 452
 - mm

- description, 71
- implementation dependencies, 71, 415, 504
- reserved values, 71
- pef
 - and FPRS.fef, 71
 - description, 71
 - See also* pef field of PSTATE register
- priv
 - access to register-window PR state registers, 66
 - accessing restricted ASIs, 411
 - description, 73
 - determining mode, 8, 9, 478
- tct
 - branch instructions, 124, 127, 128, 138, 158, 160
 - CALL instruction, 130
 - description, 70
 - DONE instruction, 148
 - JMPL instruction, 238
 - RETRY instruction, 318, 365
 - RETURN instruction, 320
- tle
 - description, 70
- PTE (page table entry), *See* translation table entry (TTE)

Q

- quadword, 9
 - alignment, 18, 79, 409
 - data format, 23
- quiet NaN (not-a-number), 42, 162

R

- R register, 10
 - #15, 37
 - special-purpose, 37
 - alignment, 257, 260
- rational quotient, 369
- R-A-W, *See* read-after-write memory hazard
- rcond instruction field
 - branch instructions, 128
 - encoding of, 491
 - move instructions, 283
- rd (rounding), 10
- rd instruction field, 134
 - arithmetic instructions, 110, 127, 128, 291, 325, 333, 369, 371
 - floating point arithmetic, 153
 - floating point arithmetic instructions, 164, 190, 230
 - floating point conversion instructions, 231, 232, 235
 - floating point integer conversion, 168
 - floating point load instructions, 246, 248, 251, 264
 - floating point move instructions, 179, 180
 - floating point negate instructions, 193
 - floating-point instructions, 152
 - jump-and-link instruction, 238
 - load instructions, 239, 255, 256, 257, 259
 - logical operation instructions, 118, 296, 390
 - move instructions, 281, 283
 - POPC, 301

RDASI instruction, 49, 51, 310

RDasr instruction, 310

- accessing I/O registers, 19
- implementation dependencies, 311, 501
- reading ASRs, 48

RDCCR instruction, 49, 50, 310

RDCFR instruction, 49, 310

RDFPRS instruction, 49, 53, 310

RDGSR instruction, 49, 54, 310

RDPC instruction, 49, 310

- reading PC register, 53

RDPR instruction, 49, 313

- accessing GL register, 75
- accessing non-register-window PR state registers, 66
- accessing register-window PR state registers, 62
- and register-window PR state registers, 62
- effect on TNPC register, 68
- effect on TPC register, 67
- effect on TSTATE register, 68
- effect on TT register, 69
- reading privileged registers, 66
- reading PSTATE register, 70
- reading the TICK register, 52
- registers read, 313

RDSOFTINT instruction, 49, 55, 310

RDSTICK instruction, 49, 57, 310, 311

RDSTICK_CMPR instruction, 49, 310

RTICK instruction, 49, 52, 310, 311

RDY instruction, 50, 311

read ancillary state register (RDasr) instructions, 310

read state register instructions, 20

read-after-write memory hazard, 413

real address, 10

real ASI, 424

real memory, 407, 408

reference MMU, 511

reg, 511

reg_or_imm, 515

reg_plus_imm, 514

regaddr, 515

register reference instructions, data flow order

- constraints, 413

register window, 33, 34

register window management instructions, 21

register windows

- clean, 64, 66, 90, 460, 465, 466
- fill, 36, 66, 90, 91, 315, 462, 465, 466
- management of, 15
- overlapping, 35–37
- spill, 36, 66, 90, 91, 322, 464, 465, 466

registers

- See also* individual register (common) names
- accessing MMU registers, 482
- address space identifier (ASI), 411
- ASI (address space identifier), 51
- CFR, 59
- CFR (compatibility feature register), 461
- chip-level multithreading, *See* CMT
- clean windows (CLEANWIN), 64
- clock-tick (TICK), 463
- current window pointer (CWP), 63
- F (floating point), 391, 450
- floating-point, 17
 - programming, 41
- floating-point registers state (FPRS), 53
- floating-point state (FSR), 42
- general status (GSR), 54
- global*, 14, 17, 32, 33, 33, 499
- global level (GL), 75
- IER (SPARC V8), 374
- in*, 33, 35, 322
- local*, 33, 35
- next program counter (NPC)
 - and PSTATE.tct, 70
- next program counter (NPC), 53
- other windows (OTHERWIN), 65
- out*, 33, 35, 322
- out* #7, 37
- PAUSE, 60, 61
- processor interrupt level (PIL)
 - and SOFTINT, 55
- processor interrupt level (PIL), 74
- program counter (PC)
 - and PSTATE.tct, 70
- program counter (PC), 53
- PSR (SPARC V8), 374
- R register #15, 37
- renaming mechanism, 413
- restorable windows (CANRESTORE), 64, 64
- savable windows (CANSAVE), 63
- scratchpad
 - privileged, 438
- SOFTINT, 49
- SOFTINT_CLR pseudo-register, 49, 56
- SOFTINT_SET pseudo-register, 49, 55
- STICK
 - and STICK_CMPR, 57
- STICK_CMPR
 - ASR summary, 49
 - int_dis field, 55, 58
 - stick_cmpr field, 58
 - and system software trapping, 57
- STICK_CMPR, 57
- TBR (SPARC V8), 374
- TICK, 52
- trap base address (TBA), 69
- trap base address, *See* registers: TBA
- trap level (TL), 73
- trap level, *See* registers: TL
- trap next program counter (TNPC), 67
- trap next program counter, *See* registers: TNPC
- trap program counter (TPC), 67
- trap program counter, *See* registers: TPC
- trap state (TSTATE), 68
- trap state, *See* registers: TSTATE
- trap type (TT), 69, 451
- trap type, *See* registers: TT
- VA_WATCHPOINT, 463, 464
- visible to software in privileged mode, 66–76
- WIM (SPARC V8), 374

- window state (WSTATE), 65
 - window state, *See* registers: WSTATE
- Y (32-bit multiply/divide), 50
- relaxed memory order (RMO) memory model, 271, 415
- renaming mechanism, register, 413
- reorder unit, 412
- reordering instruction, 412
- reserved, 10
 - fields in instructions, 108
 - register field, 32
- reset
 - reset trap, 447
- restartable deferred trap, 446
- restorable windows register, *See* CANRESTORE register
- RESTORE instruction, 35, 315–316
 - actions, 90
 - and current window, 37
 - decrementing CWP register, 35
 - fill trap, 462, 465
 - followed by SAVE instruction, 36
 - managing register windows, 21
 - operation, 315
 - performance trade-off, 315, 322
 - and restorable windows (CANRESTORE) register, 64
 - restoring register window, 315
 - role in register state partitioning, 66
- restore synthetic instruction, 517
- RESTORED instruction, 91, 317
 - creating inconsistent window state, 317
 - fill handler, 315
 - fill trap handler, 91, 466
 - register window management, 21
- restricted, 10
- restricted address space identifier, 84
- restricted ASI, 411, 423
- resumable_error* exception, 464
- ret/ret1 synthetic instructions, 517
- RETRY instruction, 318
 - and restartable deferred traps, 447
 - effect on TNPC register, 68
 - effect on TPC register, 67
 - effect on TSTATE register, 68
 - generating *illegal_instruction* exception, 463
 - modifying CCR.xcc, 51
 - reexecuting trapped instruction, 466
 - restoring gl value in GL, 76
 - return from trap, 443
 - returning to instruction after trap, 449
 - target address, return from privileged traps, 20
- RETURN instruction, 320–321
 - computing target address, 20
 - fill trap, 462
 - mem_address_not_aligned* exception, 463
 - operation, 320
 - reexecuting trapped instruction, 320
- RETURN vs. RESTORE instructions, 320
- RMO, 10
- RMO, *See* relaxed memory order (RMO) memory model
- rounding
 - for floating-point results, 43
 - in signed division, 325
- rounding direction (rd) field of FSR register, 153, 164, 190, 230, 231, 232, 234, 235
- routine, nonleaf, 238
- rs1 instruction field, 134
 - arithmetic instructions, 110, 127, 128, 291, 325, 333, 369, 371
 - branch instructions, 128
 - compare and branch instructions, 136
 - floating point arithmetic instructions, 153, 164, 190
 - floating point compare instructions, 162
 - floating point load instructions, 246, 248, 251, 264
 - flush memory instruction, 171
 - jump-and-link instruction, 238
 - load instructions, 239, 255, 256, 257, 259
 - logical operation instructions, 118, 296, 390
 - move instructions, 283
 - PREFETCH, 303
 - RETURN, 320
- rs2 instruction field, 134
 - arithmetic instructions, 110, 127, 128, 291, 296, 325, 333, 369, 371
 - compare and branch instructions, 136
 - floating point arithmetic instructions, 153, 164, 190, 230
 - floating point compare instructions, 162
 - floating point conversion instructions, 231, 232, 235
 - floating point instructions, 152
 - floating point integer conversion, 168
 - floating point load instructions, 246, 248, 251, 264
 - floating point move instructions, 179, 180
 - floating point negate instructions, 193
 - flush memory instruction, 171
 - jump-and-link instruction, 238
 - load instructions, 239, 257, 259
 - logical operation instructions, 118, 390
 - move instructions, 281, 283
 - POPC, 301
 - PREFETCH, 303
- RTO, 10
- RTS, 10

S

- savable windows register, *See* CANSAVE register
- SAVE instruction, 35, 322
 - actions, 90
 - after RESTORE instruction, 320
 - clean_window* exception, 460, 465
 - and current window, 37
 - decrementing CWP register, 35
 - effect on privileged state, 322
 - leaf procedure, 238
 - and *local/out* registers of register window, 36
 - managing register windows, 21
 - no clean window available, 64
 - number of usable windows, 64
 - operation, 322
 - performance trade-off, 322
 - role in register state partitioning, 66
 - and savable windows (CANSAVE) register, 63

- spill trap, 464, 465, 466
- save synthetic instruction, 517
- SAVED instruction, 91, **324**
 - creating inconsistent window state, 324
 - register window management, 21
 - spill handler, 323, 324
 - spill trap handler, 91, 466
- scaling of the coefficient, 186
- scratchpad registers
 - privileged, **438**
- SDIV instruction, 50, **325**
- SDIVcc instruction, 50, **325**
- SDIVX instruction, **291**
- Secondary Context ID 0, 482
- Secondary Context ID 1, 482
- Secondary Context register, 483
- self-consistency, processor, 412
- self-modifying code, 171, 172, 419
- sequencing MEMBAR instructions, 85
- sequential consistency, 409, 415
- sequential consistency memory model, **415**
- SETHI instruction, **85, 327**
 - creating 32-bit constant in R register, 19
 - and NOP instruction, 294
 - with rd = 0, 327
- setn synthetic instructions, 517
- SHA1 instruction, **328**
- SHA256 instruction, **328**
- SHA512 instruction, **328**
- shall (keyword), **10**
- shared memory, 407
- shift count encodings, 331
- shift instructions, 19
- shift instructions, 85, **331**
- short floating-point load and store instructions, 437
- short floating-point load instructions, 253
- short floating-point store instructions, 350
- should (keyword), **10**
- SIAM instruction, 330
- side effect
 - accesses, 408
 - definition, **10**
 - I/O locations, 408
 - instruction prefetching, 409
 - real memory storage, 408
 - visible, 408
- side-effect page, illegal access to, 462
- signalling NaN (not-a-number), 42, 232
- signed integer data type, 23
- signx synthetic instructions, 518
- SIMD, **11**
 - and CMASK instructions, 140
 - instruction data formats, 29–30
- simm10 instruction field
 - move instructions, 283
- simm11 instruction field
 - move instructions, 281
- simm13 instruction field
 - floating point
 - load instructions, 246, 264
 - simm13 instruction field
 - arithmetic instructions, 291, 296, 325, 333, 369, 371
 - floating point load instructions, 248, 251
 - flush memory instruction, 171
 - jump-and-link instruction, 238
 - load instructions, 239, 255, 256, 257, 259
 - logical operation instructions, 118, 390
 - POPC, 301
 - PREFETCH, 303
 - RETURN, 320
 - single instruction/multiple data, *See* SIMD
 - SLL instruction, 331
 - SLLX instruction, 331
 - SMUL instruction, 50, **333**
 - SMULcc instruction, 50, **333**
 - SOFTINT register, 49, **55**
 - clearing, 468
 - clearing of selected bits, 56
 - communication from nucleus code to kernel code, 468
 - scheduling interrupt vectors, 467, 468
 - setting, 468
 - SOFTINT register fields
 - int_level, 55
 - sm (stick_int), 55
 - SOFTINT_CLR pseudo-register, 49, **56**
 - SOFTINT_SET pseudo-register, 49, **55, 56**
 - software
 - nucleus, **9**
 - software translation table, 473
 - software trap, 365, 451
 - software trap number (SWTN), **365**
 - software, nonprivileged, 53
 - software_trap_number*, **516**
 - source operands, 161, 177
 - SPARC V8 compatibility
 - LD, LDUIW instructions, 239
 - operations to I/O locations, 409
 - read state register instructions, 311
 - STA instruction renamed, 336
 - STBAR instruction, 270
 - STD instruction, 258, 259, 353, 355
 - tagged subtract instructions, 368
 - UNIMP instruction renamed, 236
 - window_overflow* exception superseded, 462
 - write state register instructions, 374
 - SPARC V9
 - compliance, 9, 481
 - features, 13
 - SPARC V9 Application Binary Interface (ABI), 15
 - speculative load, **11**
 - spill register window, 464
 - FLUSH instruction, 91
 - overflow/underflow, 36
 - RESTORE instruction, 90
 - SAVE instruction, 66, 90, 322, 465
 - SAVED instruction, 91, 466
 - selection of, 465
 - trap handling, 466
 - trap vectors, 322, 466
 - window state, 66

- spill_n_normal* exception, 323, 464
 - and FLUSHW instruction, 174
- spill_n_other* exception, 323, 464
 - and FLUSHW instruction, 174
- SRA instruction, 331
- SRAX instruction, 331
- SRL instruction, 331
- SRLX instruction, 331
- stack frame, 322
- state registers (ASRs), 48–62
- STB instruction, 334
- STBA instruction, 335
- STBAR instruction, 374, 413, 419
- STBLOCKF instruction, 337, 436
- STDF instruction, 79, 340, 464
- STDF_mem_address_not_aligned* exception, 340, 342, 349, 464
 - and store instructions, 341, 344
 - STDF/STDFA instruction, 79
- STDFA instruction, 342
 - alignment, 79
 - ASIs for fp store operations, 437
 - causing *DAE_invalid_ASI* exception, 437
 - causing *mem_address_not_aligned* or *illegal_instruction* exception, 437
 - causing *STDF_mem_address_not_aligned* exception, 79, 464
 - for block load operations, 436
 - for partial store operations, 437
 - used with ASIs, 436
- STF instruction, 340
- STFA instruction, 342
- STFSR instruction, 42, 44, 463
- STFSR instruction (deprecated), 345
- STH instruction, 334
- STHA instruction, 335
- STICK register, 49, 52, 56
 - and **STICK_CMPR**, 57
 - RDSTICK instruction, 310
- STICK_CMPR register, 49, 57, 57
 - int_dis field, 55, 58
 - RDSTICK_CMPR instruction, 310
 - stick_cmpr field, 58
- store
 - block, *See* block store instructions
 - partial, *See* partial store instructions
 - short floating-point, *See* short floating-point store instructions
- store buffer
 - merging, 408
- store instructions, 11, 79
- store short floating-point instructions, 350
- StoreLoad MEMBAR relationship, 269, 419
- StoreLoad predefined constant, 515
- stores to alternate space, 19, 51, 83
- StoreStore MEMBAR relationship, 269, 419
- StoreStore predefined constant, 515
- STPARTIALF instruction, 347
- STPARTIALF instruction, *DAE_nc_page* exception, 461
- STQF instruction, 80, 340, 464
 - STQF_mem_address_not_aligned* exception, 340, 342, 464
 - STQF/STQFA instruction, 80
- STQFA instruction, 80, 342
- strand, 11
- strong consistency memory model, 415
- strong ordering, 415
- Strong Sequential Order, 416
- STSHORTF instruction, 350
- STTW instruction, 38, 79
- STTW instruction (deprecated), 352
- STTWA instruction, 38, 79
- STTWA instruction (deprecated), 354
- STW instruction, 334
- STWA instruction, 335
- STX instruction, 334
- STXA instruction, 335
 - accessing nontranslating ASIs, 355
 - referencing internal ASIs, 416
- STXFSR instruction, 42, 44, 356, 463
- SUB instruction, 357, 357
- SUBC instruction, 357, 357
- SUBcc instruction, 85, 357, 357
- SUBCcc instruction, 357, 357
- subnormal number, 11
- subtract instructions, 357
- SUBX instruction, 358
- superscalar, 11
- supervisor software
 - accessing special protected registers, 18
 - definition, 11
- SWAP instruction, 18, 359
 - accessing doubleword simultaneously with other instructions, 360
 - and *DAE_nc_page* exception, 461
 - hardware primitive for mutual exclusion, 418
 - identification of R register to be exchanged, 79
 - in multiprocessor system, 255, 256
 - memory accessing, 359
 - ordering by MEMBAR, 419
- swap R register
 - with alternate space memory instructions, 360
 - with memory instructions, 359
- SWAPA instruction, 360
 - accessing doubleword simultaneously with other instructions, 360
 - alternate space addressing, 18
 - and *DAE_nc_page* exception, 461
 - hardware primitive for mutual exclusion, 418
 - in multiprocessor system, 255, 256
 - ordering by MEMBAR, 419
- SWTN (software trap number), 365
- Sync predefined constant, 515
- synchronization, 271
- synchronization, 11
- synthetic instructions
 - mapping to SPARC V9 instructions, 516–518
 - for assembly language programmers, 516
 - mapping
 - bclrg, 518
 - bset, 518

- btog, 518
- btst, 518
- call, 516, 517
- casn, 518
- clrn, 518
- cmp, 516
- dec, 518
- deccc, 518
- inc, 518
- inccc, 518
- iprefetch, 517
- jmp, 516
- movn, 518
- neg, 518
- not, 518
- restore, 517
- ret/retl, 517
- save, 517
- setn, 517
- signx, 518
- tst, 517
- vs. pseudo ops, 516
- system software
 - accessing memory space by server program, 411
 - ASIs allowing access to memory space, 412
 - FLUSH instruction, 173, 421
 - processing exceptions, 411
 - trap types from which software must recover, 44
- System Tick Compare register, *See* STICK_CMPR register
- System Tick register, *See* STICK register

T

- TA instruction, 364, 491
- TADDcc instruction, 85, 362
- TADDccTV instruction, 85, 464
- tag overflow, 85
- tag_overflow* exception, 85, 362, 363, 367, 368
- tag_overflow* exception (deprecated), 464
- tagged arithmetic, 85
- tagged arithmetic instructions, 20
- tagged word data format, 23
- tagged words, 23
- TBA (trap base address) register, 69, 444
 - establishing table address, 21, 443
 - initialization, 450
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - trap behavior, 11
- TBR register (SPARC V8), 374
- TCC instruction, 364
- Tcc instructions, 364
 - at TL > 0, 451
 - causing trap, 443
 - causing trap to privileged trap handler, 451
 - CCR register bits, 51
 - generating *htrap_instruction* exception, 462
 - generating *illegal_instruction* exception, 463
 - generating *trap_instruction* exception, 464
 - opcode maps, 487, 491, 492

- programming uses, 365
- trap table space, 21
- vector through trap table, 443
- TCS instruction, 364, 491
- TE instruction, 364, 491
- termination deferred trap, 446
- test-and-set instruction, 418
- TG instruction, 364, 491
- TGE instruction, 364, 491
- TGU instruction, 364, 491
- thread, 11
- TICK register, 49
 - counter field, 52, 502
 - inaccuracies between two readings of, 502, 510
 - specification for RDPR instruction, 313
- timer registers, *See* TICK register *and* STICK register
- timing of instructions, 109
- tininess (floating-point), 48
- TL (trap level) register, 73, 444
 - affect on privilege level to which a trap is delivered, 450
 - and implicit ASIs, 83
 - displacement in trap table, 443
 - executing RESTORED instruction, 317
 - executing SAVED instruction, 324
 - indexing for WRPR instruction, 376
 - indexing privileged register after RDPR, 313
 - setting register value after WRPR, 376
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - and TBA register, 450
 - and TPC register, 67
 - and TSTATE register, 68
 - and TT register, 69
 - use in calculating privileged trap vector address, 450
 - and WSTATE register, 65
- TL instruction, 364, 491
- TLB
 - and 3-dimensional arrays, 122
 - miss
 - reloading TLB, 473
 - specialized miss handler code, 482
- TLE instruction, 364, 491
- TLEU instruction, 364, 491
- TN instruction, 364, 491
- TNE instruction, 364, 491
- TNEG instruction, 364, 491
- TNPC (trap next program counter) register, 67
 - saving NPC, 446
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
- TNPC (trap-saved next program counter) register, 11
- total order, 414
- total store order (TSO) memory model, 71, 271, 408, 415, 415, 416
- TPC (trap program counter) register, 11, 67
 - address of trapping instruction, 314
 - number of instances, 67
 - specification for RDPR instructions, 313
 - specification for WRPR instruction, 376
- TPOS instruction, 364, 491

- trailing zero count instruction (how to synthesize), 301
- translating ASI, 424
- Translation Storage Buffer (TSB), 475
- Translation Table Entry, *See* TTE
- trap
 - See also* exceptions and traps
 - noncacheable accesses, 409
 - when taken, 11
- trap enable mask (tem) field of FSR register, 449, 450, 500
- trap handler
 - privileged mode, 451
 - regular/nonfaulting loads, 8
 - returning from, 147, 318
 - user, 45, 393
- trap level register, *See* TL register
- trap next program counter register, *See* TNPC register
- Trap on Control Transfer
 - and instructions
 - Bicc, 124, 138
 - BPcc, 127
 - BPr, 128
 - CALL, 130
 - DONE, 148, 318, 365
 - FBfcc, 158, 160
 - JMPL, 238
 - tct field of PSTATE register, 70
- trap on integer condition codes instructions, 364
- trap program counter register, *See* TPC register
- trap state register, *See* TSTATE register
- trap type (TT) register, 451
- trap type register, *See* TT register
- trap_instruction* (ISA) exception, 365, 366, 464
- trap_little_endian (tle) field of PSTATE register, 70
- traps, 11
 - See also* exceptions and individual trap names
 - categories
 - deferred, 445, 446, 447
 - disrupting, 445, 447
 - precise, 445, 446, 447
 - priority, 449, 458
 - reset, 446, 447
 - restartable
 - implementation dependency, 447
 - restartable deferred, 446
 - termination deferred, 446
 - caused by undefined feature/behavior, 12
 - causes, 21, 21
 - definition, 21, 443
 - hardware, 451
 - hardware stack, 14
 - level specification, 73
 - model stipulations, 449
 - nested, 14
 - normal, 451
 - processing, 458
 - software, 365, 451
 - stack, 459
 - vector address, specifying, 69
- TSB, 11, 478
 - cacheability, 479
 - catching, 479
 - organization, 479
- TSB (Translation Storage Buffer), 475
- TSO, 11
- TSO, *See* total store order (TSO) memory model
- tst synthetic instruction, 517
- TSTATE (trap state) register, 68
 - DONE instruction, 147, 318
 - registers saved after trap, 21
 - restoring GL value, 76
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
- tstate, *See* trap state (TSTATE) register
- TSUBcc instruction, 85, 367
- TSUBccTV instruction, 85, 464
- TT (trap type) register, 69
 - and privileged trap vector address, 450
 - reserved values, 501
 - specification for RDPR instruction, 313
 - specification for WRPR instruction, 376
 - and Tcc instructions, 365
 - transferring trap control, 451
 - window spill/fill exceptions, 65
 - WRPR instruction, 376
- TTE, 11
 - context ID field, 476
 - cp (cacheability) field, 408
 - cp field, 461, 477, 477
 - e field, 408, 421, 462, 477
 - ie field, 477
 - nfo field, 421, 461, 476, 477
 - p field, 461, 478
 - soft2 field, 476
 - sz (size) field, 478
 - taddr field, 476
 - v field, 476
 - va_tag field, 476
 - w field, 478
- TTE in TSB, 475
- TVC instruction, 364, 491
- TVS instruction, 364, 491
- typewriter font, in assembly language syntax, 511
- TZCNT instruction (how to synthesize), 301

U

- UDIV instruction, 50, 369
- UDIVcc instruction, 50, 369
- UDIVX instruction, 291
- ufm (underflow mask) field of FSR.tem, 47
- UltraSPARC, previous ASIs
 - ASI_PHY_BYPASS_EC_WITH_EBIT_L, 438
 - ASI_PHYS_BYPASS_EC_WITH_EBIT, 438
 - ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE, 438
 - ASI_PHYS_USE_EC, 438
 - ASI_PHYS_USE_EC_L, 438
 - ASI_PHYS_USE_EC_LITTLE, 438
- UMUL instruction, 50
- UMUL instruction (deprecated), 371
- UMULcc instruction, 50

- UMULcc instruction (**deprecated**), 371
- UMULXHI instruction, 372
- unassigned, 11
- unconditional branches, 123, 127, 157, 160
- undefined, 12
- underflow
 - bits of FSR register
 - accrued (ufa) bit of aexc field, 48, 392
 - current (ufc) bit of cexc, 48
 - current (ufc) bit of cexc field, 392
 - mask (ufm) bit of FSR.tem, 48
 - mask (ufm) bit of tem field, 393
 - detection, 36
 - occurrence, 465
- underflow mask (ufm) field of FSR.tem, 47
- unfinished_FPop floating-point trap type, 45, 153, 164, 190, 230, 232, 234, 391
 - handling, 48
 - in normal computation, 44
 - results after recovery, 45
- UNIMP instruction (SPARC V8), 236
- unimplemented, 12
- unimplemented_FPop floating-point trap type, 392
 - handling, 48
- unimplemented_LDTW* exception, 258, 464
- unimplemented_STTW* exception, 353, 464
- uniprocessor system, 12
- unrestricted, 12
- unrestricted ASI, 423
- unsigned integer data type, 23
- user application program, 12
- user trap handler, 45, 393

V

- VA, 12
- VA_watchpoint* exception, 135, 464
- VA_WATCHPOINT register, 463, 464
- value clipping, *See* FPACK instructions
- value semantics of input/output (I/O) locations, 408
- virtual
 - address, 407
 - address 0, 421
- virtual address, 12
- virtual core, 12
- virtual memory, 308
- virtual processor, pipeline draining, 63
- VIS, 12
- VIS instructions
 - encoding, 493, 494, 495
 - implicitly referencing GSR register, 54
- Visual Instruction Set, *See* VIS instructions

W

- W-A-R, *See* write-after-read memory hazard
- watchpoint comparator, 72
- W-A-W, *See* write-after-write memory hazard
- WIM register (SPARC V8), 374
- window fill exception, *See also* *fill_n_normal* exception

- window fill trap handler, 21
- window overflow, 36, 465
- window spill exception, *See also* *spill_n_normal* exception
- window spill trap handler, 21
- window state register, *See* WSTATE register
- window underflow, 465
- window, clean, 322
- window_fill* exception, 65, 90
 - RETURN, 320
- window_spill* exception, 65
- word, 12
 - alignment, 18, 79, 409
 - data format, 23
- WRASI instruction, 49, 51, 373
- WRAsr instruction, 373
 - accessing I/O registers, 19
 - attempt to write to ASR 5 (PC), 53
 - cannot write to PC register, 53
 - implementation dependencies, 501
 - writing ASRs, 48
- WRCCR instruction, 49, 50, 51, 373
- WRFPRS instruction, 49, 53, 373
- WRGSR instruction, 49, 54, 373
- WRIER instruction (SPARC V8), 374
- write ancillary state register (WRAsr) instructions, 373
- write ancillary state register instructions, *See* WRAsr instruction
- write privileged register instruction, 376
- write-after-read memory hazard, 413
- write-after-write memory hazard, 413
- WRMWAIT instruction, 61
- WRPAUSE instruction, 373
- WRPR instruction
 - accessing non-register-window PR state registers, 66
 - accessing register-window PR state registers, 62
 - and register-window PR state registers, 62
 - effect on TNPC register, 68
 - effect on TPC register, 67
 - effect on TSTATE register, 68
 - effect on TT register, 69
 - writing the TICK register, 52
 - writing to GL register, 75
 - writing to PSTATE register, 70
 - writing to TICK register, 52
- WRPSR instruction (SPARC V8), 374
- WRSOFTINT instruction, 49, 55, 373
- WRSOFTINT_CLR instruction, 49, 55, 56, 373, 468
- WRSOFTINT_SET instruction, 49, 55, 373, 468
- WRSTICK_CMPR instruction, 49, 373
- WRTBR instruction (SPARC V8), 374
- WRWIM instruction (SPARC V8), 374
- WRY instruction, 49, 50, 373
- WSTATE (window state) register
 - description, 65
 - and fill/spill exceptions, 465
 - normal field, 465
 - other field, 465
 - overview, 62
 - reading with RDPR instruction, 313
 - spill exception, 174

spill trap, 322
writing with WRPR instruction, 376

X

XMULX instruction, 389
XMULXHI instruction, 389
XNOR instruction, 390
XNORcc instruction, 390
XOR instruction, 390
XORcc instruction, 390

Y

Y register, 49, 50
 content after divide operation, 325, 369
 divide operation, 325, 369
 unsigned multiply results, 333, 371
 WRY instruction, 374
Y register (deprecated), 50

Z

zero virtual address, 421

